MEASUREMENTS ON DELAY AND HOP-COUNT OF THE INTERNET

Aiguo Fei, Guangyu Pei, Roy Liu, and Lixia Zhang
Department of Computer Science
University of California
Los Angeles, CA 90095
{afei, pei, royliu, lixia}@cs.ucla.edu

Abstract

To find out how big the Internet is, we measured the round-trip delays and hop-counts from a UCLA host computer to a randomly selected set of three thousand Internet hosts around the world. Our results show that over 90% of these hosts in continental US are within 18 hops from UCLA, and the round-trip delays to 90% of these hosts are less than 153ms. There seems no strong correlation between the delay and hop-count, although the average delay increases with hop-count. Measurements to international hosts show that the delay and hop-count strongly depend on the countries the hosts locate. Physical distances and link speeds are the most important factors that determine the round-trip delay.

1 Introduction

The Internet has experienced exponential growth in recent years. By estimate[H98], there are about 30 million hosts connected to the Internet at the time this paper is being written, and this number is increasing everyday. To design network protocols and technologies that can scale with such rapid growth, it is important to know how big the Internet "size" is, and how fast this size grows. Although a number of measurement studies have been conducted over the last few years, most of them focus on the traffic characteristics, congestion control issues, and routing protocol stability. During fall 1997 we conducted a massive measurement effort aiming specifically at finding out how big the Internet was.

Internet measurement has a history as long as the Internet itself. Measurement experiments on the ARPANET packet delays was conducted as early as 1971 [K76]. More measurement studies [CPB93, CPB93-2, H90, PF95] were performed on the NSFNET after it replaced ARPANET in 1990 [MERIT]. More recently measurement studies showing the Internet routing instability and dynamics have also been reported in [GR97, LMF97, P97]. These measurement studies not only exposed the unexpected behavior of the current Internet protocols, but also help us better understand the dynamics of large scale systems.

End-to-end behaviors in the Internet, including delay and hop-count (number of hops along a path from one host to another), had also been the subject of a number of studies. In [K91], systematic measurements were taken to see how the network delay varied with different packet sizes, different paths, different times during the day, and different days in a week. In [B93], the author reported analysis of end-to-end packet delay and loss behavior from observing the round trip delays of small UDP packets sent at regular time intervals. In that study, compression of probe packets and rapid fluctuations of queueing delays over small intervals were observed and analysized by applying some queueing model. During spring 1996, Rautman, a UCLA graduate student, used traceroute to measure the delays and routes from UCLA to three specific sites at USC, MIT, and UCL(University College London) [R96]. Rautman's main interest was to find out how the network delay may vary with time and day as in [K76]. He reported that the variance in delay from day-to-day is not large, although weekend days tend to have less delay. Network delay varies with different time of the day, however there is no definite correlation between the time and the delay. Furthermore, the three destinations exhibited different delay variation patterns. Some end-to-end delay and hop-count measurements are also reported in some research of choosing replicated Internet servers [CC96, GS95]. Along with the examination of different approaches for locating nearby replicated Internet servers, the authors in [GS95] discussed an optimized approach for hop-count probing and presented some statistics of Internet hop-count. For example, they reported an average 17.0 hops among 8,098 Internet site pairs. In [CC96], the authors showed that empirical distributions of hop-count and round-trip time to 5,262 Internet servers are dramatically different. They were interested in how good delay or hop-count is as a distance metric in selecting replicated Internet server.

The main objective of our study is to answer a simple but fundamental question: how big is the Internet? Our definition of "big" is not measured by the population, that is how many hosts connected to the Internet, but rather by the size, that is how long is the path (in terms of hop-count) and the delay from one host to another. For example, how long and how many hops does it take to reach all the hosts out there in the Internet? What difference can one expect if one is to access two hosts that are located, say, in New York and Australia (given the source is here at UCLA)?

We measured the round-trip delays and hop-counts from a host at Computer Science Department of UCLA to 3,219 hosts in four continents. We examined the delay and hop-count distributions of hosts chosen from different US domains, different geographical locations within continental USA and different countries/areas. One of our goals was to understand how geographical distance affects the delay and hop-count, how different the delay and hop-count would be for hosts in different countries. We took both the delay and hop-count measurement at the same time to see how the delay is related to the hop-count.

The rest of this paper is organized as follows. We first describe how we did our measurements and how we picked out those hosts in the next section, then we present our measurement results and analysis in section 3. Two measurement-related issues, "Internet diameter" and Internet mapping are discussed in section 4, followed by a brief summary in section 5.

2 Measurement Method

To collect hop-count and delay data we wrote a small program based on the traceroute [J88, S94] utility originally written by Van Jacobson[J88]. Here is a short description of how it works. For any destination, it sends a 48-byte UDP packet to it with TTL (Time To Live) starting from 1 until the destination is reached. For example, if the destination is n hops away, for any TTL<n, the UDP packets cannot reach the destination, and the intermediate node which receives a packet with TTL=1 sends an ICMP(Internet Control Message Protocol) time-exceeded error message back to the source. In this way, the intermediate nodes can be tracked out. At the same time, the UDP packet uses a port number which in general will not be in use, so when the destination receives it, it will send back a port-unreachable message, thus the program knows destination is reached. When it is known that the destination is reached, our program sends a number of packets (we used 20 in our measurements) to the destination one by one in a stop-and-wait fashion, with a timeout of 5 seconds. The time from sending a probe packet to receiving the reply is the round-trip time (RTT). After all the packets are sent and replies are received (or timed out), the average is taken as the round-trip delay. One may take half of the round-trip delay as one-way delay, but since routes may be asymmetric[P97], it can only be an approximation. In this paper we report the round-trip delay only. We ran our program on a Sun Ultra Sparc-1 machine with Solaris 2.5.1 to collect all the data.

Some details in the measurement are worth mentioning. Sometimes a probe packet receives no reply. This can be caused by a number of different reasons: the probe packet or the reply may have got lost, or a router may be configured not to send back time-exceeded ICMP message, or it only generates ICMP messages at a limited rate [P96]. Without receiving a reply within the timeout period, a second packet with the same TTL will be sent. The timer we used is 5 seconds; our measurements show that delays to all hosts reached, except those in China, are far less than 5 seconds. If no reply is received for three consecutive packets with the same TTL, then that node in the route is treated as unknown and TTL for next probe packet is increased by 1. If no reply is received for 5 consecutive TTL values, our program will report a failure. It is possible that our measurement returns a failure but the destination is reachable, but that possibility should be small based on all our observation. If network- or host-unreachable messages are received for 3 consecutive probe packets with the same TTL, it is treats this as a failure too. Sometimes such error message can be generated because of administrative configuration of the intermediate or destination router, not because a network or destination really can't be reached, but one can't tell. Another detail worth mentioning is that the delay we measured is only for the packet size we used, packets of different sizes may experience different delays.

We need a set of hosts randomly selected from the global Internet as the destinations for our measurement. We found a list of DNS servers from the InterNIC ftp site[NIC], and randomly picked a set of IP addresses from that list as our study subjects. In order to study the effect of physical location on hop-count and delay, we also hand-picked a number of hosts. We divided the continental US into four regions and picked a number of web servers of universities from each region. We used web servers of universities because we know for sure the geographical locations of those universities and the web server names are easy to figure out. Hop-count and delay to hosts in China is one of our interests, unfortunately the list from the InterNIC contained only few sites in China. We visited the homepage of CERNET(China Education and Research Network)[CERNET] and found a list of Chinese universities connected to CERNET.

3 Measurement Results and Analysis

Our first measurement is delay and hop-count to US hosts. We picked out a total of 1,617 hosts from six major US domains (com, edu, net, gov, org, mil), statistics of delay and hop count are shown in table 1a and 1b. "Std." in tables stands for standard deviation and "avg." stands for "average".

Table 1a. Measurements of US Domains: Delays

domain	# of hosts	# of success	delay of 90% hosts ≤	median delay	avg. delay of low 90%	delay std. of low 90%	avg. delay of low 95%	delay std. of low 95%
com	429	334	175ms	93ms	87.8ms	37.4ms	92.8ms	44.6ms
edu	418	395	112ms	82.5ms	72.7ms	25.6ms	74.9ms	26.8ms
net	406	377	170ms	92.5ms	84.6ms	31.3ms	90.5ms	39.0ms
gov	171	135	113ms	81ms	66.5ms	28.2ms	69.1ms	29.5ms
org	100	81	139ms	93ms	82.3ms	28.7ms	86.2ms	32.6ms
mil	93	73	238ms	141ms	133.9ms	48.2ms	158.8ms	137.6ms
total	1617	1395	152.5ms	88.4ms	80.6ms	31.1ms	85.6ms	37.0ms

Table 1b. Measurements of US Domains: Hop Count

domain	avgerage	std.	median	hop count of low 90% ≤
com	13.2	3.0	14	18
edu	14.2	3.9	15	20
net	13.0	2.8	14	18
gov	13.1	2.7	14	17
org	12.9	2.5	14	17
mil	14.5	2.8	15	17
total	13.5	3.2	14	19

We put all the hop-counts and delays together and generated the distribution graphs of delay, hop-count and delay vs. hop-count as shown in Fig.1. In this and all other figures, the height of a bar represents number of hosts for a given hop-count or delay range (e.g. from 50ms to 60ms), it can be seen as a plot of pdf (probability density function). In the plot of average delay vs. hop-count, the length of the error bar is twice as much as the standard deviation, thus it is 2/3 significance level of delay for hosts of a given hop-count.

We can see from the above tables that there is some difference among different domains. Since we didn't take sample hosts from different domains according to their real percentages (e.g. there are more hosts in .com domain than that in .edu domain, but we took about same number of .com and .edu hosts in our measurement), strictly we couldn't simply add them all together as in Fig.1. But it should be close to the distribution with samples taken according to their statistical percentages if we do this way, because basic shapes of distributions for different domains look similar despite minor difference in average or median. Distribution graphs for different domains are shown in Fig.2. This similarity suggests that the Internet in US is "homogeneous" with respect to domains. Some detailed difference exhibited in our measurement may partly be due to our rather limited sample size.Another concern is that .net and .mil domains may have hosts outside continental US, thus we also made graphs of distributions without results from these two domains (not shown here). The resulting graphs look very similar to the graphs including all domains, indicating that few hosts we picked are outside of continental US. We can also see that the difference in hop count among domains is far less significant than that in delay.

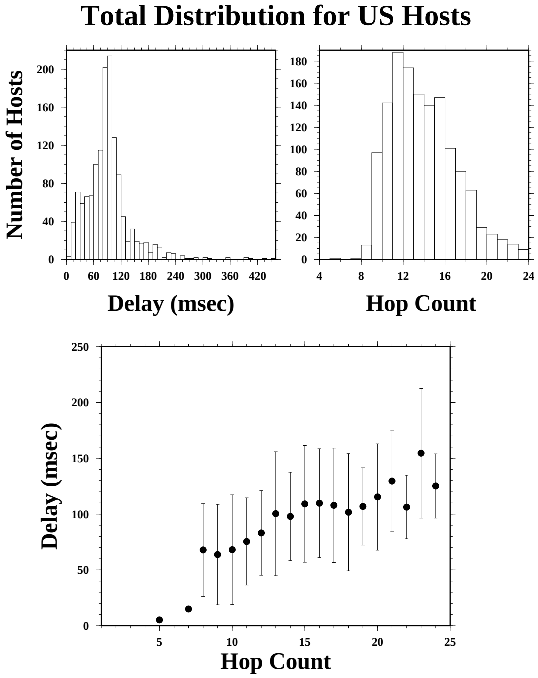

Fig.1 Distribution of delay, hop-count and delay vs. hop-count for hosts in US

Table 2. US Regional Measurements

region	#of hosts	# of success	average hop count	hop count std	average delay	delay std.	median hop count	median delay
West	25	19	10.7	2.2	33.6ms	19.4ms	12	26.5ms
Mountain	32	32	13.7	2.7	60.2ms	24.9ms	14	56.0ms
Central-east	91	84	14.4	3.2	94.8ms	25.8ms	16	88.9ms
East	65	60	15.5	3.3	99.3ms	15.4ms	16	96.0ms

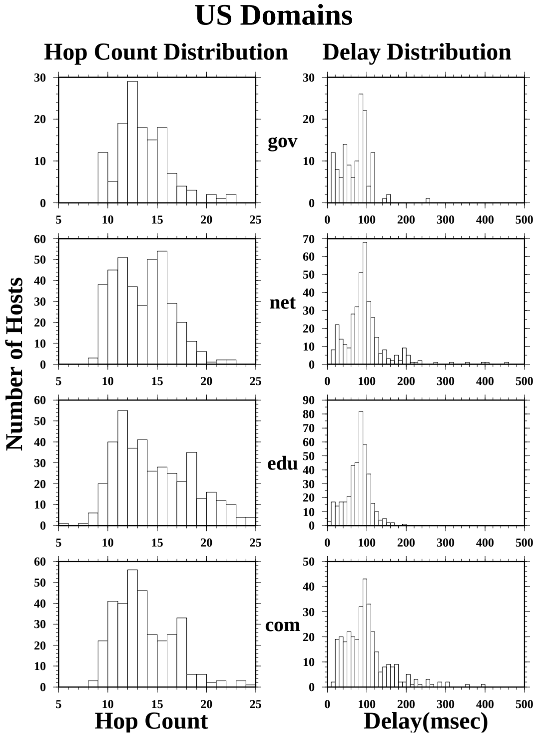

Fig.2 Distribution of hop-count and delay for four major US domains

From Fig.1, one can see, in general, that average delay increases with the increase of hop-count, though the relation is not linear. At the same time, the standard deviation of delay is comparable with the mean delay. This means there is no strong correlation between hop count and delay. In other words, one can't accurately predict the delay to a host given the hop-count, as suggested in [CC96]. Our observation shows that delay is not simply determined by number of hops, it depends on a lot of other factors, including physical distance, distance to the backbone, link capacities and traffic conditions along the route. This also demonstrates the great heterogeneity of the Internet.

As pointed out by Rautman[R96], delay varies with time-of-day and day-of-week. The results we show here were obtained during weekdays, most measurements last from afternoon to night, some were done at midnight. Because our measurements have a large sample space, it takes a long time to finish (more than one hour for 100 US hosts, longer for international hosts), we were unable to do measurements at some specific time and compare. However, according to Rautman[R96], the variation is not very large, especially when the delay is long (e.g. about 100ms to MIT), it is about 10% to 20%. Although the delay to a given host fluctuates, our large sample space should minimize the effect of a single event. So we believe the time of conducting our measurements should not affect the validity. To examine how delay may vary with time-of-day, we randomly picked a subset of 200 hosts from .com domain, did one measurement during the day which lasted from noon to 3:00pm and did another measurement at midnight which lasted from 0:00am to 2:30am. The difference between the mean delays from these two measurements is about 10%. The mean and standard deviation on hop count remained about the same (only a 0.1% difference on mean). This result agrees with observations in [P97,R96] that, though there is certain dynamic variation of routing, the route change doesn't happen very often (a dominant route exists), and the variation of hop-count is minimal. For international measurement, because of time zone difference, the effect of time-of-day should be even less.

The results for US regional measurements are shown in Table 2 and Fig.3. We call the four regions West, Mountain Area, Central-East, and East. Using standard state name abbreviations, West contains WA, OR and CA. Mountain area contains MT, ID, WY, NV, UT, CO, AZ and NM. Central-East contains ND, SD, NE, KS, OK, TX, MN, IA, MO, AR, LA, WI, IL, MI, IN, OH, KY, TN, MS and AL. East contains ME, VT, NH, NY, PA, WV, VA, NC, SC, GA and FL. They are going from west coast to east coast with increasing distance to our measurement starting point. A number of hosts from each region were picked as described in the previous section.

The distributions of hop count and delay are shown in Fig.3. From the table above and that figure, it is clear that both hop-count and delay increase with the increase of physical distance. It suggests that, at least inside the US, physical distance is an important factor on hop-count and round trip delay. At the same time, one also observes that the physical distance has a bigger effect on delay than on hop-count. This can be attributed to the fact that, most wide-area traffic is routed through the backbone, a few hops on the backbone can route traffic from west coast to east coast, while propagation delay is what one can never beat. Taking a signal propagation speed $2 \times 10^8 m/s$ (2/3 of light speed, for signal in fiber), the round-trip delay is $40ms$ for a distance $2,500miles$(from Los Angeles to New York).

Table 3. International Measurements

country/area	#of hosts	# of success	average hop count	hop count std.	average delay	delay std.	median hop count	median delay
Canada	106	90	14.80	2.2	116.74ms	72.3ms	15	105.4ms
Australia	213	160	13.76	1.8	399.68ms	181.6ms	15	410.4ms
Germany	97	86	14.50	2.8	211.84ms	110.2ms	14	186ms
France	104	88	23.70	3.4	202.38ms	97.5ms	26	182.5ms
UK	248	211	16.53	2.2	269.68ms	156.7ms	18	219.5ms
Italy	104	100	17.73	1.7	270.72ms	77.8ms	18	255ms
China	126	114	20.27	1.6	1537.57ms	1257ms	21	946ms
Japan	289	247	18.07	3.4	317.97ms	456.8ms	19	266ms
Taiwan	55	53	16.26	1.4	304.06ms	29.6ms	17	296ms
South Korea	49	36	13.39	2.2	254.89ms	62.7ms	14	219ms

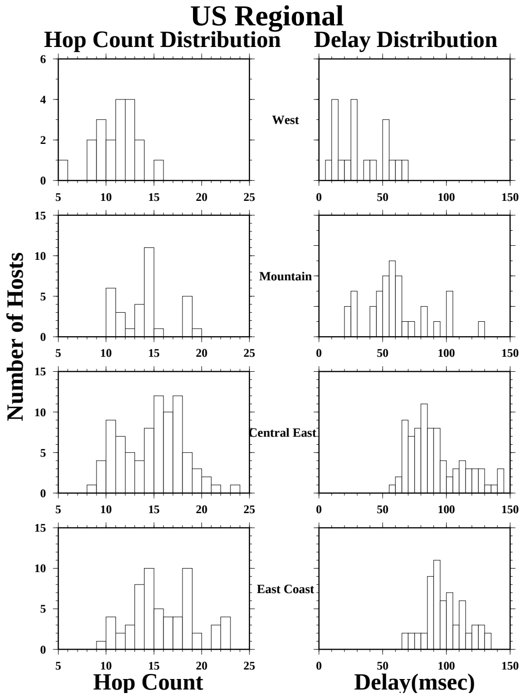

Fig.3 Measurements of US regional

We did our measurements to 10 countries and areas outside of USA including Canada, Australia, four in Asia, and four in Europe. The results are shown in table 3, Fig. 4 to Fig. 7.

All these show that hop-count and delay to an international host heavily depend on the specific country the host is in, and vary over a wide range from country to country. While network condition inside that country plays an important role, physical distance and the connection between US and that country are also important factors. As seen from the difference between Canada and Australia in Fig.4, hop counts to these two countries look similar, but there is a dramatic difference in delay due to the difference in physical distances. Comparing the results of China and Japan (Fig.6), these two countries have similar physical distance from US, the hop counts are close too, but the delays to these two countries are dramatically different. We found out from CERNET homepage[CERNET], there were only two links between China and the US at the time of our measurement, one is 128Kbps and the other 2Mbps, while links inside China were of very low speed too.

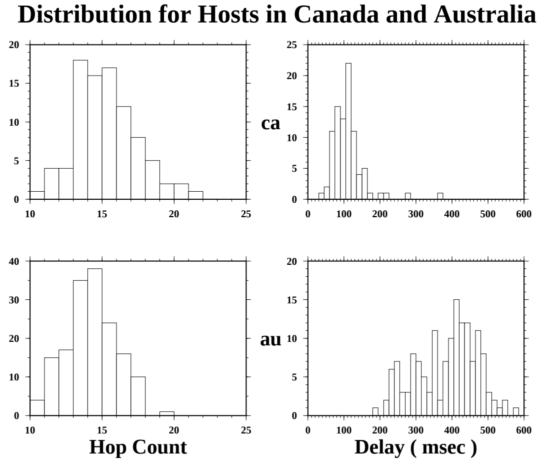

Fig.4 Measurements of Canada and Australia

The data for the four European countries shows a very interesting phenomenon. Almost all hosts with hop count 10 are UK hosts, and these hosts have a longer delay than other hosts with hop-count over 20. At the same time, most hosts with hop-count 25 or 26 are in France, and the delay is even shorter than hosts in other countries with hop count less than 20. This suggests that, there is a link to UK with few hops but pretty slow, while there is a link to France with a number of nodes but pretty fast. Examining the traceout data, we identified a common path of 9 hops (from UCLA to mci.net then to demon.net) shared by 8 UK hosts and all of them have a total hop count of 10 and a round-trip delay around 500ms. While the delay up to MCI's last hop is about 80ms, the delay up to ermin-router.router.demon.net is about 500ms. It is clear that demon.net was the ISP shared by those hosts and introduced the long delay. We also identified a common path of 12 hops shared by 79% UK hosts which has a delay about 200ms up to the last hop JANET-gw.Teleglobe.net which should be in Europe because there is a 100ms gap between it and the hop before it. We also found 87.5% of France hosts shared the same path as long as 16 hops, from UCLA to MCI backbone to Sprintlink backbone and then to pennsauken-hssi.eurogate.net and 193.55.152.65, the delay up to 193.55.152.65 is about 150ms. Hop-count and delay distributions and delay vs. hop-count in Fig.7 were generated by putting measurements of four European countries together. As stated above, because of the difference among those countries, the result is not very meaningful. The main purpose of having them here is to show the delay/hop-count anomaly observed and compare the distributions with those from US.

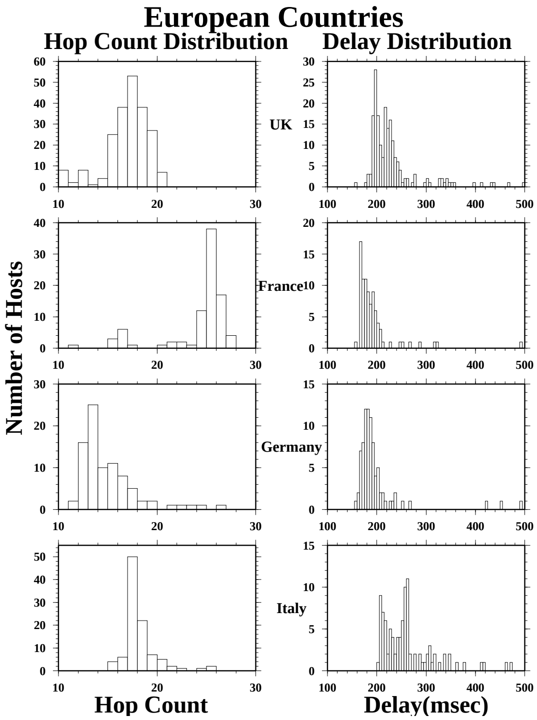

Fig.5 Measurements of four European countries

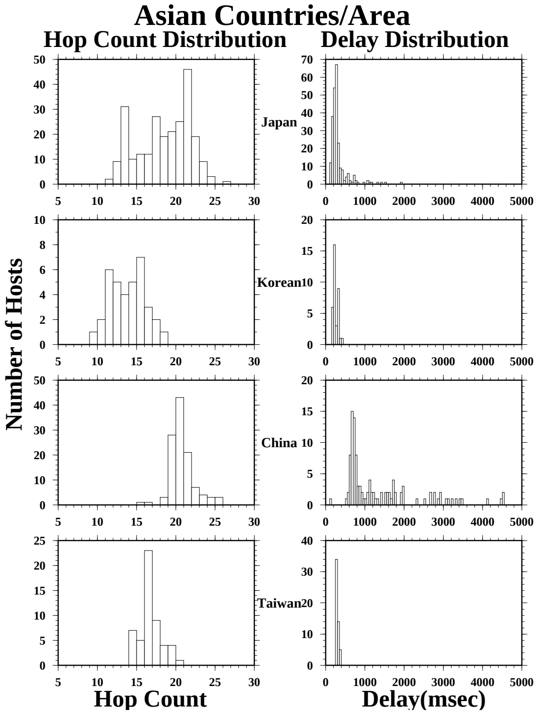

Fig.6 Measurements of Asian Countries/Area

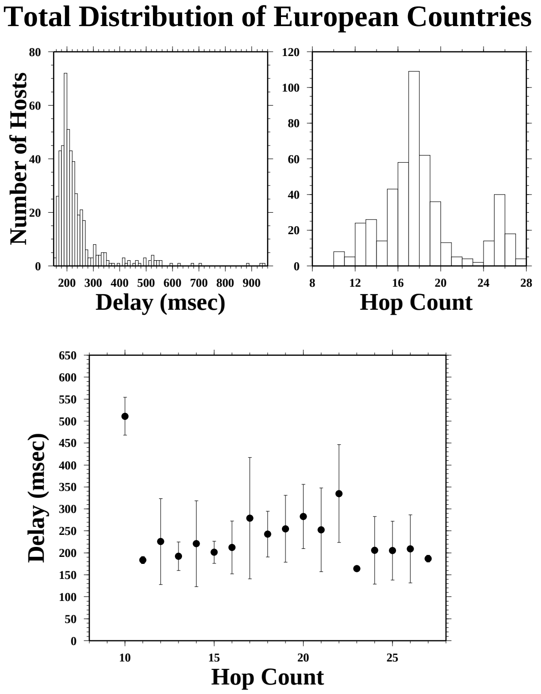

Fig.7 Measurements of four European countries put together: delay distribution, hop-count distribution and delay vs. hop-count

4 Discussions

One limitation of our measurements is that the starting point is only at UCLA. So if one does measurement from a different place, delay and hop-count could be different. However, we believe the statistics should still be similar. Consider the delay, it has three parts: delay from source to the backbone, delay on the backbone, delay in the destination site. The first part depends on the topology at the source site, how the site is connected to the backbone. The third part most likely does not depend on the location of the source. The second part depends on the relative locations of source and destination. With a large sample space, the second and third parts would yield similar statistics respectively with measurements from different starting point and the main difference between measurements from different sources is the difference of the first part. The same holds for hop-count measurement.

Our measurement may help answer what's "the Internet diameter"[P96], which is interesting to some people and would be an important parameter to consider in large scale simulations and topology modeling of the Internet. Intuitively one may think "the diameter of the Internet" is the number of hops of the longest route needed to connect two far-apart hosts. In our measurement, the longest route recorded has a hop-count of 27. In measurement of [P97], the longest route is 32 hops. Such a diameter gives us the impression of how big the Internet is in size. However, since one can hardly claim a long route he or she observed is indeed the longest one, the above intuitive definition may not be easily measurable. A sound definition should also be a description of an executable empirical method to measure such a metric. An alternative metric to represent "how big the Internet is" would be the average hop-count because of some good properties of hop-count distribution. As shown in the previous section and in [CC96], hop-count is pretty much evenly distributed with the mean at the center. So average hop-count should give us a fairly good impression on how "deep" "vertically"[GR97] the Internet is. However, the Internet is so highly heterogeneous, and our measurement shows that routes to different countries are really country-specific, one has to be careful at choosing sample hosts when doing such measurement.

Another interesting problem is how the "diameter" or average hop-count grows with the growth of the Internet. The growth trend concerns the scalability of network protocols. Some people believe that it grows as the logarithm of the size (number of nodes) of the Internet. When Internet size grows exponentially, we expect its "diameter" to grow linearly. There is evidence that average hop-count does increase with the growth of the Internet. In [P97], the author reports that, the mean hop-count of routes measured during November and December 1994 is 15.6, while that measured during November and December 1995 is 16.2. To carefully study how the mean hop-count or "Internet diameter" increases with the growth of the Internet, systematic measurements have to be carried out from time to time on a regular basis.

People had been enthusiastic about "information superhighway" after the widespread use of Internet. In comparison with the widely-available freeway maps, some people including ourselves have been interested in "maps" of the Internet. There are maps which have geographical locations of switching systems, NAPs(network access point), subnetworks and their interconnections, and there are logical maps which show topology and interconnections only. Maps for NFSNET backbones can be found at Merit homepage[MERIT] and maps for the new vBNS can be found at vBNS homepage[VBNS]. There is an ongoing research effort to visualize backbones of different ISPs(Internet Service Provider)[MAPNET]. They have a database of backbone topologies of a number of ISPs and provide a Java applet to visualize the maps. However, an ISP needs to provide its backbone topology or peering information and commercial ISPs may be hesitant to provide such information about their networks. Also, their database only has topology information of different backbones, but doesn't have any information about interconnections and subnetworks. During our measurement, we collected thousands of routes at the same time. One interesting future research problem is to discover topology and interconnection information of different ISPs' networks and other subnetworks from the routes collected. If we can succeed in discovering the Internet topology, one will be able to discover information, which ISPs are unwilling to provide to public, by sending probe packets and collecting routes. A new research project, the Internet Distance-Map Service(IDMaps) [F97], aims at discovering such topology information (and more) to meet applications needs. Our measurement project could be considered an early experimental step in that direction. As shown in the last section, we successfully identified the common route shared by most UK hosts and that by most France hosts, and we also identified a slow path to some hosts in UK.

5 Summary

With the rapid growth in recent years, the Internet has become the biggest "lab" mankind has ever made. In this paper, we reported our measurement experiment conducted in this "lab". We measured hop-count and round-trip delay from our host computer at UCLA to more than 3,000 hosts worldwide, and examined the relation between delay and hop count. Our results show that, in the continental US, more than 90% of hosts can be reached within 18 hops and the round-trip delays to more than 90% of hosts are less than 153ms for our measurement packets. We also observed no strong correlation between hop-count and delay, although the average round-trip delay does increase with the hop-count. We also observed that the hop-count and delay to hosts in different countries demonstrate country specific patterns.

Our measurement on the network round-trip delay may provide useful information to a number of applications that need an estimate of delay of the underlying network. Delay and hop-count knowledge together may help researchers in choosing parameters for large-scale simulation and modeling of the Internet. Meanwhile, it is still an open research problem if we can discover Internet topology from routes collected through traceroute. We are confident that active probing utilizing the ICMP protocol is an effective way to do a lot of Internet measurements, but we also want to point out that it shouldn't be abused since it generates network traffic which can be significant sometimes. We realize that no measurement results can hold forever or everywhere because of rapid change and great heterogeneity of the Internet. Our measurement is best viewed as a snapshot of the Internet at the time. We expect that systematic approaches to Internet size measurement will be developed soon which will lead to periodic snapshots, giving us an accurate picture of both the Internet size and its growth rate and patterns.

References

[B93] J-C. Bolot. End-to-end packet delay and loss behavior in the Internet. In the Proceedings of ACM SIGCOMM 1993.

[CC96] R. Carter, and M. Crovella. Dynamic Server Selection Using Bandwidth Probing in Wide-Area Networks. TR-96-007, Boston University Computer Science Department, March 18, 1996. Available from http://www.cs.bu.edu/faculty/crovella/papers.html.

[CPB93] K. Claffy, G. Polyzos, H-W. Braun. Traffic characteristics of the T1 NSFNET backbone. In the Proceedings of IEEE Infocom 1993.

[CPB93-2] K. Claffy, G. Polyzos, H-W. Braun. Applications of Sampling Methodologies to Network Traffic Characterization. In the Proceedings of ACM SIGCOMM 1993.

[CNET] CERNET homepage.
See http://www.cernet.net/ and
http://www.edu.cn/cernet/.

[F97] P. Francis. Second outline of post-cannes Internet distance-map service.
See http://www.ingrid.org/maps/.

[GR97] R. Govindan, and A. Reddy. An analysis of Internet inter-domain topology and route stability. In the Proceedings of the IEEE Infocom 1997.

[GS95] J. Guyton, and M. Schwartz. Locating nearby copies of replicated Internet servers. In the Proceedings of ACM SIGCOMM 1995.

[H90] S. Heimlich. Traffic characteristics of the NSFNET national backbone. In the Proceedings of Winter USENIX Conference 1990.

[H98] C. Huitema. How many hosts in the Internet.
See ftp://ftp.bellcore.com/pub/huitema/stats/size-today.html.

[J88] V. Jacobson. Traceroute program, 1988. Available from ftp://ftp.ee.lbl.gov/traceroute.tar.Z.

[K76] L. Kleinrock. Queueing Systems, Vol2 2: Computer Applications, Chapter 6. Wiley-Interscience, New York 1976.

[LMF97] C. Labovitz, G. R. Malan, and F. Jahanian. Internet routing instability. In the Proceedings the ACM SIGCOMM 1997.

[MAPNET] Mapnet
See http://www.caida.org/Tools/Mapnet/.

[MERIT] MERIT homepage.
See http://www.merit.edu/.

[NIC] FTP site of InterNIC.
See ftp://ftp.rs.internic.net/domain/inaddr.zone.gz.

[P96] V. Paxson. Towards a framework for defining Internet performance metrics. LBNL-38952, June 20, 1996.

[P97] V. Paxson. End-to-end routing behavior in the Internet. IEEE/ACM Transactions on Networking, Vol.5, No.5, pp. 601-615, October 1997.

[PF95] V. Paxson, and S. Floyd. Wide-area traffic: the failure of Poisson modeling. IEEE/ACM Transactions on Networking, 3(3), pp.226-244, June 1995.

[R96] E. C. Rautman. Delay measurements on the Internet: some observations and analyses. CS219 Project, UCLA, Winter 1996. Available from http://www.cs.ucla.edu/ afei/inetmeasure97/er96.ps.gz.

[S94] W. R. Stevens. TCP/IP Illustrated, Vol 1: The Protocols, Addison-Wesley Publishing, 1994

[VBNS] vBNS maps.
See http://www.vbns.net/net_maps.html.