

Analysis of End-to-end Delay Measurements in Internet

C. J. Bovy, H. T. Mertodimedjo, G. Hooghiemstra, H. Uijterwaal and P. Van Mieghem

Abstract— Parts of the RIPE NCC measurements of the end-to-end delay between a source and destination (on a specific day) have been analysed. A classification of the numerous histograms of the end-to-end delay of a fixed path demonstrates that about 84% are *typical* histograms possessing a Gamma-like shape with subexponential (or heavy) tail. Further, the deterministic component of the end-to-end delay is investigated and measured in a specific, unloaded environment. From the latter measurement, the delay caused by one router is estimated and used to produce the summarizing table VI below. The method of combining a travel-planner and the trace-routes to compute an estimate of the length of a path seems reasonable as no contradictions with the computed and measured end-to-end deterministic delay are encountered.

Keywords— Internet, end-to-end delay, measurements

I. INTRODUCTION

The interest in accurate end-to-end delay measurements is twofold. First and most important, deployment of real-time services (such as VoIP [8]) necessitates delay constraints to be met. The knowledge of the delay distribution along paths in the Internet allows us to verify whether these QoS requirements can be met and, alternatively, how to re-dimension paths or redesign network parts that currently fail to offer acceptable delay bounds. Second, from a more academic point of view, these end-to-end delay measurements may learn about the underlying properties of the current Internet, in particular, the topology and traffic pattern.

Here, results of an analysis of end-to-end delay measurements performed by RIPE NCC (the RIPE Test Traffic Measurements (TTM) Project) are presented. The RIPE measurement configuration is described in detail in [4]. About 40 testboxes are scattered over Europe (and a few in the US and New Zealand), measuring accurately (within 10 μ s) the one-way delay between each pair of boxes and the trace-routes. These measurements are performed regularly over time. The objective of this article is to understand the various different end-to-end delay histograms over fixed paths between a pair of boxes. First, a taxonomy of different end-to-end delay distributions (normalized histogram) is presented in section III. The main observation is that about 84% of these distributions are *typical*, possessing a Gamma-like shape and a heavy tail. A model and more detailed analysis of these typical end-to-end delay distributions is found elsewhere [6]. The remaining 16% consists of

a-typical or anomalous end-to-end distributions. A second part (section IV) is devoted to a study of the minimum delay. Before turning to both parts, we briefly overview and define the several components of the end-to-end delay in section II.

II. END-TO-END DELAY COMPONENTS.

In the RIPE measurement configuration, fixed size IP probe-packets of 100 bytes are sent from a source to a destination measurement box. The difference between the timestamps of departure at the source and arrival at the destination box are termed as the end-to-end delay of the IP probe-packet. A histogram contains about 2160 measurements per day per path.

Broadly, the components of the end-to-end delay D can be divided into four main categories:

- Processing delay D_p : The time needed to process a packet at each node and prepare it for (re)transmission. The processing delay is determined by the complexity of the protocol stack, the computational power available at each node and the link driver (or interface card). Processing delays are determined by the available hardware and are hardly affected by the traffic. However, the processing delay D_p is in general a stochastic random variable because it is not precisely the same for each probe-packet due to variability in tasks performed in the router. Therefore, we split the processing delay into a deterministic D_{pd} and a stochastic component D_{ps} , as $D_p = D_{ps} + D_{pd}$.
- Transmission delay D_t : The time needed to transmit an entire packet (or frame, bit train), from first bit to last bit, over a communication link. The transmission delay D_t is dictated primarily by the link speeds or capacity. For each probe-packet, the transmission delay can be regarded as being the same.
- Propagation delay D_{ew} : The time to propagate a bit through the communication link. The propagation delay D_{ew} is determined by the travel time of an electromagnetic wave through the physical channel of the communication path and is independent of actual traffic on the link. The propagation delay D_{ew} can be significant as, for instance, in satellite links or trans-Atlantic links.
- Queueing delay D_q : The waiting time of packets in the buffer of routers before transmission. The queueing delays depend on the details of the switching fabric in routers (or lower layer switches). The queueing delay is typically stochastic in nature due to interference of the probe-packet with other IP packets on (parts of) the path.

From a measurement point of view, the end-to-end delay over a fixed path is best separated into two components, a deterministic (or fixed) delay D_d and a stochastic de-

C.J. Bovy, H.T. Mertodimedjo, G. Hooghiemstra, P. Van Mieghem are with the Delft University of Technology, Information Technology and Systems, P.O. Box 5031, 2600 GA Delft, The Netherlands. Email: (C.J.Bovy, H.T.Mertodimedjo, G.Hooghiemstra, P.VanMieghem)@its.tudelft.nl. H. Uijterwaal is with RIPE NCC, Amsterdam, The Netherlands.

lay D_s . This separation is illustrated in Figure 1 where the histogram of a typical end-to-end delay on a link from Amsterdam to London is shown.

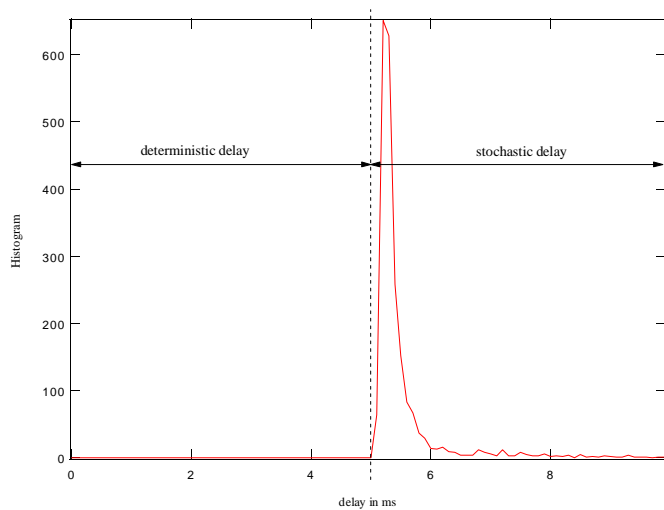


Fig. 1. A typical delay histogram with the separation into D_d and D_s

By combining the two different methods of separation into delay components, we have

$$\begin{aligned} D_d &= D_{pd} + D_t + D_{ew} \\ D_s &= D_{ps} + D_q \end{aligned}$$

The sequel $D_{ptd} = D_{pd} + D_t$ is treated as a single deterministic component because it is difficult to decompose a deterministic delay into the above three components, but relatively easy into two components D_{ptd} and D_{ew} . It is also possible to split D_{ptd} into two other components: the processing- and transmission delay at the end-points D_{tt} and the total processing- and transmission delay in the intermediate routers D_l , also called latency. The last split of $D_{ptd} = D_{tt} + D_l$ will be used in the next paragraphs. Section IV will further concentrate on D_d , while D_s is treated in [6].

III. TAXONOMY OF END-TO-END DELAY DISTRIBUTIONS.

We have separated the measured data from 5 March 2001 per fixed path (i.e. same path vector extracted from the trace-routes) for each test box pair (same source/target test boxes). For each test box pair, only the normalized delay distribution of the dominant path was considered. The dominant path is followed by most IP probe-packets during that day. In total 963 normalized delay distributions have been taken into account. The basic observation is that most of the distributions show a *typical* Gamma-like shape with a heavy tail. These typical distributions are studied elsewhere [6]. Those with a-typical behavior are often occurring in paths with either a same source or target test box. The underlying physics for these anomalous end-to-end delay distributions is still being investigated. Below we categorized those a-typical distributions into several classes.

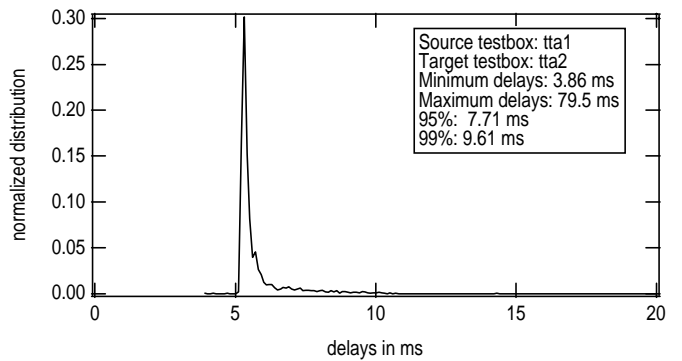


Fig. 2. A class A typical end-to-end distribution

Class A: Gamma-like shape with heavy tail.

For this typical class, the pdf (probability density function) delay decays slower than an exponential. There are indications that the tail follows a power law, $\Pr[D > x] \sim x^{-b}$. We found that 84% of total 963 normalized delay-distributions belong to this class. Figure 2 is an example of this distribution and clearly resembles a gamma-like shape. On double-log axes the tail of the (normalized) histogram or pdf can be approximated by a straight line, with slope lying between -1 and -3. The fitting of 30 paths of this class on a log-linear scale with a hyperbolic shape of $a|x-c|^{-b}$ leads to shape factors b around 0.6 - 1.0.

This heavy tail behavior seem to agree with early reports on self-similarity of WAN traffic (see [1] and for the influence of TCP on the self-similar structure [9]). Perhaps, not surprisingly because most WAN traffic [7] consist of most popular Internet applications like WWW and FTP.

Class B: Gamma-like with Gaussian or triangle lob

As exemplified in Figure 3, the pdf has a minimum, raises fast with a high exponential rate to reach the maximum and then decays. The decay is slower than exponential rate. At the end it raises again and decays with a slow slope. The second maximum is lower than the first one. The second bubble could be a Gaussian or triangle shape. We found that 6.3% of 963 normalized delay-distributions belong to this class B. Our first impression was that we had to face a new delay behavior. However, Figure 4 shows that the bubble in the tail is caused by packets sent between 9 AM and 9 PM. Trace route records don't show path switching around that time. However, the trace-routes are performed ten times per hour. If there occur (short) path variations in between two consecutive trace-route measurement, delay measurements of two different paths can get mixed, presumably leading to class B anomalies.

Class C: 2 gamma-like distributions

Figure 5 looks like 2 class A distributions shifted over some time. About 2.8% of 963 normalized delays distributions are of this class. From examination of the paths we suspect errors in the data. The traceroute record shows that the routing path is changing frequently resulting in a

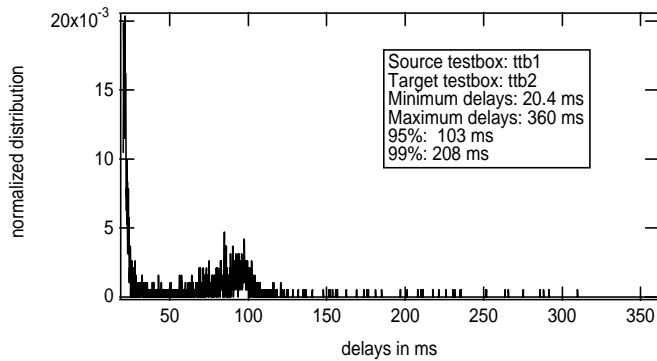


Fig. 3. An example of class B

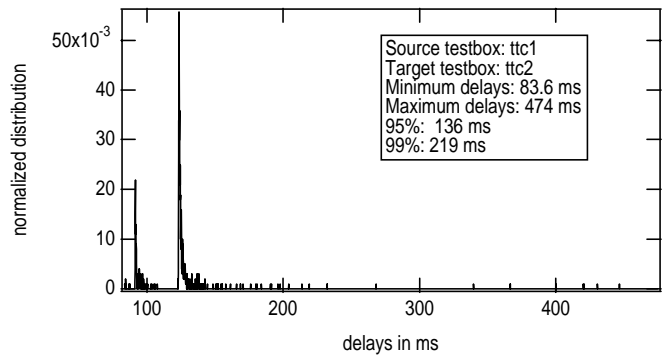


Fig. 5. An example of class C

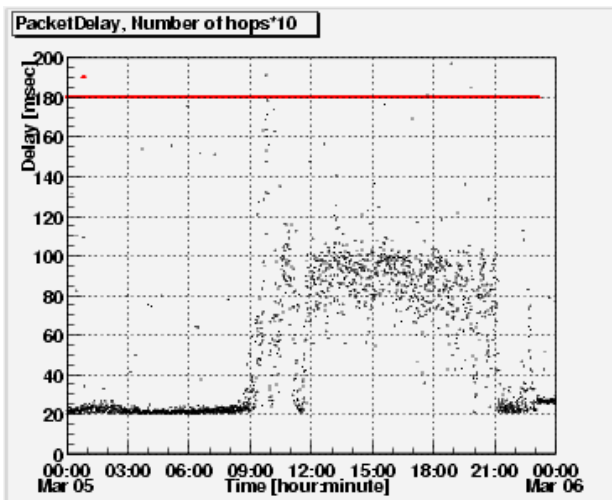


Fig. 4. Explanation of behavior of class B

non-stable switching from a path with 28 hops to a path with 22 hops in short time span. Because the path vector was updated at rate around 10 minutes while the IP probe-packets are sent at rate of 40 seconds, we suspect that some of the data are marked with the wrong trace-route path vector.

Class D: many peaks

An example of class D in Figure 6 shows that the location of peaks seems random, like white noise. In some case we can recognize a vague gamma like shape in it. About 5% of the 963 normalized delay distributions show this distribution. It occurs, for example, on path from box d1 to box d2 as observed in Figure 6. The variance of the delay seems very big. The losses on paths of class D as illustrated in Figure 7 may increase up to 40%. Although the precise reason for this behavior is still unknown, we suspect that it is caused by one specific router that may be overloaded (or under dimensioned).

Rest Class

The rest class of about 1.5% of all 963 cases contains the uncategorized distributions. Most of them only occurred once and do not fit in any of the classes described above.

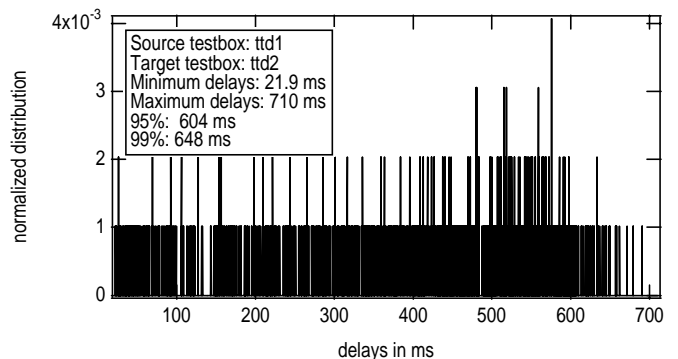


Fig. 6. Example of class D

This class may be attributed to spurious effects which do not deserve at this point more attention.

IV. DETERMINISTIC DELAY

In order to verify the separation of the end-to-end delay D into a deterministic D_d and stochastic D_s component, we have measured D_d independently. Besides the data from the RIPE TTM, the end-to-end delay between two similar RIPE measurement boxes (with same hardware and measurement software/procedure) has been measured under unloaded (or labo) conditions as shown in Figure 8 and detailed below.

Hence, instead of connecting two measurement boxes by the Internet, we have connected the two measurements boxes by a well controlled environment excluding largely the stochastic component D_s . Recall that the processing and transmission delay D_{ptd} can be split in D_{tt} and D_l where D_l in the latter measurements is the time to process a frame/packet at each node and prepare it for (re)transmission. D_{tt} is the total time to process a frame/packet at source and destination testbox. The processing time at the source and target node is augmented by the additional time between the timestamping of the probe-packet and the actual appearance on the link.

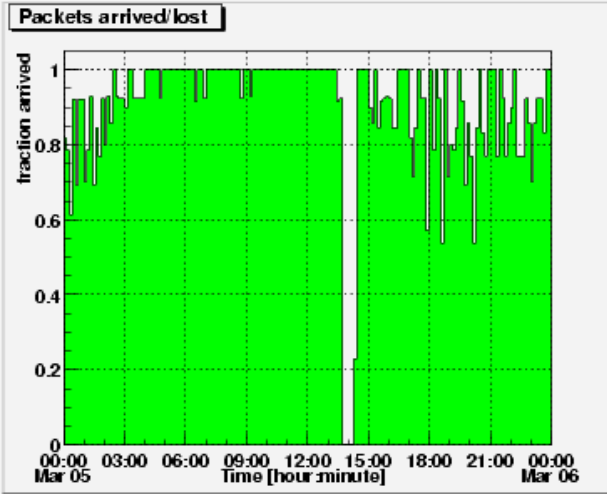


Fig. 7. Losses occurring in Class D

A. The deterministic processing and transmission delay D_{ptd} .

As stated in section II we split D_{ptd} into D_{tt} and D_l . To measure these parameters we have performed four tests. The measurement configuration is sketched in Figure 8. The first test is to measure the delay D_{tt} introduced by the end-points (testboxes). The other tests measure the latency D_l of the intermediate routers. In all these tests the propagation delay D_{ew} can be neglected because of the small distance between the two testboxes (half a meter).

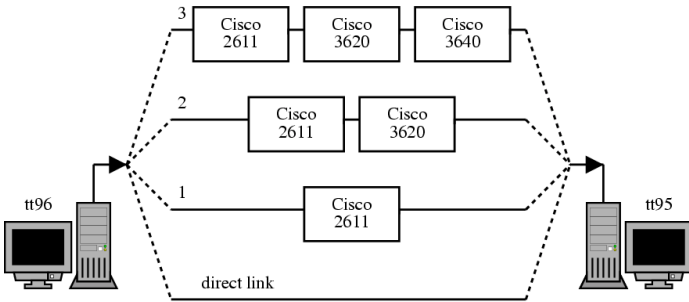


Fig. 8. The measurement configuration with routers between test boxes

A.1 Direct link D_{tt}

The first test measures the time needed by the testboxes itself, without any devices attached between the two testboxes, but connected via a short (half a meter) cross-link cable. The delay that would be measured can only be caused by the testboxes themselves and not by any router or propagation through the cable. The propagation-delay D_{ew} is negligible.

The time measured by the process can be split in three components:

1. Time between timestamping and the actual appearance on the link

2. Time needed by the packet to travel through the medium (in this case a UTP cross-link cable)
3. Time between the receipt at the destination-interface and the comparison of departure timestamp with the current GPS-time.

In Figure 8 the test-setup is shown while Figure 9 plots the result of the measurement.

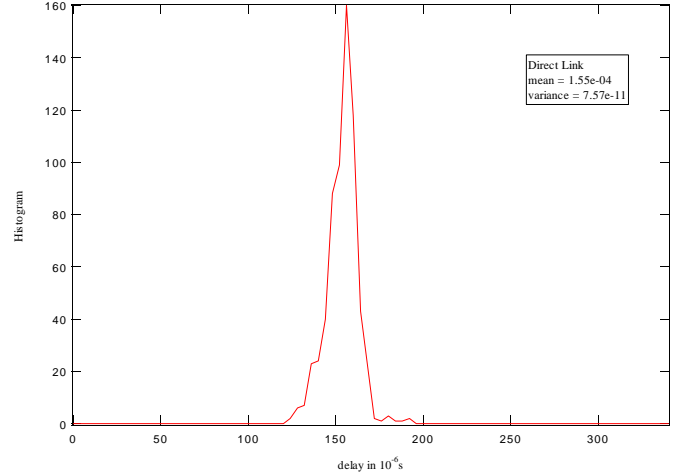


Fig. 9. Histogram of the delay (in s) measured between two test boxes connected by a direct link (cross cable)

A.2 Latency D_l

For the measurement of the latency of the intermediate routers, the same setup is used as shown in Figure 8. This measurement consists of actually three separated tests, each with a different number of intermediate routers (one, two and three routers, all different types). Note that these measurements return the end-to-end-delay of the setup and not the latency per router. The processing and transmission delay D_{tt} is included in the results. Out of the end-to-end delay we have to extract the per-router latency. For processing a probe packet at an intermediate node (router), the packet has to pass upwards through the lower OSI-layers to the IP layer 3 and the same way backwards through the layers. The histogram of the measurements is shown in Figure 10. The mean, variance and standard deviation $\sigma = \sqrt{\text{variance}}$ of all these tests are summarized in the table below

	Mean [s]	Variance [s ²]	σ [s]
Direct link	$1.55 \cdot 10^{-4}$	$7.57 \cdot 10^{-11}$	$8.87 \cdot 10^{-6}$
One router	$5.31 \cdot 10^{-4}$	$1.01 \cdot 10^{-10}$	$10.0 \cdot 10^{-6}$
Two routers	$8.36 \cdot 10^{-4}$	$1.86 \cdot 10^{-10}$	$13.6 \cdot 10^{-6}$
Three routers	$10.6 \cdot 10^{-4}$	$1.44 \cdot 10^{-10}$	$12.0 \cdot 10^{-6}$

We used a test-setup of three different routers (Cisco 2611, 3620 and 3630). Although these routers are used in small or medium office environments and are not used in a core IP-network, these measurements give us an insight in the order of the latency of routers.

Looking at the increase of the mean between the 'Direct link' and 'One router', we notice an increase of 376 μs .

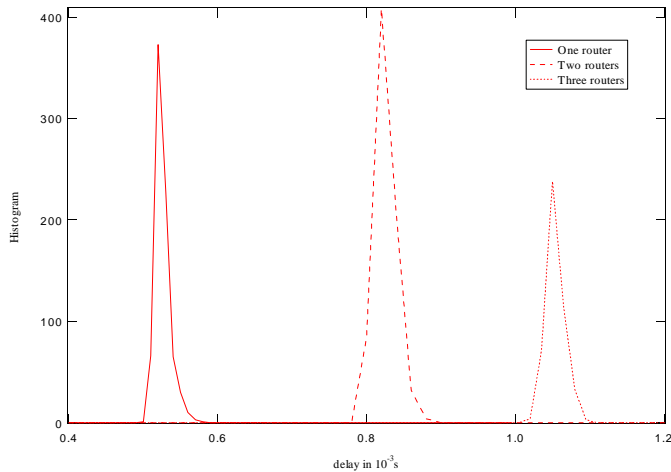


Fig. 10. The three histograms of the end-to-end delay between two test boxes with 1, 2 and 3 routers in between.

We denote this as the latency of the added router (Cisco 2611). The same is valid for the increase between 'One router' and 'Two routers' and respectively 'Two routers' and 'Three routers', an increase of $305 \mu\text{s}$ resp. $224 \mu\text{s}$. As expected, the last added router (Cisco 3630) has the smallest latency. This router is more advanced than the others. In further calculations we use a per-router latency of $224 \mu\text{s}$. Because we want a lower bound for the latency, we use the smallest per-router latency we measured. On a path with h hops (routers), we multiply the number of routers with the per-router latency and get $D_l = h \times 224 \mu\text{s}$. Lightreading, a company specialized in Network Testing [5] has published latency results with core routers¹ shown in Figure 11. Especially the Cisco router values agree well with our estimate of $224 \mu\text{s}$.

B. Propagation delay D_{ew}

The propagation delay D_{ew} is dependent on the physical length of the communication path. To determine the path-length of a link, we have to know the locations of the testboxes and the intermediate routers. The coordinates of the RIPE testboxes are known. From the coordinates of the RIPE testboxes, we can determine the line of sight (birds eye view) distance between the two nodes. Unfortunately, neither the location of the intermediate routers nor the exact length of the communication path between two nodes is known, which complicates the calculation of the propagation delay. Here, we propose to investigate the trace-routes and try to determine the location of the intermediate nodes. One way to determine the location of nodes or host is to check the DNS-record of the given node or host. This heavily depends on the owner of the node whether he/she updates the LOC-field in the DNS. The LOC-field stores the coordinates of the location of the node. We noticed that the LOC-fields are not widely used. Using this method to determine the location of the intermediate

¹Cisco 12416 Internet Router\12.0(14)SX, NetIron Internet Router\7.1.00 and M160 Internet Backbone Router\4.2R2.4 / 4.2E

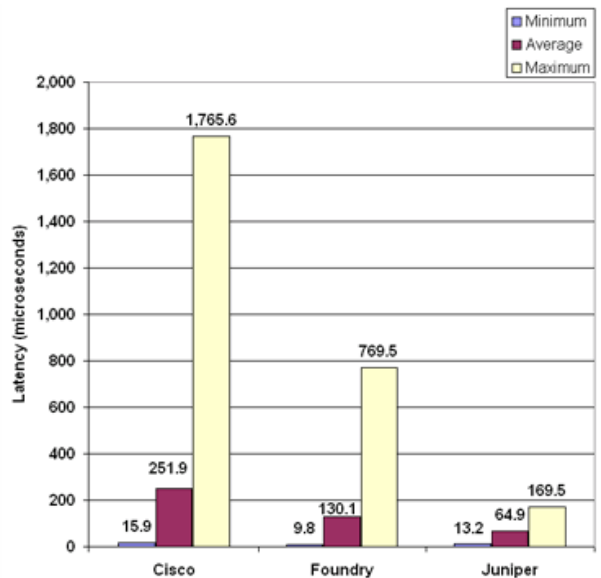


Fig. 11. Latency per router for Cisco, Foundry and Juniper

routers seems not effective. Another method is to investigate the naming of the routers/nodes. Some of the router hostnames refer to an abbreviation of a city that hints to the location. With a few of such nodes in a path, a better representation of the real physical path can be determined. To estimate the distance between those nodes, a travel-planner was used. Most of the communication links (fibers and copper wires) are buried along highways and railways. However, this method does not provide the real path length, only a rough estimate. As shown below, this rough estimate is a better representation of the real distance than the line of sight distance between the two endpoints.

For example, a traceroute from testbox A in Bratislava to testbox B in Munich, on a specific day (March 5, 2001) is specified² below as

Hop	IP Address	Hostname
1	x.x.23.1	BTS-Core-sw.y.y
2	x.x.243.133	skbra302-ta-f6-0-0.y.y
3	x.x.71.81	czpra103-tc-s4-1.y.y
4	x.x.70.45	debln302-tc-p2-0.y.y
5	x.x.70.37	debln301-tc-p5-0.y.y
6	x.x.71.122	Unknown address
7	x.x.70.14	bebru420-tc-p0-0.y.y
8	x.x.244.134	Unknown address
9	x.x.44.1	Cisco-M-XII-S1-0.y.y
10	x.x.0.116	Cisco-M-XI-Vlan1.y.y
11	x.x.1.121	cisco3000.y.y
12	x.x.44.25	ripe-testbox.y.y

The hostnames of the routers contain abbreviations of cities that can be recognized. For example in router *debln302-tc-p2-0.y.y* the abbreviation *de* stands for Deutschland and *bln* stands for Berlin. In this traceroute

²For confidential reasons, the precise IP numbers and domains cannot be given

a few other locations can be recognized. The locations of the other unknown nodes are ignored and we assume these nodes will not effect the path-length considerable. Figure 12 displays the geographical traceroute from ttA (Bratislava) to ttB (Munich). We observe that the distance is 5 times larger than the Line of Sight distance.

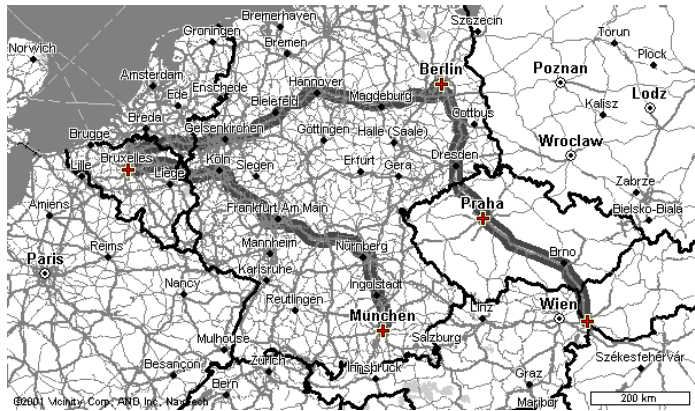


Fig. 12. Trace route from textbox A (Bratislava) to textbox B (Munich)

This method of calculating the distance is not possible for all links, because not all routers have logical names. Interestingly to mention is that some (European) links are routed via the US, for example from Testbox C (Stockholm) to Testbox D (Madrid). This introduces a considerable portion of propagation delay on that link.

For calculating the propagation-delay from the distance, we assume that all links are optical fibres with refraction index $n = \frac{c}{v} = 1.5$, where c is the speed of light in a vacuum ($3 \cdot 10^8$ m/s) and v is the speed of light in the medium. This results in a propagation delay of $5 \mu\text{s}/\text{km}$ in optical fibres (the same value is used in [8]). The resulting D_{ew} are also shown in Table VI at the end.

C. Summary

The deterministic (or minimum delay) here consists of

$$D_d = D_{ew} + D_l + D_{tt} \quad (1)$$

where $D_{tt} = 155 \mu\text{s}$ as determined from the direct link measurement.

In Table VI the results of the calculations are presented for 35 paths out of about 900 that can be mapped on a geographical map. Due to the precise GPS location of the measurement boxes, the line of sight L_{LoS} in the third column can be computed very accurately. Each delay-component is specified in a column. The eighth column contains the minimal delay measured by the RIPE TTM Project. We try to compare this column with our calculated D_d values. We expect, if our assumptions and calculations are correct, that the minimal delay measured by RIPE is equal or larger than our calculated D_d . The ratio of the third and fourth column and the second column is plotted in Figure 13 which illustrates that the paths are about 2.5 times (on average) longer than the shortest possible path along the

line of sight. Hence, from a network design point of view, an improvement of a factor of about 2.5 (on average) is possible by more appropriate routing or linking of routers. Figure 13 indicates that this improvement stems from a few badly routed paths and that the majority of paths (those with $L/L_{LoS} < 2$) is already quite economically routed. Also, there does not seem to exist a correlation between the hopcount and the length of a path.

The last column in Table VI presents a performance-factor $\eta = D_{min}/D_d$. This factor η provides an insight of the quality of the path. A large value of η means a larger measured minimal delay (provided by RIPE TTM) than calculated.

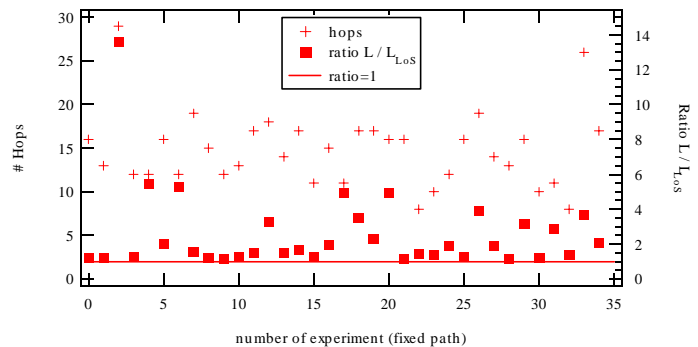


Fig. 13. The number of hops and the ratio L/L_{LoS} for each of the 35 fixed paths. The average (over the 35 paths) is $E[L/L_{LoS}] = 2.6$ whereas the standard deviation is $\sigma_{L/L_{LoS}} = 2.3$.

A performance-factor $\eta < 1$ implies that errors have been made in the method and that the measured minimal delay is smaller than the calculated minimum. All packets travelled over a specific path must have a delay larger than the calculated minimal delay for that specific path. All investigated paths meet this requirements (as can be verified from Table VI). The methods used to calculate D_d contain a few assumptions. The first assumption we made, is that the routers in a path are all of the same brand and type, and have a per-router latency of $224 \mu\text{s}$. If we use a smaller per-router latency, the performance-factor increases which means that no exceptions are introduced. On the other hand, if larger values are used, possible exceptions may occur. Hence, the estimate of the per-router latency of $224 \mu\text{s}$ seems a rather sharp overall design number. The second assumption consists of the route of the fibers or copper wires. We stated that they are buried along highways and used a travel-planner for calculating the distance. The exact route is not known, but this method gives us a rough estimate. A third assumption concerns the medium of the transportation link. We assumed that all links consist of optical fibres or that a refractive index n of 1.5 is appropriate. Furthermore, the location of the intermediate routers is not known exactly. The travel-planner returns routes, starting and ending at the center of a city. The last assumption concerns the extraction of the location out of the hostnames. The provider can for example install a router on a totally different location than mentioned in the host-

name of the router or a moved router's hostname has never changed. Finally, a few excessive performance-factors (bottom rows of Table VI) can be explained as unreliable paths, satellite connections or the use of slower routers. Satellite connections do not satisfy the method of calculation of the propagation delay.

V. CONCLUSIONS

We have analyzed the RIPE NCC measurements of the end-to-end delay between a source and destination. A classification of numerous histograms of the end-to-end delay of a fixed path demonstrates that about 84% are *typical* histograms possessing a Gamma-like shape with subexponential (of heavy) tail. Further, the deterministic component of the end-to-end delay is investigated and measured in a specific, unloaded environment. From the latter measurement, the delay caused by one router is estimated and used to produce the summarizing table VI below. The method of using the travel-planner seems reasonable as no contradictions with the computed and measurement end-to-end deterministic delay are encountered.

REFERENCES

- [1] M. E. Crovella and A. Bestavros, "Self-Similarity in World Wide Web Traffic: Evidence and Possible Causes", IEEE/ACM Transactions on Networking, Vol 5, No. 6, pp. 835-846, December 1997
- [2] RIPE NCC, <http://www.ripe.net>
- [3] RIPE Test Traffic Measurements, <http://www.ripe.net/ripenc/mem-services/ttm/>
- [4] Fotis Georgatos, Florian Gruber, Daniel Karrenberg, Mark Santcroos, Ana Susanj, Henk Uijterwaal and René Wilhelm, "Providing Active Measurements as a Regular Service for ISP's", Proceedings of Passive and Active Measurement (PAM2001), Amsterdam, The Netherlands, April 23-24, pp. 45-56, 2001. (also <http://www.ripe.net/ripenc/mem-services/ttm/Notes/PAM2001.ps.gz>)
- [5] Core Router test, <http://www.lightreading.com/testing>
- [6] G. Hooghiemstra and P. Van Mieghem, "Delay Distributions on fixed Internet Paths", Delft University of Technology, report20011031.
- [7] V. Paxson and S. Floyd, "Wide-Area Traffic: The Failure of Poisson Modeling", IEEE/ACM Transactions on Networking, Vol. 3 No. 3, pp. 226-244, June 1995.
- [8] A. Van Moffaert, D. De Vleeschauwer, J. Janssen, M.J.C. Büchli, G.H. Petit, "Tuning the VoIP Gateways to Transport International Voice Calls over a Best-Effort IP Backbone, Proceedings of the 9th IFIP Conference on Performance Modelling and Evaluation of ATM & IP Networks 2001 (IFIP01), pp. 193-205, Budapest (Hungary), 27-29- June 2001.
- [9] Andras Veres and Miklos Boda, "The chaotic nature of tcp congestion control", Proceedings of the IEEE Infocom, 2000, pp. 1715-1723

VI. TABLE CONTAINING THE DETERMINISTIC DELAY ANALYSIS OF 35 PATHS

	# Hops	Distance L_{LoS} in km (LoS)	Distance L in km (RoutePlan)	Propagation delay D_{ew} in ms (RoutePlan)	Processing delay D_l in ms routers	Total delay D_d $D_{ew} + D_l + D_{tt}$ in ms	Minimal delay D_{min} in ms. (RIPE)	Performance D_{min} / D_d
1	16	665.71	823	4.1	3.58	7.85	7.87	1.00
2	13	444.57	546	2.7	2.91	5.80	6.28	1.08
3	29	1190.44	16154	80.8	6.50	87.42	98.59	1.13
4	12	1301.41	1634	8.2	2.69	11.01	12.60	1.14
5	12	408.41	2218	11.1	2.69	13.93	16.00	1.15
6	16	1305.62	2604	13.0	3.58	16.76	19.32	1.15
7	12	408.42	2163	10.8	2.69	13.66	16.00	1.17
8	19	1301.41	2032	10.2	4.26	14.57	17.33	1.19
9	15	1122.04	1378	6.9	3.36	10.41	12.49	1.20
10	12	1305.76	1543	7.7	2.69	10.56	13.41	1.27
11	13	448.64	560	2.8	2.91	5.87	7.50	1.28
12	17	1624.04	2422	12.1	3.81	16.07	20.58	1.28
13	18	645.56	2097	10.5	4.03	14.67	18.92	1.29
14	14	985.29	1472	7.4	3.14	10.65	14.08	1.32
15	17	1120.51	1900	9.5	3.81	13.46	17.82	1.32
16	11	1656.37	2059	10.3	2.46	12.91	17.36	1.34
17	15	800.38	1576	7.9	3.36	11.40	15.32	1.34
18	11	2589.60	12708	63.5	2.46	66.16	89.44	1.35
19	17	469.38	1656	8.3	3.81	12.24	17.03	1.39
20	17	1620.27	3710	18.6	3.81	22.51	32.00	1.42
21	16	805.11	3951	19.8	3.58	23.49	34.32	1.46
22	16	1732.21	2047	10.2	3.58	13.97	20.95	1.50
23	8	687.45	1000	5.0	1.79	6.95	10.87	1.56
24	10	1242.25	1685	8.4	2.24	10.82	18.63	1.72
25	12	647.36	1240	6.2	2.69	9.04	15.95	1.76
26	16	659.79	820	4.1	3.58	7.84	13.93	1.78
27	19	445.32	1741	8.7	4.26	13.12	24.38	1.86
28	14	1304.81	2469	12.3	3.14	15.64	29.77	1.90
29	13	1495.48	1700	8.5	2.91	11.57	23.77	2.06
30	16	805.12	2523	12.6	3.58	16.35	34.32	2.10
31	10	356.52	437	2.2	2.24	4.58	9.73	2.12
32	11	861.62	2494	12.5	2.46	15.09	32.52	2.16
33	8	314.28	429	2.1	1.79	4.09	13.88	3.39
34	26	439.60	1614	8.1	5.82	14.05	190.86	13.59
35	17	450.88	944	4.7	3.81	8.68	189.43	21.82