Proc. 3rd International Conference on Management of Multimedia Network and Services (IFIP/IEEE MMNS), Fortaleza, Brazil, September 2000, Kluwer Academic Publishers, Boston, MA, 2000, pp. 47-59

Automating Internet Routing Behavior Analysis Using Public WWW Traceroute Services

Selim Gurun and Boleslaw K. Szymanski
Department of Computer Science, RPI, Troy, NY 12180-3590

Key words: Internet, route tracing, routing anomalies, path stability, path symmetry.

Abstract: Growing dependence of commerce, industry and academia on the Internet underlines importance of analyzing its behavior. Changing Internet structure makes the gathered data quickly obsolete, yet the results of analyses are useful in validating network management software and network protocols, as well as in planning extensions and enhancement. In this paper we describe tools for data analysis and an experiment conducted during the summer of 1999 using traceroute for analyzing Internet routing anomalies. We compare our results with Paxton's analysis of 1995's Internet.

1. INTRODUCTION

The Internet performance measurements are useful in validating network management software or protocols. They can also assist in planning new facilities and enhancing existing ones. Finally, such measurements can be used to gather statistics about the traffic, packets and routers. One of the very first studies of this kind investigated the dynamics of routing in the NSFNET backbone network [Ch93]. Data collected during a 12-hour period on August 1992 from all NSFNET nodes on the T1 backbone were used to analyze the frequency of routing table updates. Another study analyzed network inter-

domain routing information [LMJ97]. The authors collected BGP routing messages generated by border routers located at five public exchange points in the U.S. over the nine months. An important analysis of routing pathologies in the Internet was presented in [P97]. The author made end-to-end measurements using *traceroute* program between nearly 30 sites and used the results to study routing pathologies, path stability and routing symmetry in the Internet. The author also analyzed Internet packet dynamics by initiating and tracing TCP transmissions between the investigated sites. Data were collected during a few week periods in 1994 and 1995.

In this paper we describe tools for data analysis and an experiment conducted during the summer of 1999 during which *traceroute* measurements were collected and used to analyze the Internet routing anomalies. We closely followed methodology of analysis used in [P97] to easily compare routing problems of 1995's Internet with the problems existing today.

2. DESIGN AND IMPLEMENTATION OF MEASUREMENTS

2.1 *Introduction*

traceroute utility, available on many platforms, traces the hops on a route from the host to the destination using Time-To-Live (TTL) field of transmitted packets. It relies on the IP requirement that a packet with TTL field smaller than 2 is discarded by all routers except the destination. If a router discards a packet, it also sends an ICMP reply to the host that sent this packet. *traceroute* repeatedly sends packets to the traced destination with increasing values in TTL field, starting with 1. From the ICMP replies that it receives, *traceroute* reconstructs the route followed by the packets. The process stops when the packets sent in one step reach the destination and therefore do not trigger the ICMP reply. *traceroute* also calculates round trip time, and collects errors such as unreachable host, no replies, or unreachable network. *traceroute* output shows routing loops, and other routing pathologies, like fluttering paths. It is also possible to evaluate path symmetry or the average lifetime of paths using this information. The detailed description of *traceroute* utility is given in [JA89] and [P97].

Traditionally, tracing a route from a host machine to a destination requires telnet access to the host machine. Currently, some web sites have installed a CGI interface, which enables anybody to trace routes from these sites to any other host using WWW. A list of traceroute servers accessible via WWW is given in [T1] [T2] and [T3]. The price for using the existing

and easily accessible servers was limited control over parameters of the call. Most often, WWW traceroute servers allow only passing the name or IP address of a remote site. Some of the sites allow measurements only between themselves and the host that has directly connected to their web server making them useless to our analysis. Most of the sites use a script that restricts the measurement time to 45 seconds. The adverse effect of this restriction is discussed in the sections analyzing the results in detail. The results of measurements are displayed as HTML pages that are inconvenient for analysis, so we provided a parsing utility.

2.2 Basic Design of Measurements

We created a small model of the Internet, which is formed from hosts at many locations around the world. The selected hosts provide WWW traceroute service that we used to measure the paths to other selected hosts at random times for a long period of time. To make this model represent the Internet well, we selected the sites from geographically diverse and heavily wired locations. Thanks to this selection, even a small number of hosts can span a large number of paths, routers and autonomous systems because the number of paths grows with the square of the number of sites. We refer to the selected sites as *traceroute servers*. The site that initiates *traceroute* is referred to as a *source site* and the destination of this *traceroute* is called a *destination site*.

We wrote a few utilities to facilitate the measurements. *ndc* is a utility that contacts a traceroute server and tells it to measure a route to a destination site and then collects the results of measurements and save them for later analysis. *ndc* is installed on a single workstation in our computer laboratory, creating a centralized design with a single point of failure. In case of a network failure between our site and the world, *ndc* cannot contact traceroute servers. However, the same centralized design was used in [P97], making the results of both studies comparable. The collected results are preprocessed using *preproc* utility to remove measurements having errors and problems discussed later. They are analyzed using *trazer* utility

2.3 Data collecting (ndc) and analysis (trazer) utilities

ndc is a utility that we wrote to automate process of assigning measurements to traceroute servers. It selects a pair of servers, then contacts each of them to request measurements of routes between them. It also

receives, parses and saves the results returned by the servers. There is no single way to communicate with a WWW traceroute server. Some servers use ISINDEX field of HTML page, some others use forms and name fields in a different way. Some servers use POST while others use GET method to communicate with their CGI interface. The results are embedded into HTML pages with a design differing from site to site. These differences forced us to make *ndc* re-configurable for each single site.

The utility was designed to be immune to various network problems, because it must work continuously for quite long periods, a few weeks at a time. The main process creates child processes and tells them where they suppose to make measurements. A child makes a connection with the assigned traceroute servers, and when problems are detected, it cleans its leftovers, writes an error log and then exits.

ndc can make measurements in two different time scales, one with the large average inter-measurement time between two specific sites and the other with this time set to the much smaller value. The reason for having two different scales is to keep the size of data collected reasonable without jeopardizing the quality of analysis of results. A short interval between measurements is useful in analyzing route stability and duration of routing errors. However, the number of traced paths is proportional to the square of the number of participating sites, so a short inter-measurement time yields huge amounts of data and introduces significant load to the network.

After preprocessing is complete, the data is analyzed, using a utility named *trazer*, for path stability, routing pathologies and path symmetry. *traceroute* sends three packets at each hop. Sometimes these packets traverse different paths, causing different routers to send a reply. This is not an error or pathology but often a sign of load balancing. However, this also causes a problem during analysis of routing symmetry and stability. One option of *trazer* selects and removes the traceroutes with *ambiguous* paths. Another enables the user to replace router addresses with substitutes, which aids in locating routing errors.

3. ANALYZING RAW DATA AND ERRORS

3.1 Measurements having no hop information

We collected routing data during a two weeks from June 1 to June 14, 1999. We have used 29 sites, which are located in North American, European and Asian-Pacific countries. These sites and their locations are listed in [GS99]. All these sites, except `us4`, continued giving traceroute service until the end of our experiment. Starting on Jun 7-8 midnight, site

us4 canceled traceroute service. During the experiment, a total of 13172 traceroutes were attempted. Among these, 754 measurements failed by returning no hops data. The rest of measurements, a total of 12418, were collected and saved successfully. We analyzed 754 failures further using our measurements of path symmetry made for each traceroute.

23 traceroutes were lost because of errors in concurrent writing of the results by *ndc* children to a single file. Fortunately, only 0.17% of all traceroutes were affected by this problem.

104 traceroutes were lost because "connection timed out. Nearly a third of these failures were from three different time intervals and related with three different sites. In the longest of these periods, on June 3, from early morning to noon, a series of 23 connection attempts to us5 had ended in connection time out error. Surprisingly, during this period all the traceroutes conducted from the other sites to us5 have ended successfully! This makes us to believe that an Internet connectivity failure between our site and us5 lasting several hours had caused this problem. Another such period was on June 11, lasting for several hours and causing a series of 7 connection time out failures from tr1. Also the symmetric traceroutes to tr1 failed just on the border of metu.edu.tr domain during this period. However, except the failures discussed above and a series of 7 failures with us7 on June 13, all the other failures were distributed fairly evenly among the sites and persisted for much shorter time periods.

There were 6 failures due to "connection refused by the remote host" error. 45 connection attempts failed with "no route to host" error, out of which 23 were from ca1 and 20 were from de3. Our data shows that from June 6 to June 11, 20 connection attempts to ca1 returned this error. Our observation is that the sites/domains go into an unreachable period, which lasts for a few hours at a time and then recover. One interesting observation is that the connectivity problem generally appears again in less than 24 hours later. As an example, ca1 was unreachable from 21:45 to 23:55 on June 6, then it was unreachable from 13:00 to 18:00 on June 7, and this was repeated on June 8. We discuss this phenomenon again from different perspective in the routing error section.

Table 1. Distribution of Failures

Unsaved/lost	23	Concurrent writing in *ndc* utility
Connection timed out	110	Mostly 3 sites in specific intervals
No route to host	45	Mostly from ca1 and de3
Traceroute service stopped	343	Site stopped supporting *traceroute*
Unknown host	18	IP address of the destination not found
CGI script restrictions/other	215	45 second limitation in most scripts
Total	754	Measurements contained no data

As mentioned in the introduction part of this section, us4 has stopped *traceroute* service on June 7. Because of this stoppage a total of 254 traces had failed. There was also a temporary service outage with us2 that started on June 5 and ended on June 8. During this time, 89 traceroutes failed.

18 traceroutes failed with an error of “unknown host”. This was most probably caused by the DNS service problems.

The remaining 215 failures were mostly caused by the 45 second time limit in the CGI script.

3.2 Interrupted Measurements

The limit on the measurement time causes the trace to stop. Such stopped measurements can be identified by the name of the last hop being different from destination. This group of errors contains also traces of routes that are more than 30 hops long, because this is the length limit of *traceroute* tracing capabilities. Hop data of interrupted measurements is useless in analysis of path symmetry and lifetime, but can be used in analyzing routing loops.

Only 175 traces or 1.3% of the total number, suffered from this problem (see table 3). The effect of the time limit is small, except of site tr1. From a total of 454 traceroutes attempted from this site, 152 failed or returned no data and 74 were interrupted because of timing limit, which for this site was 345 seconds, much higher than for the most of other sites. This indicates a serious congestion and delay at this site.

Some hosts are located behind firewalls, which drop ‘undefined packages’. These hosts are seen as unreachable to *traceroute* utility. For example, among the 422 measurements directed to se1, 417 ended in *ITcenter-gw.DMZ.Umea.SE* and the remaining 5 were interrupted because of the time limit. Since the circular references after firewall node do not indicate routing loops we excluded them from the routing loop analysis.

Table 2. Distribution of Failures (Interrupted Measurement) by site

Site	*#*	*% of total*	*Site*	*#*	*%of total*	*Site*	*#*	*%of total*
us4	256 (1)	33.95	us7	9 (4)	1.19	au1	3 (2)	0.40
tr1	152(74)	20.16	hk1	9 (3)	1.19	at1	3 (8)	0.40
us2	91 (3)	12.07	se1	7 (5)	0.93	us9	2 (2)	0.27
de1	67 (1)	8.89	su1	5 (3)	0.66	us3	2 (3)	0.27
ca1	31 (2)	4.11	sp1	4(10)	0.53	us1	2 (2)	0.27
de3	25 (2)	3.32	nz1	4 (1)	0.53	ch1	2 (2)	0.27
us5	24 (3)	3.18	us12	3 (5)	0.40	us8	1 (5)	0.13
it1	11 (6)	1.46	jp1	3 (8)	0.40	us13	1 (1)	0.13
us6	10 (7)	1.33	cz1	3 (2)	0.40	us10	1 (3)	0.13

Table 3. Trace statistics (* denotes exclusion of 23 unsaved attempts, ta=traces attempted, f=traces failed, i=traces interrupted and %e=percentage of incomplete traces)

Site	ta*	f*	i	%e	Site	ta*	f*	i	%e	Site	ta*	f*	i	%e
at1	466	3	8	2.4	jp1	437	3	8	2.5	us5	477	24	3	5.7
au1	448	3	2	1.1	nz1	496	4	1	1.0	us6	442	10	7	3.8
ca1	402	31	2	8.2	se1	460	7	5	2.6	us7	481	9	4	2.7
ch1	538	2	2	0.7	sp1	459	4	10	3.1	us8	485	1	5	1.2
cz1	422	3	2	1.2	su1	419	5	3	1.9	us9	499	2	2	0.8
de1	459	67	1	14.8	tr1	454	152	74	49.8	us10	412	1	3	1.0
de2	423	0	5	1.2	us1	448	2	2	0.9	us11	468	0	2	0.4
de3	378	25	2	7.1	us2	404	91	3	23.3	us12	445	3	5	1.8
hk1	468	9	3	2.6	us3	481	2	3	1.0	us13	433	1	1	0.5
it1	446	11	6	3.8	us4	499	256	1	51.5	total	13149	731	175	6.9

ROUTING PATHOLOGIES

3.3 Permanent routing loops

Assume that we are tracing a path from a site A to a site B and the routers in the path are named R_1 to R_n. We assume that there is a **persistent** routing loop if the reported path is $R_1, R_2, \ldots R_i, R_{i+1}, \ldots R_k, R_i, R_{i+1}, \ldots R_k, \ldots$ and the trace ends before reaching site B because of the 30 hops limit. In other words, a **persistent** routing loop is a loop that does not resolve during a *traceroute* time. A loop in *traceroute* does not necessarily mean that there is a real forwarding loop. In [P97], a loop is classified as a real forwarding loop if the cycle is repeated three times. We have used the same rule in our analysis.

There were 12 persistent routing loops. In all twelve cases, a routing loop in a traceroute from site A to B corresponded to failures of measurements from site B to A with error messages, like “no route to host” and “connection timed out”. Most of the loops were directed to site ca1. This site became unreachable between June 3 and 11. Three of the permanent loops affected only one measurement. The measurements just before and after these have no loops. One of the loops, from su1 to jp1, has lived for two traces measured 10 minutes apart. None of the loops spanned two autonomous systems.

A temporary routing loop is the one that is resolved before *traceroute* ends. It may cause congestion by producing duplicate packets in the system [P97]. We found 11 temporary routing loops in our test data, 7 of them were either started from or directed to tr1.

3.4 Fluttering Paths

Fluttering refers to rapidly varying routing [P97]. More than one router name for a hop indicates that the path has changed during the trace time, which usually lasts tens of seconds. Fluttering may be a side effect of load balancing. However, it might introduce problems described in [P97] and [V97].

In our study we found that fluttering is very frequent in the Internet. Approximately 25% (3011 traceroutes) of all the traces contains at least one hop with two different routers. However, the distribution of fluttering paths is not equal. Nearly half (1501 traceroutes) of the fluttering paths either have jp1 or se1 as source or destination site. We also found that most of the affected paths have only one or two hops involved in fluttering, a case which we call a *minor fluttering*. To understand the distribution of fluttering, we assigned weights to fluttering paths. If a hop had more than one router reported, we assigned 1 to it, otherwise it was assigned the weight 0. The weight of a traceroute is the sum of weights of all hops. The distribution below shows that only a small percentage of paths suffer from *heavy fluttering.*

Table 4. Distribution of fluttering paths

Weight	=1	=2	=3	=4	>4	Total
%fluttering	2208 73%	607 20%	119 3.9%	44 1.5%	33 1%	3011 100%
%all traces	17%	4.8%	0.95%	0.35%	0.26%	24%

The negative effects of fluttering increase if the paths are significantly different from each other. Hence, we studied fluttering in Autonomous System (AS) granularity. This means that we assigned the weight 0 to a fluttering hop if the fluttering routers were in the same AS. Replacing router numbers with AS numbers in traceroute weight computation yielded the results shown in Table 5.

Table.5. Distribution of fluttering paths in AS granularity

Weight	=1	=2	=3	=4	>4	Total
%fluttering	448 78%	95 16%	25 4.3%	2 0.3%	3 0.4%	573 100%
%all traces	3.6%	0.76%	0.2%	0.01%	0.02%	4.6%

3.5 Unreachable messages

173 traceroutes returned message "*access administratively prohibited*". All of these traceroutes were directed to site se1. After removing failed and interrupted measurements, 61 traceroutes returned either "host unreachable" or "network unreachable" message with the following distribution among sites: ca1 24, de3 19, su1 4, us7 3 and others 11.

4. ROUTING STABILITY

Following [P97], we analyzed routing pathology using two basic notions: *prevalence* and *persistence* of routing. The first one defines the probability that a route *r* observed at time *t* is observed again at time t+s. In other words prevalence defines how likely a route is to be used again. The persistence defines the average lifetime of a path and it is not easy to measure based on information collected over small time periods. We evaluated stability in two different granularities to distinguish between minor and major path changes. Using bare IP names, we conducted *host granularity* analysis and then replacing IP names with AS numbers we made *AS granularity* analysis. A path rapidly switching between autonomous systems indicates much more serious problems then a path switching within the same AS.

Before analyzing path stability we removed incomplete traceroutes missing one or more hop data (eliminating 660 traceroutes) and ambiguous traceroutes with more than one router for any hop (removing another 2862 traceroutes). That left 8896 traceroutes for host granularity analysis. In AS granularity the number of ambiguous traceroutes decreases, leaving 11226 traceroutes for the corresponding analysis.

4.1 Prevalence

Our study includes measurements of 812 potential paths, i.e., *source-destination pairs*. After eliminating irrelevant and ambiguous traceroutes we had data on 727 paths. For each path, we calculated the ratio of number of times the dominant route was used to the total number of time the path was traversed. For example, assume that the path from A to B was traversed 9 times along the following routes:

R_1, R_2, R_2, R_1, R_3, R_1, R_1, R_1 R_1.

Then the prevalent route is R_1 and the prevalence of R_1 is 6/9 or 0.66.

The result of our analysis shows that 421 of 727 (58%) paths used the prevalent route at least 50% of time.

When we repeat the analysis at the AS granularity, results change dramatically. From 787 analyzed paths, only 3% (30 paths) had used the prevalent route less than 50% of time.

4.2 Persistence

We analyzed persistence in five different time scales: 60 sec., 10 min., 1 hour, 10 hours and more than 10 hours. For each category, we selected routes that changed in this time, and then calculated the ratio of varying paths to stable ones. We also studied which source/destination pairs are likely to switch frequently and which pairs are more long-lived. We again used two different granularities, host and AS granularity, as we did for prevalence. We used the same data that we used for prevalence. The results, shown in Table 6, are cumulative, i.e., 10-min. scale also covers the data for 60-sec. scale. The number of route changes counts number of changes from one router to another in a route. About a third of route changes at any scale are inter-AS route changes.

At 60 sec. scale all changes involved (at1,us11) and (au1,us5) pairs with low use of prevalent paths. Out of 60 changes for 10-min. scale, 44 routes started at at1 or de2 and 20 ended at us11. Nearly half of the route changes at 1 hour scale started at five sites (at1 14%, ch1 11%, us5 10%, de2 8% and au1 8% with a total of 54% of all changed routes), that were involved in a few routes, observed. The same sites were responsible for 23% of all route changes in AS granularity.

Table 6. Persistence in Host (AS) granularity

Scale	#of Routes	#of Route Changes
<=60 sec	47 (68)	2 (0)
<=10 min	405 (560)	60 (9)
<=1 hour	1900 (2551)	475 (117)
<=10 hours	5717 (8180)	1879 (525)
10 hours+	8897 (12013)	3311 (1023)

5. ROUTING SYMMETRY

Following [P97], we define two routes between a pair of sites as symmetric if they have the same number of hops, n, and $\forall i, 1 \le i \le n : r_i = r_{n-i}$, where r_i denotes a router.

We studied routing symmetry only in AS granularity, using only unambiguous traceroutes (i.e., those without errors or missing hops, with single router traversed by each hop and uninterrupted). This left us with 6624 traceroutes or 3312 traceroute pairs.

Table 7 with the results of analysis shows that only 36% of all measurements were symmetric in AS granularity. If we include the paths having only one hop difference, this ratio increases to 42.8%. In addition,

47% of pairs have different number of hops and are not included in this count. Another conclusion is that generally each pair maintains a fairly constant level of asymmetry. For example, we have 15 observations with asymmetry level equal to 7 and all of them included de1 either as source or as destination. To support the above conclusion, for each level of asymmetry we computed percentage of asymmetry associated with 50% of pairs that have most asymmetric paths. This value is about 80% at all levels of asymmetry, indicating that path asymmetry is quite unevenly distributed among hosts. Hence, paths are stable and asymmetric and each pair produces pretty constant level of asymmetry.

Table 7. Classification of pairs in terms of symmetry of their paths

pairs with different #of hops	1561	47%
asymmetric pairs with equal #of hops	557	17%
symmetric pairs	1194	36%
total number of pairs	3312	100%

Table 8. Details of pairs having equal number of hops (w=weight of asymmetry, N=number of measurement, and P=percentage of total measurements)

w	N	P	Comments
w=0	1194	36	Symmetric pairs
w=1	424	12.8	
w=2	83	2.5	
w=3	24	0.72	w>=3 is less than %1
w=4	3	0.09	
w=5	8	0.24	
w=6	0	0	
w=7	15	0.45	site de1 in all w=7
total	1751	52.86	

Table 9. Details of asymmetric pairs

difference between #of hops	#of pairs	% of all pairs
h=1	566	17
h=2	560	16
h=3	250	7
h>3	185	5.5

6. SUMMARY OF RESULTS

During a two weeks period, we collected 13172 traceroutes from 29 sites. These 29 sites introduced 812 possible paths. After eliminating the erroneous

754 traceroutes, we were left with 12418 measurements to analyze routing pathology, stability and symmetry.

We found 12 persistent and 11 temporary routing loops. However, our results may be well underestimating the real number of errors, because we have lost 390 measurements (215 with no data and 175 measurement interrupted in the middle) as a result of a time limit that restricts the *traceroute* measurement time to 45 seconds. Unfortunately, in case of an Internet connectivity problem, a *traceroute* is likely to take more time, which means some of interesting routing problems are lost with these 390 measurements.

We studied fluttering and path stability in two different granularities. First, we made our analysis based on each host visited and then on each autonomous system visited. We called the former *host granularity* and the latter *AS granularity*. Generally, a problem affecting more than one AS is much more serious then a problem inside a single AS.

We found that the ratio of fluttering paths in host granularity is 24% and it drops in AS granularity to 4.6%. However, not all the paths suffer evenly from fluttering. Some sites are much more likely to introduce varying paths. Nearly half of the fluttering paths had jp1 or se1 either as a source or as a destination.

We studied path stability in two different aspects, prevalence and persistence. The most frequently used route is called a prevalent route and the probability of using prevalent path is referred to as prevalence. The persistence measures the average lifetime of paths. The results of prevalence analysis in host granularity show that for half of the paths, the probability of using the prevalent path is 58%. In AS granularity this ratio increases to 97%.

Calculating the average lifetime of paths seems a much more difficult problem. We separated the data into different scales and then studied each group. The ratio of route changes to number of routes is an increasing function. We found that in 60 seconds scale there were 2 route changes among the 47 routes observed. In one hour scale, there were 475 route changes among 1900 routes and in 10 hour scale, there were 1879 route changes among 5700 routes.

We studied routing symmetry only in AS granularity level, because in host granularity almost all routes are highly asymmetric. We found 36% of all routes are symmetric. 12% of routes differ only in one AS and 5% of them differ in more than one AS. For 47% of traceroute pairs each traceroute had different number of hops so we did not calculate the asymmetry weight for them. We also noticed an interesting asymmetric path between two sites located in Germany and Italy in which packets from Germany go to USA before reaching Italy!

Our motivation for this study was to evaluate the Internet in terms of routing behavior and to compare our results to the previous studies [P97].

7. ACKNOWLEDGEMENT

This work was supported in part by DARPA grant F19628-98-C-0057. The content of this entry does not necessarily reflect the position or policy of the U.S. Government---no official endorsement should be inferred or implied.

8. REFERENCES

[CH93] B. Chinoy, "Dynamics of Internet Routing Information," Proceedings of SIGCOMM'93, pp. 45-52, September,1993.

[GS99] S. Gurun, B. Szymanski "Internet Routing Behavior Analysis Using Traceroute," Technical Report, Dept. Computer Science, Rensselaer Polytechnic Institute, Troy, NY, TR99-11, http://www.cs.rpi.edu/~szymansk

[JA89] V. Jacobson, traceroute, ftp://ftp.ee.lbl.gov, 1989.

[LMJ97] C. Labowitz, G.Malan, F. Jahanian, " Internet Routing Instability," Proceedings of SIGCOMM '97, September 1997.

[P97] Vern Paxson, "Measurements and Analysis of End-to-End Routing Internet Dynamics," Ph.D. dissertation, University of California, Berkeley, April 1997.

[T1] http://www.traceroute.org

[T2] http://www.netaxs.com/~rmsolino/tech/tracert.htm

[T3] http://www.geocities.com/Athens/4273/traceroute.html

[V97] Vern Paxson, "End-to-end routing behavior in Internet," IEEE/ACM transactions on Networking 5, page 601-615.