Review of An Empirical Evaluation of Wide-Area Internet Bottlenecks

Lijie Heng
9032366867

Contributions of the Paper

This paper investigates potential non-access bottlenecks which may arise after fixing the bottlenecks at the edges of the network by updating the access links in the future. The first contribution of this paper is to propose a good method for measuring the non-access Internet bottleneck links by choosing representative sources and destinations. Besides, a new measurement tool BFind is developed specially for testing the bandwidth and location of bottlenecks.

The second achievement of this paper is to classify and characterize the bottleneck links according to their location, latency and bandwidth. The paper shows that nearly 50% of the paths having a non-access bottleneck link with available capacity less than 50Mbps. Moreover, the percentage of paths with bottlenecks grows when coming to lower-tier destinations.

Potential Areas for Improvement

One drawback in this paper is the maximum sending rate of BFind is limited to 50Mbps, in order to avoid using too much bandwidth of the network. However, if the bandwidth of a network is higher than 50Mpbs, bottlenecks can not be found out using the method in this paper.

Another point which can be improved is that in order to test the capacity of the bottleneck, BFind sends a large amount of data into the network, which could cause congestions in the transit path. The high overhead of BFind is an undesirable feature for a general-purpose probing tool.

Finally, the available bandwidth is a time-varying quantity. BFind is just to estimate the average available bandwidth. If the available bandwidth is considered as a stationary random process, the other statistics, like the variance of the marginal distribution and the autocorrelation function, are needed for a more complete characterization of the available bandwidth process.

Implications on the Future Internet

The key observations in this paper have a great significance for both ISPs and customers. It helps ISPs to evaluate their providers and peers, and helps customers to pick ideal upstream providers. It is a performance advantage to using tier-1 providers. And compared with tier 2, choosing tier-3 may be more beneficial based on the conclusion of this paper.

Another implication of this work is to guide the design of systems and choice of routers in the future. The conclusions of this paper tell us which routers should be chosen to achieve a better performance, which tier of ISPs should be used to get a biggest bandwidth and so on.