Analysis of End-to-End Measurements to Investigate Delays On the Internet

Rakesh Kushwaha and Mario Garzia
Service and Architecture Performance, AT&T,
101 Crawfords Cr Rd, Holmdel, NJ

rkushwaha@att.com *garzia@att.com*

Extended Abstract

End-to-End delays on the Internet have been a major concern for those wanting to provide quality of service in today's Internet. In this work, we analyze 36,000 end-to-end traceroute measurements to study delays in the current architecture of the Internet . The measurements were collected by Vern Paxson of Lawrence Berkeley Laboratory using repeated traceroutes between 37 Internet sites [PAX 96]. The data was gathered primarily to analyze the routing behavior on the Internet for pathological conditions, routing stability, and routing symmetry. The study concluded that Internet paths are heavily dominated by a single prevalent route and about 2/3 of the Internet paths had routes persisting for either days or weeks. Furthermore, by the end of 1995 the likelihood of encountering a major routing pathology (erroneous routing, infrastructure failures, and temporary outages) was about 3.4%.

Since the traceroute data observed by Paxson plausibly represent the Internet as a whole [PAX 96], and the data also provides rough estimates of round-trip delays between the originating site and each intermediate hop , we further analyzed the data to investigate the delay behavior on the Internet. We aggregated end-to-end traceroute measurements based on the number of hops. For different routes, with same number of hops we compute "per hop delays". The per hop delay distribution, allows us to categorize the hops as Class A,B and C hops. The hops with small delays are classified as class A hops, the hops with large delays as class C hops and hops with intermediate delays as class B hops. Next, from the data we estimate the percentages of each class of hops for every route. Analysis reveals that as the number of hops increase, so does the number of class C hops. Given the percentages of class A,B, and C hops, their average delays, and total number of hops in a route, we estimate end-to-end delay on a multi-hop route on the Internet. The estimated delays are compared with the measured delays.

An understandable criticism of our study is that traceroute does not provide accurate round trip delay measurements on the internet. There are several reasons for this shortcoming worth noting. The first difficulty is that in traceroute, TTL expiration is an exceptional event, hence some routers may take longer time in sending the ICMP reply than others. This means that the differential between the times for two different hops is biased by the difference in the times taken by the two routers to generate the ICMP, and that can be large. (We see negative delays per hop using this computation). However, our analysis reveals that the data does differentiate hops with large delays and hops with very minimal delays. We do not include negative delays in our analysis. The second concern is, if the traceroute delay data sample size is large enough to establish claim about average delays. In our analysis, we aggregate different geographical routes with the same number of hops in one category to compute per hop delay distribution, this gives a large sample to corroborate the claim. Lastly, the analysis here is not used for establishing a solid theory, but to analyze the delay behavior and build up an intuition for estimating end-to-end delay in the promiscuous Internet.

For end-to-end delivery on a multi-hop route, a packet passes through many autonomous systems (AS) which are administered by Regional Internet Service Providers (R-ISPs) or Local Internet Service Providers (L-ISPs) or LAN managers. For international routes, such AS are more in number. Within an AS itself, there could be several types of hops, of particular interest are the types located on the border of AS, which runs BGP/EGP and handle all the incoming and outgoing traffic for the AS. Heuristically, the more autonomous systems a packet passes through, the more will be its end-to-end delay. The observations of our analysis support this heuristic and also estimates end-to-end delay. Such estimates provide insight

into the quality of service (QoS) that can be provided today on the Internet and also establishes end-to-end delay benchmarks for all upcoming improvements on the Internet.

The end-to-end measurements for different routes, gathered on different days and times, were aggregated based on number of hops. The number of hops range from 3 hops to 29 hops. Fig 1 shows the distribution of per hop delay for 10 hop routes, and fig 2 shows similar distribution for 20 hop routes.

Fig 1: Per hop delay distribution for 10 hops (one way delays)

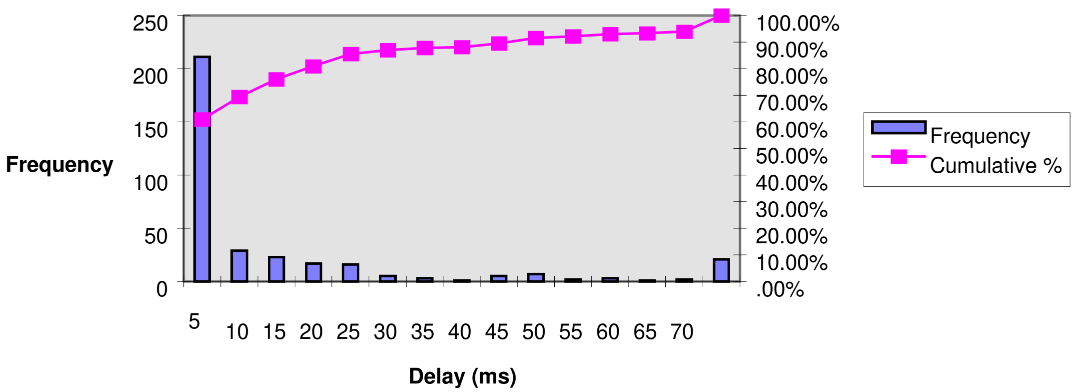

Fig 2: Per hop delay distribution for 20 hops (one way delays)

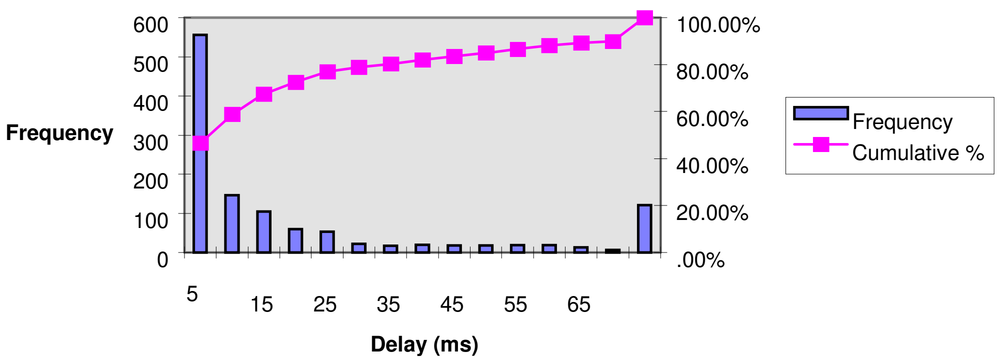

Based on the per hop delay distribution, we categorize the hops (routers/gateways) as class A, class B and class C hops. The hops with delays of 10 ms or less are categorized as class A hops, and hops with delays 70 ms or above are categorized as class C hops. The routers/gateways in neither of the above two categories are classified as class B hops. The classification is further summarized in Table 1.

TABLE 1 : Categorization of routers/gateways on the internet based on per hop delay measurements and functionality

Category	Sample Average Delay (ms)	Delay Range (ms)	Conjecture
Class A	5	0-10	Access Routers, hops running static routes, destination net-id matches net-id of directly connected network. (scenario of Ethernet LAN).
Class C	100	70-160	Border Routers, runs EGP/BGP, handles all outgoing and incoming traffic for an AS, performs packet filtering , layer 3 processing and perhaps other tcp/udp applications
Class B	35	10-70	All routers that do not fall in Class A and C; routers typically connecting few networks

From the data analyzed it can be seen that the percentage of class C routers increase from nearly 5% in 10 hop routes, to about 10% in 20 hop routes. The analysis further reveals that the percentage of Class C routers increase steadily as the number of hops increase in today's Internet. The percentage of class A, B and C routers for varying hops is shown in fig 3.

The measured end-to-end delay as a function of number of hops is shown in fig 4. From fig 3 and 4 it is evident that class C routers, the routers with large delays, play a prominent role in determining the end-to-end delays.

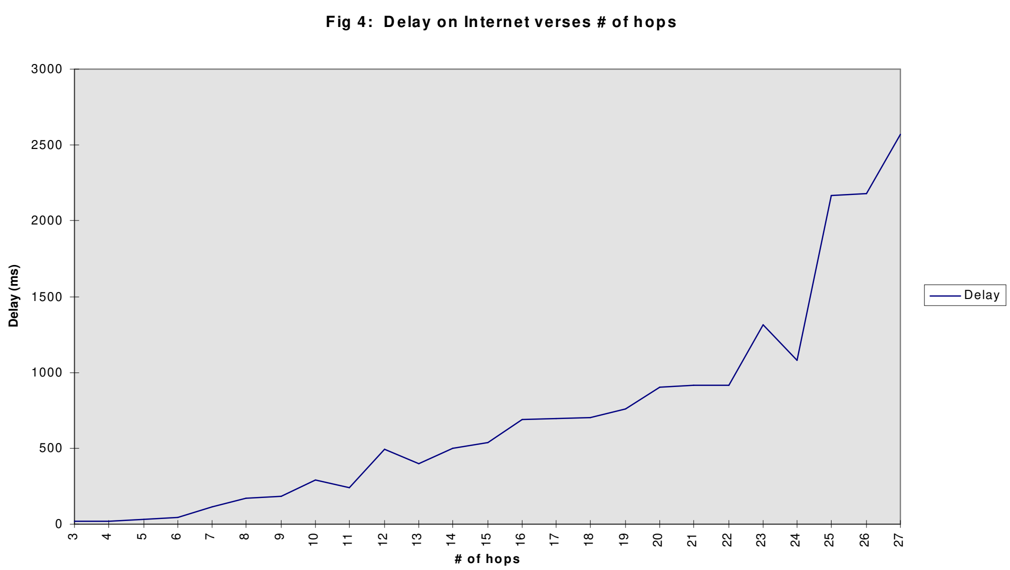

The end-to-end delay in a multi-hop route clearly depends on per hop delay. If hops in a internet route can be categorized as proposed class A, B and C hops, we can estimate end-to-end delay as :

$$\textbf{Est . Delay (end-to-end) = Total number of hops} * \sum_{i=classA,B,C} \textbf{(\% of } (i) \textbf{ hops * Ave. delay of } (i) \textbf{ hop)}$$

Assuming, the average delays of class A, B and C hops as shown in Table 1 and percentage of each class from the fig 3, we estimate end-to-end delay using the above equation for varying number of hops. The estimated delay with the above approximation is shown in fig 5, along with the measured delay from fig 4.

The result suggests that the proposed estimations predict end-to-end delays on the Internet within some reasonable range.

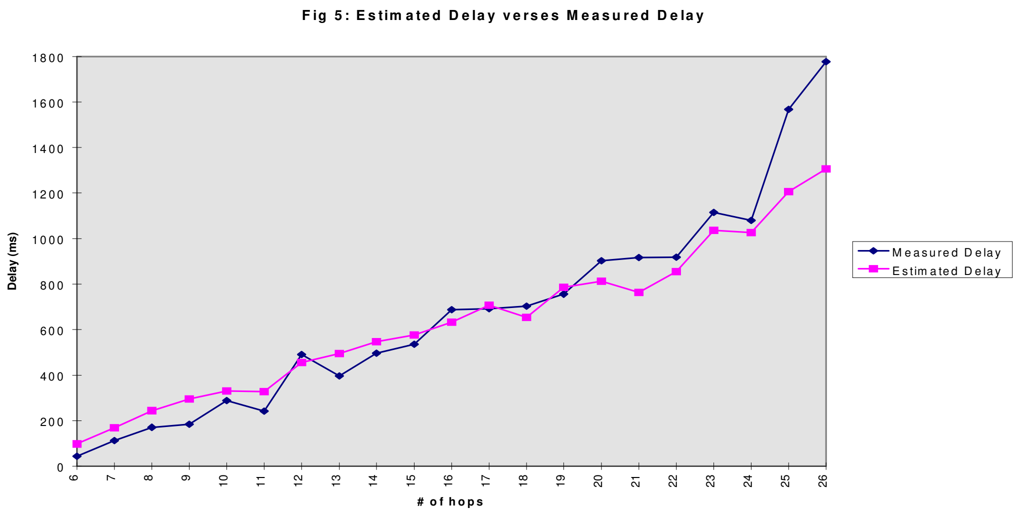

In summary, the study investigates the delay behavior on today's Internet, provides guidelines and a first order approximation to estimate end-to-end delays on the Internet. Given the categorization of the routers/gateways, their average delays and percentage numbers along the route, end-to-end delays can be estimated on the Internet.

Acknowledgments

Authors like to thank Vern Paxson of Lawrence Berkeley Laboratory for providing the traceroute raw data used for analysis, helpful suggestions and criticism throughout this study.

REFERENCES

[PAX-96] V. Paxson, "End-to-End Routing Behavior in the Internet," Proceedings of SIGCOMM'96, August 1996. (http://town.hall.org/Archives/pub/ITA/html/contrib/NPD-Routes.html)