Analysis of Network Infrastructure Based on Active Measurement Data

Author name withheld

January 31, 2001

Abstract

In this paper, we study the Round Trip Times of around 10.000 traceroutes between various sites of the internet. Questions we tried to answer in our work include, what is the distibution of Round Trip Times on a hop basis. How stable are these Times, and does the shortest path, relative to the number of hops, gives us always a better result than if we used a longest path to the destination. We tried to answer all these questions keeping in mind that the network is a chaotic system, and the observations that we make may be true for only the time space that we analyzed, and might not be true in a different time period, especially if this period happen to be a transitive period for the network.

1 Introduction

Our project focuses on the router level of the Internet. In a previous work by [4], they found that despite the apparent randomness of the internet some surprisingly simple power-laws hold for the Internet topology. In this current project we are interested in finding out what is happening with the Round Trip Times of the end-to-end paths. Mainly we are interested in investigating aspects of RTT such as the stability of that time, the distribution we have for different hop lengths of various paths. Also we were interesting in finding out if the shortest path between a source and a destination always gives us a better result.

There has been previous work in this subject [2, 6]. These papers focused on the end-to-end behaviors and analyzed the behavior of the paths. They investigated issues like path Assymetries, out-of-order delivery and packet corruption. Another similar work [3] they investigated issues of routing based on performance and reliability and they showed cases where the shortest path is not always better than an alternative longer path.

The rest of this project is structured as follows. In Section 2 we present the data that we used and we also mention the preprocessing we had to do. In Section 3 we present the first result of our analysis which presents the Cummulative RTT function on a hop basis. In Section 4 we study the stability of the RTT time. In Section 5 we study if the shortest Path always gives us the better RTT time. Finally in 6 we present some conclusions and some future work.

2 Data

In our analysis we used the data provided from NLANR [5]. In their project, which is called Active Measurement program, they use 133 sites mainly from educational institutes. These sites traceroute each other every 10 minutes. The total size of data per day is around 1Gbyte. The format is the usual traceroute format which is the following,

```
Sun Oct 29 00:01:07 PDT 2000
1 137.229.71.33 0.528 ms 0.348 ms 0.341 ms
2 137.229.2.1 0.232 ms 0.216 ms 0.214 ms
3 198.32.40.129 0.265 ms 0.255 ms 0.251 ms
4 198.48.91.33 40.884 ms 40.868 ms 40.881 ms
5 198.32.170.2 40.924 ms 40.807 ms 40.817 ms
6 198.48.91.78 40.874 ms 41.127 ms 40.810 ms
7 198.32.8.50 69.518 ms 69.095 ms 69.116 ms
8 192.80.43.21 87.319 ms 87.551 ms 87.324 ms
9 192.80.43.137 87.719 ms 87.605 ms 87.582 ms
10 128.196.128.52 87.804 ms 87.525 ms 87.341 ms
```

In this project we analyzed the data for all possible end-to-end pairs for one day only, specifically for the October 29th of 2000.

Note that in our analysis we had to remove some data that were caused by network failures. For example we removed traceroutes in which we had routing loops inside, and also traceroute that for some specific hop we had a star, that is that the router didn't answer. We removed them if we had more than 3 such consecutive failures.

Also before analyzing the data we did a preprocessing and for each path we removed the 1% of the smallest RTTs and the 1% of the largest RTTs in order to remove possible errors or pathological conditions that would have affected our analysis. Note however that we look the 98% of the data.

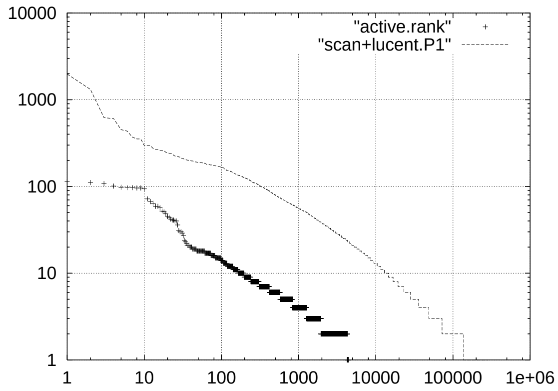

Figure 1: Comparison between the Real router topology and the active data paths

In order to evaluate the topology that was formed from the Active Router program we looked the topology that was formed from the traceroutes and we compared it with the a real router level topology gathered by ISI [1]. In figure 1 we can see the rank plot [4], which is used as a metric to see the properties of a graph. As we can see from the figure both graphs show the same properties, i.e, the same slope of the power-law, so we can argue that the active measurement topology we have reflects but on a different scale the properties of the real router topology.

3 The cumulative RTT distribution

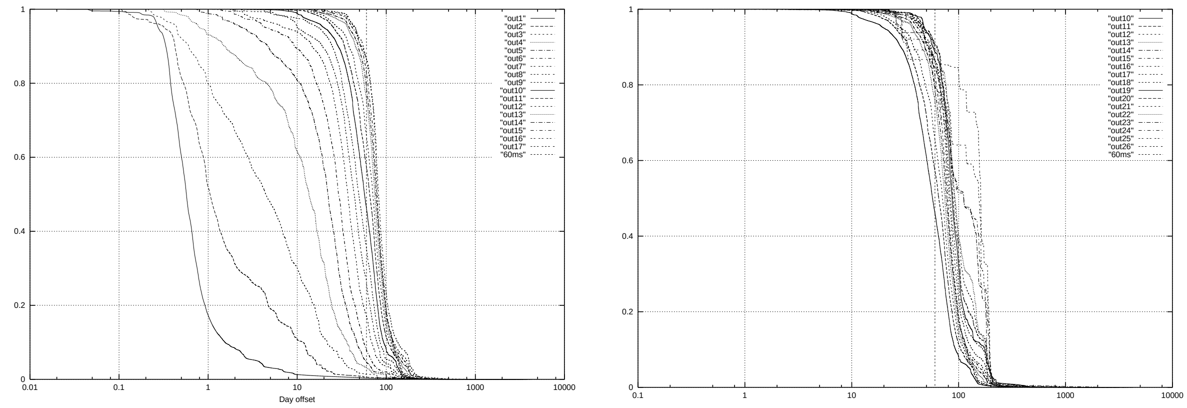

Figure 2: The Cummulative RTT distribution per hop. (a) for hops 1-17 (b) for hops 10-27.

In this section, we study the distribution of the RTTs of different number of hops. Recall that we define the D_{time}, of a path, to be the percentage of nodes that have RTT greater or equal to the RTT time. For example, for the first RTT that we find, we have all the nodes so the value is one, $D_{first} = 1$. For the next RTT, we have all the nodes except the ones with the first RTT, and the value is $D_{second} = 1 - first$.

We plot the D_{time} for all the hop values that we have, versus the RTT log-linear scale. In figure 2.a we can see these plots for number of hops from 1 to 17, while in figure 2.b we plot the cummulative RTT distribution from hop equal to 10 till 30 . Note here that we do not plot only the source - destinations pairs, but we also do plot that for all the intermediate hops that we have. In total each line represents more than 2 and half million points.

The major observation is that the plot is as we expected, since paths which are lengther have a bigger probability to have a higher RTT time. From these graph we can found out valuable information if our application has time constraints. By that we mean that we must deliver, or receive, our packets on a specific time interval, otherwise we consider them lost. Examples of this applications are multimedia traffic or real-time applications. Especially for interactive traffic in [7] they mention that in order to give sufficient time to the higher level protocols to process the multimedia traffic in order to have interactive sessions, the network delay must be around 30ms. Since we have RTTs this value is 60ms. As we can see from the figures, in order to have a good probability for an interactive session you must be around 10 hops away from your destination.

4 Stability of RTT

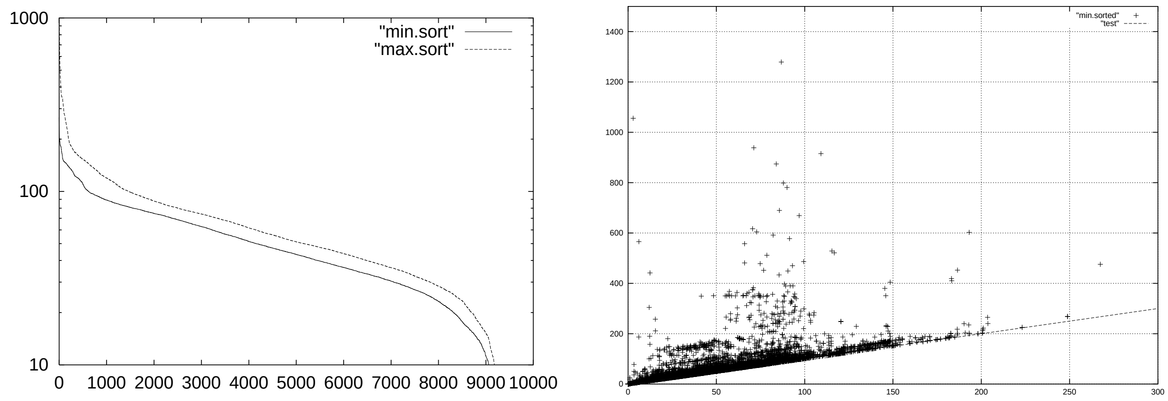

Figure 3: a. average min and max RTTs b. minimum versus maximum for each path

In this section we want to study how stable are the Round Trip Times we get for a particular path. In order to do that we manipulate the min and max Round Trip Times on a specific path, which shows us the variations we can have. First we present the average case where we plot the average difference between all the max and min Round Trip Times and then we plot for each specific path the min time and the max time, in order to see the variatons on a per path base.

In figure 3.a, we plot the distribution of the average min and max Round Trip Times. We do that by sorting the files so that we have the biggest values first and then we plot these sorted files. In the x axis we simply have the number of such min or max instances we have. As we can see the distributions look similar and we have a given gap between them as expected. As we can see from the figure it seems like the difference between the minimum and the maximum to be equal, except for the very high maximum numbers of RTT.

In the first figure it seems that the difference remains on average the same, but we also want to see whether we have cases that have the Round Trip Time difference smaller than others and at which Round Trip Times value this happens. In figure 3 we plot the minimum RTT seen for a path against the maximum RTT found for that path. The closer the points in the straight line, the closer the max is from the min value. As we can see we have a significant variaton and the variaton doesn't seem to be depended on the value that we have for the min value, since we can find big spikes even for small min values.

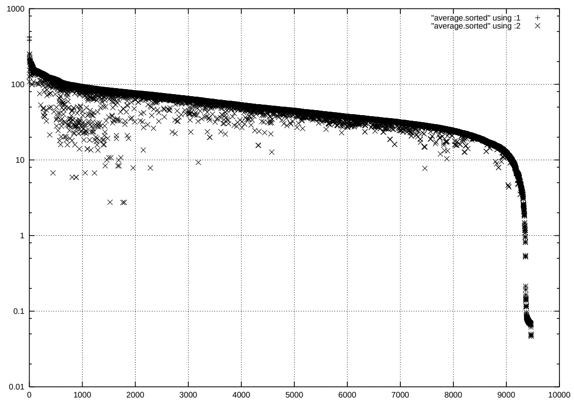

Figure 4: Comparison between the Real router topology and the active data paths

5 Triangular Inequality

As we have so far we have computed the minimum and maximum RTT distribution, and we have seen how stable are the RTT for a path from the minmax plot. It would be interesting to find out if the shortest path routing that we use will give us a worse result than if we use a detour.

We can see in picture 5 our idea of investigation. We have that the average RTT time from a source X to a destination Z is k ms. Is there any detour we can take that would give us a better result. That is are there any paths from source X to middle destination Y and from Y to the final destination Z so that the RTT time would be smaller.

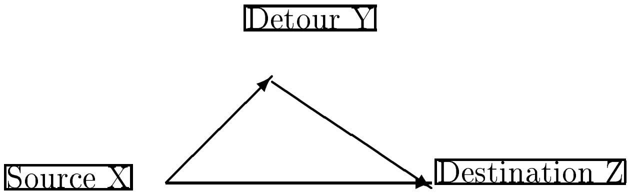

`Picture 5: Paths between a source and a destination, the shortest X>Z and the lengther X->Y->Z.`

We investigate this problem exhaustively. That is for every pair of source - destinations , which are around 10000, can we find a detour if we choose as a middle destination another point. We can see the result in figure 4. In this figure we plot the original average RTT and we also plot the new better value that we can find if we take the detour. As we can see from the figure there are a lot of points where we have a better result, especially for cases where the original RTT is quite large. In particular we found out that out of the 9470 end-to-end paths that we have, we can have a better average RTT in the 2969, which is a noticeable big proportion of the original number. Actually that number showed as the unique paths improvements we can have. In total we had 13415 better paths to follow than the original. This can happen since for one path we can find more than one better paths.

This result is of particular interest for network applications that have time constraints. If an application can not quarantee the in time delivery of the packets for a path, we find out that there is a high probability to find a lengther path that can give lower Round Trip Times, thus higher probability for acceptable performance. Note however that most of the sites that we test are academic institutions so this result might not hold in general for the commercial part of the internet.

If we want to use this result we must develop on-line algorithms that could trace continously the state of each path and be able to inform the routers or the applications for the existence of a better path.

6 Future Work

The contribution of our work is an average analysis of the Round Trip Time of a significant large number of paths and also the discovery that with high probability we can find indirect paths that can give us a better on average performance than the direct paths.

As we saw in the previous section, there exists a significant number of paths that we can have a smaller RTT time if you use not the shortest path, but instead using an indirect path. It would be very interesting to create an adaptive algorithm that would scan these cases and find out the stability of these results.

Another interesting area to investigate is the RTT time not on a router level but rather on an AS-domain level. We can find from BGP-protocol data the AS that each router belong to and find out if there are specific domain that cause the problem. Also in general it would be valuable if we could cluster routers/AS-domains in a way that describe areas of the network that cause problems.

A final future work will be to analyze and validate the results we obtained not only for one day but for other days too, and find out whether this is a transitive period that's why we observed that results, or this is something we can with high probability observe.

References

[1] Router level topology of the internet. http://www.isi.edu/scan/mercator/maps.html.

[2] M. Allman and V. Paxson. On estimating end-to-end network path properties. Sep 1999.

[3] Savage S. Anderson T. Aggarwal A. Becker D. Cardwel N. et all. Detour : Informed internet routing and transport. *IEEE Micro*, 1999.

[4] M. Faloutsos, P. Faloutsos, and C. Faloutsos. Power-laws of the Internet topology. *ACM SIGCOMM. Computer Communication Review.*, Sep 1-3, Cambridge MA, 1999.

[5] NLANR. Active measurement program by nlanr. http://watt.nlanr.net/.

[6] Vern Paxson. End-to-end internet packet dynamics. *IEEE/ACM Transactions on Networking 7.3 pp. 277–292*, 1997.

[7] H. F. Salama, D. S. Reeves, and Y. Viniotis. Evaluation of multicast routing algorithms for real-time communication on high-speed networks. *IEEE Journal of Selected Areas in Communications*, 15(13):332–345, April 1997.