

Measurements of the Hopcount in Internet

F. Begtašević and P. Van Mieghem¹

Delft University of Technology
Information Technology and Systems
P.O. Box 5031, 2600 GA Delft, The Netherlands

submitted to PAM2001, 6 Nov. 2000

Abstract

End-to-end quality of service is very likely to depend on the hopcount, the number of traversed routers. The hopcount also enhances our understanding of properties of the Internet topology. In this paper we present the results of the measurements of the hopcount from a source at Delft towards several destinations spread over three continents. Fitting the data with our theoretical model illustrates that the number of routers based on our measurements is around 98400. The dynamics in the change of the hopcount is investigated, it was found that though there is virtually no short-term change in the average hopcount the individual hopcounts do change significantly. When investigating the changes in the paths to the sites within these continents, rapid changes in the paths were observed to barely affect the hopcount.

1 Introduction

In order to better understand the current behavior of the Internet, the properties of some quantities such as hopcount between two arbitrary nodes, the corresponding end-to-end QoS (delay, jitter and packet loss) may enhance our insight. In this paper, we confine ourselves to the hopcount, apparently the simplest of the above mentioned quantities.

Results of the measurements of the distribution of the hopcount are important to unravel the current topology and to dimension and propose a more efficient network infrastructure than the current Internet. These results are also beneficial to simulate more realistic network topologies. By combining the delay estimates and results of the hopcount measurements, we would be able to specify guaranteeing the Quality of Service in current Internet.

In recent years there were few attempts to get a better picture of the Internet topology. End-to-end performance studies were rare. Most of the topology related measurements were based on the hop-limited probes with the aim to create a router level map of Internet. Performance related measurements mainly involved measurements of RTT's between pairs of hosts in order to investigate how the RTT changes depending on the path, time of day, amount of traffic and other parameters.

In 1995 Paxson [5] analyzed 40,000 end-to-end path measurements, made by repeating the traceroute utility among 37 Internet sites. His goal was to study the routing behavior including pathologies, stability and path symmetry. His data showed that the paths for the most of the sites (almost two-thirds) stayed the same over longer period of time. About half of the paths were asymmetric.

¹ Corresponding author, P.VanMieghem@its.tudelft.nl

Vanhastel *et al* [7] published their results of the measurements performed at the University of Gent in 1998, but the number of destination sites they used is rather small. Moreover, they did not study the changes in hopcounts and paths.

PingER is a large end-to-end performance monitoring project [4], organized by the researchers from different institutes working on High Energy Physics experiments (HEPnet). PingER performs active measurements using ICMP echo requests sent at low frequency to approximately 500 sites. It collects information on packet loss and RTT.

Other projects on end-to-end performance measurements are Surveyor and NLANR's (National Laboratory for Applied Network Research) Active Measurement Program (AMP). These programs are more community-specific, which means that number of sites participating in these measurements is limited and that the measurements are done over same set of sites.

The subject of this paper are the measurements of hopcount for different regions, especially the distribution of the hopcount for the designated regions and possible differences between regions were investigated. Furthermore we studied the change of the average hopcount and the changes of the paths to the sites. These measurements were done using the traceroute program. Although this tool has already been used in most of the above mentioned studies, our approach to the measurements complements these results. The changes of the average hopcount over different periods of time were examined and the dynamics of these changes were investigated. Also the changes of the paths to the sites were studied and their influence on the results for the hopcount measurements were studied.

First, we present the set up of the measurement configuration and give an explanation of the procedure of gathering the addresses of the destination sites. In section three we present the results of the measurements of the short-term changes of the hopcount, hopcount per continent and the probability distribution function (pdf) of the hopcount (per continent and for the total sample). That section also contains the study of the changes in the hopcount and the paths during these measurements. Finally there is a short discussion of our results.

2. Measurement set-up and destination sites

2.1. Set-up

For the measurements of the hopcount we used the standard UNIX tool traceroute [6] on *dutetva* (130.161.40.13) SUN-workstation at the Delft University of Technology. To reach the public Internet from *dutetva* packets always make the same three hops as illustrated in Figure 1.

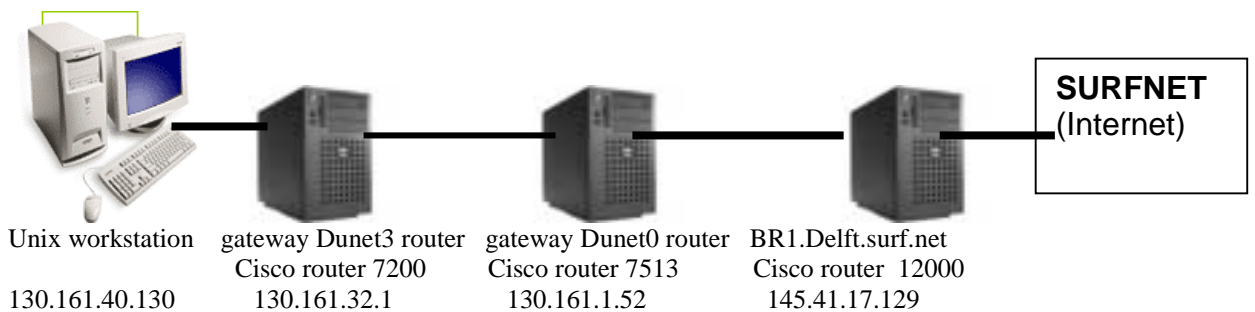


Figure 1: The measurement set-up

Two of these machines belong to the network of the University of Technology (speed of their connection is 622 Mb/s), while the third (though physically in Delft) belongs to the SURF-Net, the network of the Dutch Universities (the TU-network is connected to the SURFNET via 155 Mb/s link).



Figure 2 Capitals of the countries covered in the European set of addresses

2.2. Destination sites

We decided to concentrate on performing the hopcount measurements for three continents – Europe, Northern America and Asia-Pacific (Far East, Australia and New Zealand). This choice was made on the importance of these regions from the point of view of a European user. The most time-consuming part of these tests was finding enough suitable destination addresses. While collecting addresses of suitable sites, preference to large companies, Internet Service Providers and universities was given. If possible, universities were chosen because they provided a better geographical distribution of the sites and more responsive collaboration.

While collecting addresses for our sets, we encountered a number of problems.

The *first* problem was that many of the sites turned out to be unsuitable for the measurements due to the fact that many of the routers are programmed to discard ICMP messages used by the traceroute program.

The *second* problem was the geographical distribution of the sites. We discovered that a certain number of the sites belonging to the European and Asian institutions (BBC for example) were actually situated in other parts of the world (mostly in the United States). Hence they must be excluded from the measurements. This was the case with a limited number of sites. Despite the fact that some of the hostnames referred to a certain country, some sites were actually situated in another (e.g. sonera.be a name referring to Belgium is actually situated in Finland and many Chinese sites are placed in Hong Kong or Singapore). The problem of geographical distribution was also solved after consulting the databases of the local Internet registries. After checking every address from our sets in one of the databases, approximate geographical locations of sites were established. RIPE (Réseaux IP

Européens) is one of the three regional Internet registries that exist today. It provides allocation of Internet resources in Europe (and surrounding countries). Its database contains a list of more than ten million hosts in this area together with their IP numbers and addresses.




Figure 3 Average hopcount European set of addresses measured between 10th and 12th of April 2000

Other regional registries are APNIC (Asia Pacific Network Information Center) and ARIN (American Registry for Internet Numbers) and their databases were used during the creation of the sets of addresses for these regions. For European regions we tried a certain number of hosts (up to twenty) per country in order to get a better uniform spread of the average hopcount in Europe. The choice of the countries was made with respect to size and economical power. We mention that the set of European addresses does not contain any site in Netherlands. A similar approach was made while gathering addresses for the Pacific region, important countries (Australia, New Zealand, Japan, China) were represented with more sites than other countries. While collecting the addresses for North American we tried to include in our set at least four sites from every state in the US and Canada with (naturally more important states like California and New York were represented with more sites than states like Delaware or Idaho).

The *third* problem was caused by the fact that one hostname can refer to the different machines and that after some time the original hostname could be assigned to a different machine, which could influence our hopcount measurements. We avoided this problem by using IP-numbers instead of hostnames while using *traceroute* program.

The *final* problem we encountered after completion of our samples. After some time some of the sites (though initially there were no problems while running *traceroute* for them) needed to be discarded from the sample set because they did not respond anymore to ICMP message triggered by *traceroute*. The most probable cause for this problem were changes in (or installment of) the firewalls for these machines so that they also started to discard the ICMP messages. This problem was solved by finding other sites to replace them in order to keep the size of the sample large enough. In order to maintain a spread of the sites as uniform as possible, the replacement sites were in the same area (state or country) so that geographical

distribution wouldn't be significantly altered. Fortunately the number of sites prone to these changes was not large (two per set of addresses).

The sample for the European region consisted of 210 addresses distributed over 11 European countries. Figure 2 shows the cities in each country represented with the most sites in our sample. The Asian sample contained 186 addresses distributed over the area. The resulting set of American sites contained 274 addresses and was significantly larger than the other two samples.




Figure 4 Average hopcount and the variance for the European sample (22nd - 30th of October 2000)

3. Average Hopcount

In order to investigate the behavior of the hopcount in the current Internet we studied the short term changes in average hopcount (measurements repeated at certain intervals). The second point of interest was to observe the dynamics of the change of the hopcount for individual sites. As third point, we investigated the changes of the paths for individual sites from our sets.

3.1. Short Term Changes

To study short term changes of the average of the hopcount a short UNIX script was written which ran traceroute program consecutively for every site from the list and calculated the average hopcount (saving also the time of the start of the measurement) for that sample of sites. The total running time was two days and during this time the measurement was repeated 35 times.

We expected that some of the routers would at the moments of the increased traffic choose to transfer packets via different routes in cases of congestion and that this would increase the value of hopcount because load balancing would result in slightly longer paths.

However, as verified from Figure 3, there is virtually no increase in the number of hops in the times of the (expected) increased traffic on the network. The results from this initial test suggest that there is no significant relation between the time of the measurement and measured hopcount and that, in general, increase in the traffic due to busy hours on the Internet doesn't lead to an increase of the hopcount. Still there

was a certain increase of the average hopcount as the test progressed. In order to further investigate this phenomenon a similar test was repeated over a longer time interval.




Figure 5 Average hopcount and its variance from the three designated continents and for the total set of addresses

The second time measurement has been done from 3rd until 10th of August 00. That second measurement (Figure 4) leads to results consistent with the initial one. Again no increase of the average hopcount is observed, not even during the busy hours. We have repeated this procedure for the other two sets of addresses and the results were almost the same. This leads us to the conclusion that routing does not change due to the increase of the load to the network. This could be caused either by routers who use static routing tables or by the fact that the updates of the routing tables do not affect the value of the hopcount on average.

Still there is a significant change in variance for those measurements, which does not affect the measured mean nor depends on the time of the day at which the measurement was performed. In order to study this change further we followed how the hopcounts for individual sites change (section 3.4).

3.2. The hopcount over the continents

As expected the results for the American and Asian set of sites are slightly higher than for European set. This is caused by the geographical proximity of the sites in Europe, although some of the routes to European sites ran through United States. For the same reason results for the pacific region are highest (almost all routes to that region run through U.S.). Results from these measurements (August 2000) were used to calculate the mean and variance for the total sample, resulting in the average hopcount of 16,51 (with the variance of 14,96). The mean and the variance of the total sample were calculated as a weighted mean as described in formula (1).

$$E_{tot}[h] = \frac{\sum_{i=1}^3 E_i[h] \cdot S_i}{\sum_{i=1}^3 S_i} \quad (1)$$

Here $E_i[h]$ are averages for the sets for each continent and S_i represent the number of sites in the set i .

Indeed, the sets of addresses for three regions are not of the same size (set of addresses from Asia Pacific region is smallest) which could have influenced the result if the average of the hopcount of the different continents was taken equally important. On the other hand most of the traffic from Europe does go to sites in Europe and United States and this could be good estimate of the average hopcount in the Internet today.




Figure 6 Distribution of the hopcount for observed sets (measured in August 2000)

3.3. The pdf of the hopcount

The distributions of the hopcount for those regions are illustrated in Figure 5 and the total sample of the world in Figure 6. Results show reasonable agreement with the results published by Vanhastel *et al* (1998) According to these results the average hopcount in Internet in past two years decreased by approximately 1 hop.

We also compared our results with recent works of Van Mieghem *et al.* [8] in which the probability that the hopcount of the shortest path in the Internet consists of k hops was given as:

$$\Pr[h_N = k] = \frac{(1 + o(1))}{N} \sum_{k=1}^{N-1} c_{m+1} \frac{\ln^{k-m} N}{(k-m)!} \quad (2)$$

where c_k are the Taylor coefficients of $\frac{1}{\Gamma(x)}$ listed in [1], and N is numbers of nodes

in the graph. As we mentioned before every time the traceroute was run from our test machine the first three hops were always the same, also we established that in the average the two of the hops at the other end of the route were belonging to the same network as the destination site. Taking this in account in order to estimate the number of the routers in the Internet we decreased the results for the hopcount for 5 hops. The

distribution calculated from the equation (1) that is the closest to the actual results was for the $N=98400$ (Figure 7). This number represents the number of nodes in the graph and in our case the number of routers.




Figure 7 Distribution of the hopcount for the total set of addresses fitted with the limit law (2)

3.4.Changes in hopcount

In order to study the dynamics of the change in hopcount we recorded hopcounts for every site from the set and for every run of the procedure (162 times). We analyzed data for the North American set of addresses. Also for this set of addresses the procedure was run over the period of one week. The average hopcount and its variance did not change significantly during this period of time (Figure 8) nevertheless the individual hopcounts for these sites change for a large number of sites. When we compared all these files pair-wise we discovered that there is almost a steady rate of the change between consecutive runs. That change is presented in the Figure 9, where percentages of sites for which hopcounts changed between the referent measurement (runs 1, 88 and 156) and remaining runs. Difference between first and the last run in this trial is 25% which means that the hopcount changed in one week for 25 % of the sites from North American set of addresses. Still this change did not influence the average or variance of the hopcount.

This inspired us to study the correlation between the values of the hopcount for the individual sites in our sets. We used the files with hopcounts from the mentioned measurements, treated each measurement as vectors and then calculated the correlation coefficients for its elements.

Furthermore, the mean and variance of the hopcount of individual sites was computed. As expected from previous results, there are a significant number of sites for which the hopcounts have remained the same for all the measurements. For North American set of addresses 75 sites (or 27%) have constant hopcount for all of the 162 measurements. Further, the variance of the hopcount for most of the sites remains

limited, for 196 sites variation is less than .5, which indicates that rapid changes in the hopcount are exceptional. Also the hopcount for these sites has one value that prevails and that could be called ‘characteristic hopcount’. Only on the short-term basis, changes occur more often. Since we already expected the value for variance to be small, results for mean and variance alone were not that interesting but they were useful for the calculations of the correlation coefficients.




Figure 8 Average and the variance of the hopcount for the North American region (25 Aug - 2Sep2000)

The correlation matrix of 274x274 is too large to be shown in here. For large number of sites correlation coefficients were 1. Several groups of sites had a correlation coefficient of 1 with every other site in the group. This means that the hopcount for those sites changed at the same time (for the same measurement). Routes to some of these sites did share some of the links at the other end of the connection though they weren't in the physical vicinity of each other (they are evenly distributed over both of the American coasts) and they did not belong to the same sub-networks. Hopcounts to these sites were affected by the changes in the routing that occurred at the one of the SURF-NET routers (which could have been caused by failure of one of the routers or re-entry of its routing tables). This change affected only limited number of sites, which means that re-routing happened only for that particular part of the Internet.

In addition, if the hopcount for two sites have high correlation coefficient, then the routes to those sites go (at least to some extent) over the same machines. Because of this the distribution of the correlation coefficient (and especially limited number of sites with high correlation coefficient) could indicate the extent of diversity of the sites in the set used for this report. The fact that the hopcounts of the individual sites are mostly independent and the probable number of routers in the current Internet of around 98000 leads us to conclusion that in the Internet today the number of connections is so large that overlap very rarely occurs.




Figure 9 the percentage of the different hopcounts in all measurements from hopcounts in the designated measurements

3.5. Changes in paths

Next topic of this paper is the changes in paths that occurred during these tests. For every site and every measurement the path to that site was recorded. These paths were compared in order to see how many different paths were used for these sites in the course of these measurements. Furthermore we investigated how often paths change, how significantly the hopcount is influenced and where changes mostly occur.

Problem we encountered was that certain records of the paths contained '*' instead of the IP address of the hop. This means that a traceroute timeout occurred, because it is impossible to say at which machine this timeout occurred we decided to exclude all the files where this timeout occurred. In measurements of the hopcount we did not exclude such files because they did not influence the results. The resulting distribution is represented in Figure 10. There are few sites for which the number of routes is large, for the site 138.87.133.3 (site at the Illinois State University) there were 107 different routes in 162 tests. After studying routes of these sites we discovered that the variation of the hopcount for these sites remains limited (0.38 for mentioned site). Still for the most of the sites from our set the number of routes did not exceed value of 7 (241 sites had 7 or less routes). It is also interesting to see that there are 32 sites for which in all of the measurements we found only 1 route (we have to note that for only 2 of them in all tests a timeout did not occur).




Figure 10 Distribution of the sites on the number of the routes for each site during the measurements. We also studied the change in the paths to the sites for which the hopcount remained constant during these tests. Of the 66 sites to only eight of them packets went over one route. For one site there were 11 different routes although the hopcount remained the same. Hence, a constant hopcount does not give precise indications about the number of paths over which packet are forwarded to the destination. Yet, we have to add that in most of cases the hops that are different between paths belong to the same sub-network.




Figure 11 Distribution of the routes for the sites with constant hopcount value

Conclusion

Though traceroute program is rather simple to use and the hopcounts are easily extracted from its output, making a reliable hopcount measurement for certain region is not an easy task. One of the major problems is the fact that not all the routers send the ICMP messages through and that some of the firewalls discard these messages. This has limited the number of addresses suitable for our tests and complicated the completion of the test sets. However due to a size of the samples used in these measurements, results presented in this report do represent the average hopcounts for these three regions.

Regarding the dynamics of the hopcount, the value of the average hopcount did not exhibit large short-term changes but the difference in variation was quite large. During the course of these tests the hopcount for the addresses from the North American region, which resulted in 25 % of the hopcounts for these addresses to be different after one week of measurements.

Studying the correlation coefficients we discovered that the hopcount for certain groups of addresses are strongly correlated – though these sites are not in geographical vicinity of each other nor on the same sub-network.

Further we studied the change in the paths. Although the value of the hopcount often remains the same, there are a rather large number of the routes the packets used to reach the designated sites. Even for the sites where hopcount remained constant for each measurement it cannot be concluded that the routes remain the same.

References

- [1] M. Abramovitz and I.A. Stegun. *Handbook on Mathematical Functions*. Dover Publications, Inc., N.Y., 1968.
- [2] B. Flury. *A First Course in Multivariate Statistics*. Springer-Verlag, Inc., N.Y., 1997.
- [3] A. Leon-Garcia. *Probability and Random Processes for Electrical Engineering*. Addison-Wesley, Reading, M.A., 1994
- [4] W. Matthews and L. Cottrell. “The Pinger Project: Active Internet performance Monitoring for the HENP Community”, submitted to *IEEE Communications Magazine on Network Traffic Measurements and Experiments*
- [5] V. Paxson. End-to-end Routing Behavior in the Internet. *IEEE/ACM Transactions on Networking*, 5(5):601-615, October 1997
- [6] W.R. Stevens. *TCP/IP Illustrated, Volume 1: The Protocols*. Addison-Wesley, Reading, M.A., 1994.
- [7] S. Vanhastel, B. Duysburgh, and P. Demeester. “Performance Measurements on the Current Internet.” *7th IFIP ATM&IP Workshop, Antwerp (Belgium), June 28-30*, session 2:1-11, 1999
- [8] P. Van Mieghem, G. Hooghiemstra, R. Van der Hofstad. “A Scaling Law for the Hopcount in Internet.” *Delft University of Technology, report 2000125*, October 2000, also submitted to *IEEE Transactions on Networking*
- [9] P. Van Mieghem, G. Hooghiemstra, R. Van der Hofstad. “Stochastic Model for the Number of Traversed Routes in Internet,” *submitted to PAM 2001*.