

Measuring Bandwidth between PlanetLab Nodes

Sung-Ju Lee¹, Puneet Sharma¹, Sujata Banerjee¹, Sujoy Basu¹, and Rodrigo Fonseca²

¹ Hewlett-Packard Laboratories, Palo Alto, CA
{sjlee, puneet, sujata, basus}@hpl.hp.com

² University of California, Berkeley, CA
rfonseca@cs.berkeley.edu

Abstract. With the lack of end-to-end QoS guarantees on existing networks, applications that require certain performance levels resort to periodic measurements of network paths. Typical metrics of interest are latency, bandwidth and loss rates. While the latency metric has been the focus of many research studies, the bandwidth metric has received comparatively little attention. In this paper, we report our bandwidth measurements between PlanetLab nodes and analyze various trends and insights from the data. For this work, we assessed the capabilities of several existing bandwidth measurement tools and describe the difficulties in choosing suitable tools as well as using them on PlanetLab.

1 Introduction

The lack of end-to-end QoS and multicast support in the underlying best effort networking infrastructure has spurred a trend towards providing application level intermediaries such as web-caches and service replicas to mitigate the performance issues. It is not only important to provide intermediary services but also to connect the end-clients to the intermediary that can meet the client QoS requirements and provide the best performance. For instance, web applications might want to select the nearest content cache while the online multiplayer game players might want to choose the least loaded game server. There also have been attempts to build overlays by connecting application level intermediaries for composable and personalized web and media services. Normally, such services have QoS requirements such as bandwidth and delay. Hence, building and maintaining such overlays requires periodic or on-demand measurement of end-to-end paths. Motivated by these trends, there have been significant research studies on active and passive network measurement techniques, and measurement studies from many large scale networks [1]. Clearly, periodic or on-demand measurement of all possible network paths will incur a high overhead and is inefficient. Thus a key concern is the development of scalable measurement and inference techniques which require minimum probing and yet provide the required measurement accuracy.

The primary network metrics of interest are end-to-end latency, bandwidth and loss rates while application level metrics are HTTP response times, media streaming rates, and so on. Many studies have focused on scalable network distance estimation, mainly using the triangular inequality heuristic [17]. However, similar triangular inequality does not apply to bandwidth, and hence it is much more difficult to identify the nodes that provide the maximal bandwidth without probing to each node.

Although there is a plethora of bandwidth measurement techniques, there have been only a few large-scale bandwidth or bottleneck capacity measurement studies. In fact, most of these studies are in conjunction with the validation of a new bandwidth measurement tool. New bandwidth measurement tools continue to be developed, aiming for better accuracy with faster measurements. In this paper, we present results from a large scale bandwidth measurement study on the PlanetLab infrastructure. Our goals are (i) to understand the bandwidth characteristics of network paths connecting PlanetLab nodes and (ii) to ultimately obtain insights into potential trends that will enable scalable bandwidth estimation. We primarily focus on the first goal in this paper. We do not develop a new bandwidth estimation tool nor evaluate and compare the accuracy of various tools. Rather, we assess the capabilities of a number of available tools from a PlanetLab deployment standpoint and report our findings in the hope that it will help other researchers to make an informed choice of a tool.

In the next section, we describe our methodology and the tools we assessed for this study, followed by an analysis of the data we collected. Section 3 concludes the paper.

2 Measurement Study

PlanetLab is an attractive platform for bandwidth measurement as it is an open, globally distributed network service platform with hundreds of nodes spanning over 25 countries. PlanetLab has gained the status of the de facto standard for conducting large scale Internet experiments. Although the interdomain connectivity of the PlanetLab hosts may not represent the global Internet [2], the characterization of PlanetLab topology is still of utmost importance for designing experiments on PlanetLab and drawing informed conclusions from the results. Several measurement studies have been conducted on PlanetLab topology, mostly focusing on the connectivity and the inter-node latency. In this paper, we study the bottleneck capacity between the PlanetLab nodes.

2.1 Methodology

Our methodology consisted of deploying the bandwidth³ measurement tool on a selected set of responsive nodes on PlanetLab using standard PlanetLab tools and then executing a script to run the measurements. The collected data is then shipped back to a central node on our site for analysis.

We performed two sets of measurements at two different time periods. The first set (referred to as *Set 1* in rest of the paper) was measured and collected starting in August of 2004, and the second (referred to as *Set 2*) in January of 2005. Between the two measurements periods, PlanetLab went through a major version change. The second set of experiments were performed after the version 3 rollout on PlanetLab.

Although there are over 500 deployed nodes on PlanetLab, only a little over half the nodes consistently responded when we started the measurement process. A crucial first step was to select a tool; we describe the selection process below. Conducting pair-wise latency measurements for a few hundred nodes is a relatively quick process for which measurements can be run in parallel and finishes in the order of minutes. However,

³ We use the terms bandwidth and capacity inter-changeably throughout the paper.

pair-wise capacity measurements for a few hundred nodes needs to be well coordinated because the capacity measurement tools often do not give accurate results when cross traffic is detected. Thus the measurement process for all pairs can take much longer and is of the order of days to even weeks.

There are a large number of bandwidth measurement/estimation tools available, with several new tools recently introduced. This in itself is an indication that accurate bandwidth measurement/estimation remains a hard problem even after many years of research and there is room for further improvements [9]. For the details of the various tools and their measurement accuracy, please use our bibliography or available survey articles [15, 18]. For the purposes of our study, our goal was to find a reasonably accurate but low overhead tool that is easily deployable on the PlanetLab platform. Note that the purpose of this study is *not* to do an accuracy comparison of these tools. After some narrowing of the choices, we evaluated the following tools as described below. We merely present our experiences with different tools in the evaluation process.

We were hesitant to use per-hop capacity estimation tools as they generate excessive probing traffic overhead. Moreover, we could not build pathchar [10] or pchar [14] as they can not be built on newer Linux systems. Currently, PlanetLab runs kernel version 2.4.22, but pathchar supports up to 2.0.30 and pchar up to 2.3. When we tested Clink [5] on PlanetLab, the experiment were simply “hung” without making any progress. We suspect this is also because of a Linux version compatibility issue.

As for end-to-end capacity estimation tools, bprobe [3] works only on SGI Irix. SProbe [21, 22] is an attractive tool as it only requires to be run on the source machine, and hence can measure capacities to hosts where the user does not have account access. In addition, SProbe is included in the Scriptroute [23] tool that runs as a service on PlanetLab hosts. One key feature of SProbe is that when it detects cross traffic, instead of making a poor estimate of the capacity, it does not report any value. When we ran SProbe between PlanetLab hosts, less than 30% of the measurements returned the capacity estimate. The authors of the tool had a similar experience on their trials with Napster and Gnutella peers. As we have access to all the PlanetLab hosts, we can deploy and run pathrate [4]. Unless the network hosts are down or we could not login to the hosts for various reasons, we were able to measure capacity between PlanetLab nodes using pathrate. It was the only capacity estimation tool we could successfully run and obtain estimates on PlanetLab.

We also tested several available bandwidth estimation tools. Similar to bprobe, cprobe [3] does not run on Linux. One of the most popular tools is pathload [8]. When we tested pathload on PlanetLab nodes however, we ran into an invalid argument error on `connect`. This very issue was also recently brought up in PlanetLab user mailing list. We were able to run IGI (Initial Gap Increasing) [7] without any run-time errors. However, the tool showed poor accuracy with high variance in the estimation of the same pair on sequential attempts, and also reported unusually high estimates (ten times larger than the estimated capacity by pathrate). Spruce [24] has shown to be more accurate than pathload and IGI. However, Spruce requires the knowledge of the capacity of the path to estimate available bandwidth. We also tested pathChirp [19] and it ran successfully with reasonable accuracy in our first set of measurements performed in August 2004. However, after the version 3 rollout of PlanetLab, pathChirp, along with STAB [20], developed by the same authors of pathChirp, failed to work on PlanetLab. After a few

Table 1. End-to-end capacity statistics.

	Set 1	Set 2
Number of nodes	279	178
Measurement period	8/11/04~9/6/04	1/5/05~1/18/05
PlanetLab version	version 2	version 3
Number of pairs	12,006	21,861
Minimum capacity	0.1 Mbps	0.3 Mbps
Maximum capacity	1210.1 Mbps	682.9 Mbps
Average capacity	63.44 Mbps	64.03 Mbps
Median capacity	24.5 Mbps	91.4 Mbps
Standard deviation	119.22 Mbps	43.78 Mbps

chirps, the tool stops running and hangs. We are communicating with the authors of pathChirp to resolve the issue.

In our future work, we are planning to test tools such as ABwE [16], CapProbe [11], pathneck [6], and MultiQ [12].

2.2 Measurement Analysis

For the first set of measurements, we show the capacity measurements from pathrate (version 2.4.0) as it was the only capacity estimation tool we were able to successfully run in a consistent manner. Each pathrate run on average took approximately 30 minutes. Pathrate returns two estimates, a high estimate and a low estimate, for bottleneck capacity between a pair of source and destination nodes. In the first experiment pathrate returned negative values for low capacity estimate in certain measurements. When we reported this to the authors of the pathrate tool, they kindly debugged the calculation error and the modified version (v2.4.1b) was used in the second set of measurements. To avoid this calculation error, we only report the high capacity estimate of the pathrate in this paper.

The first set of measurements was initiated on August 11th, 2004 and completed on September 6th, 2004. The second set was measured between January 5th, 2005 and January 18th, 2005. On our first attempt in August 2004, we tried measuring capacities between all PlanetLab nodes of the then nearly 400 nodes. However, many of the nodes did not respond consistently, and many of the pathrate capacity estimates were not returned. Ultimately, in the first set, we collected bottleneck capacity data on 12,006 network paths from 279 nodes. In the second set of experiments performed in January of 2005, we prefiltered 178 nodes (and no more than two nodes per site) that consistently responded. It could be one of the reasons why the experiments were finished in a shorter time compared with the first set of measurement experiments. In the second set of measurement we managed to collect data on 21,861 paths.

We first look at the statistics of the end-to-end capacity over all paths (source-destination node pairs) measured. It is important to note that given two nodes A and B, capacity measurements in both directions, i.e., source destination node pairs (A,B) and (B,A) may not both be available. Table 1 shows that the average bandwidth between

Table 2. End-to-end capacity distribution.

Capacity (C)	Set 1		Set 2	
	Number of paths	Percentage (%)	Number of paths	Percentage (%)
$C < 20$ Mbps	4013	33.42	6733	30.8
$20 \text{ Mbps} \leq C < 50$ Mbps	4246	35.37	1910	8.74
$50 \text{ Mbps} \leq C < 80$ Mbps	674	5.61	1303	5.96
$80 \text{ Mbps} \leq C < 120$ Mbps	2193	18.27	11744	53.72
$120 \text{ Mbps} \leq C < 200$ Mbps	207	1.72	139	0.64
$200 \text{ Mbps} \leq C < 500$ Mbps	392	3.27	21	0.096
$500 \text{ Mbps} \leq C$	281	2.34	11	0.05

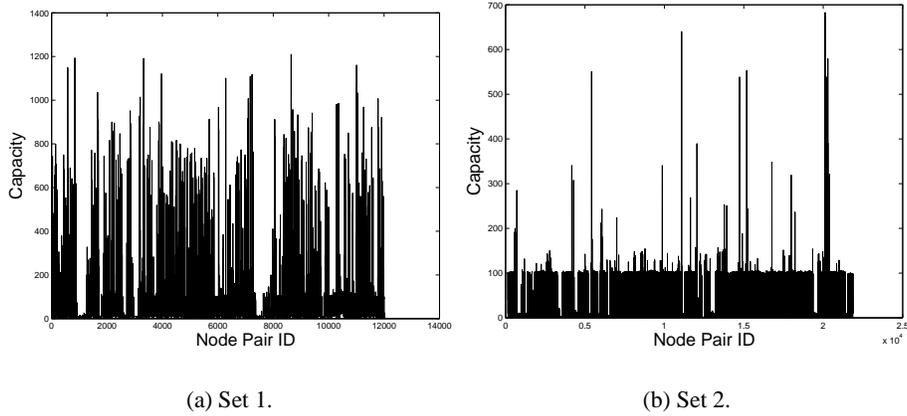


Fig. 1. Bandwidth capacity for all pairs measured.

PlanetLab hosts is nearly 64 Mbps. Table 2 shows the distribution, Figure 1 visualizes this distribution (notice the different scaling of y-axis between two subfigures) and Figure 2 shows the cumulative distribution function (notice the different scaling of x-axis between two subfigures). On further analysis, we observed that when certain nodes were the source or the destination, the bandwidth measured was very low. In the first set of measurements for instance, for paths with freedom.ri.uni-tuebingen.de as the source, the average bandwidth was 4.61 Mbps. We noticed a similar behavior in the second set as when 200-102-209-152.paemt7001.t.brasiltelecom.net.br was the source, the average capacity was 0.42 Mbps and when it was the destination, the average capacity was 0.41 Mbps. On the other hand, when planetlab1.eurecom.fr was the destination, the average bandwidth was 3.85 Mbps, with the path from planetlab1.inria.fr having 199.2 Mbps of bandwidth. Without this measurement of 199.2 Mbps, the average bandwidth with planetlab1.eurecom.fr as the destination is 2.13 Mbps. We also noticed nodes with high

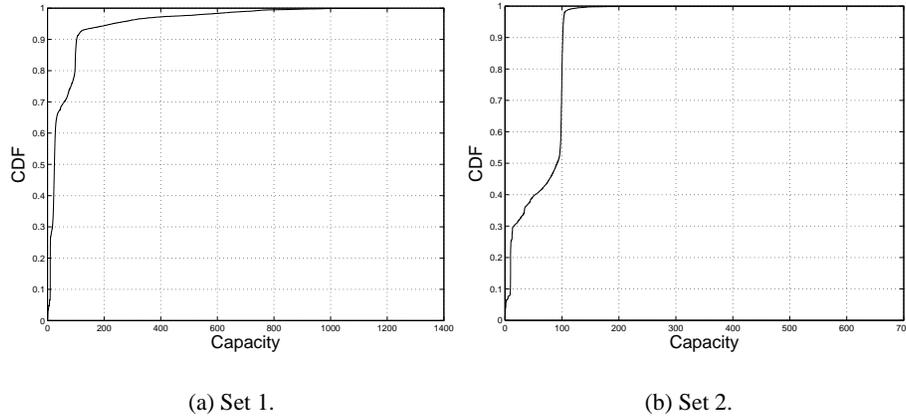


Fig. 2. Cumulative distribution function of bandwidth capacity.

Table 3. Capacity by regions (Mbps).

Source	Destination									
	North America		South America		Asia		Europe		Oceania	
	Set 1	Set 2	Set 1	Set 2	Set 1	Set 2	Set 1	Set 2	Set 1	Set 2
North America	60.67	66.28	8.34	0.41	55.74	60.8	71.78	68.11	N/A	79.64
South America	7.94	0.42	106	N/A	N/A	0.43	6.1	0.41	N/A	0.4
Asia	69.13	55.15	N/A	0.41	73.5	62.36	72.28	42.94	N/A	99.52
Europe	69.94	66.43	30.67	0.4	61.38	47.54	74.82	66.69	N/A	13.15
Oceania	N/A	37.25	N/A	0.4	N/A	22.54	N/A	7.03	N/A	50.9

average bandwidth. For instance, measurements from `planet1.ottawa.canet4.nodes.planetlab.org` showed an average of 508.46 Mbps.

Some PlanetLab nodes have imposed outgoing bandwidth limit, ranging from 500 Kbps to 10 Mbps. We observed interesting interplay between the traffic shaper for bandwidth limiting and the pathrate probing scheme. In some cases we measured end-to-end capacity of 100 Mbps even though the source was bandwidth limited to 500 Kbps. We are further exploring this interaction.

The standard deviation for the second set is much smaller than the first set. We believe the prefiltering of the nodes for the second set is the main reason as the nodes that showed extremely low or high capacities in the first set were relatively unstable, and could have been removed from our second experiments. In the second set, we have limited the number of nodes per site to at most two nodes, and hence we have less number of high capacity local paths than the first set. We can also observe that in the second set, more than 99% of the paths show the capacity of less than 120 Mbps.

Table 4. Asymmetry factor distribution.

Asymmetry factor (α)	Set 1		Set 2	
	Number of pairs	Percentage (%)	Number of pairs	Percentage (%)
$\alpha < 0.01$	132	6.08	1843	21.49
$0.01 < \alpha < 0.05$	395	18.19	3237	37.74
$0.05 \leq \alpha < 0.1$	165	7.6	817	9.52
$0.1 \leq \alpha < 0.2$	243	11.19	880	10.26
$0.2 \leq \alpha < 0.5$	328	15.1	1111	12.95
$0.5 \leq \alpha$	909	41.85	870	10.14

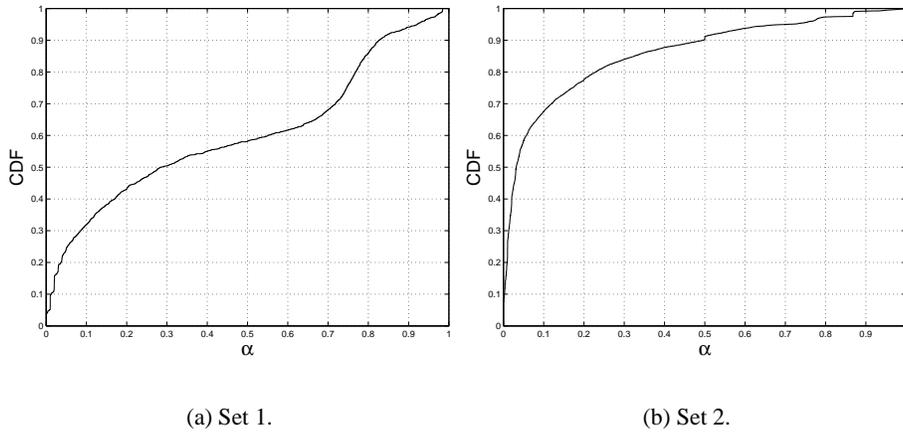
**Fig. 3.** Cumulative distribution function of asymmetry factor.

Table 3 shows the capacity measured region by region. We categorize each node into five regions: North America, South America, Asia, Europe, and Oceania. In our first measurement set we did not have any node from the Oceania region part of PlanetLab. On other entries of the table with N/A, no estimates were returned. There were only two nodes from Brasil in South American region in the second set, and as mentioned earlier, capacities of the path to and from these nodes were very low. One might think that pairs between the same region will have higher capacity than those for inter-regions. The table however does not show any strong confirmation of that belief. It is also interesting to see that although the paths from Asia to Oceania show high capacity, the same cannot be said for the reverse direction.

We now investigate whether the paths between PlanetLab hosts are symmetric in terms of capacity. For the first set, among 12,006 measurements, only 2,172 pairs (4,344 measurements) reported capacity estimates in both directions, and for the second set, 8,758 pairs (17,516 measurements) out of 21,861 measurements returned estimates for

Table 5. Comparison of End-to-end capacity statistics of common node pairs in Sets 1 and 2.

	Set 1	Set 2
Number of common pairs	3,409	
Measurement period	8/11/04~9/6/04	1/5/05~1/18/05
PlanetLab version	version 2	version 3
Minimum capacity	0.1 Mbps	0.5 Mbps
Maximum capacity	1014.1 Mbps	152.6 Mbps
Average capacity	55.79 Mbps	71.09 Mbps
Median capacity	24.3 Mbps	97.3 Mbps
Standard deviation	109.94 Mbps	39.32 Mbps

both directions. To understand the path asymmetry, we define asymmetry factor between two nodes i and j , $\alpha_{i,j}$, as follows:

$$\alpha_{i,j} = \frac{|BW_{i,j} - BW_{j,i}|}{\max(BW_{i,j}, BW_{j,i})}$$

where $BW_{i,j}$ is the bottleneck bandwidth from node i to node j .

When the capacity of the forward path equals the capacity of the reverse path (i.e., complete symmetry), the asymmetry factor is zero. A high asymmetry factor implies stronger asymmetry.

The distribution of the asymmetry factor for both sets are reported in Table 4 while the CDF is plotted in Figure 3. In the first set, only 132 pairs (6%) showed α of less than 0.01 and 692 pairs (32%) are less than 0.1. Although about 60% of the pairs have asymmetry factor below 0.5, there are a significant number of pairs with high asymmetry factor. We further investigated the reason for high asymmetry in 328 pairs that have α larger than 0.5. The PlanetLab imposed artificial bandwidth limit was reason for asymmetry in 189 of these pairs.

In the second set however, surprisingly large portion of the paths showed high symmetry. Nearly 60% has the asymmetry factor of less than 0.05. We believe that the main reason is, as shown in Table 2, more than half of the capacity estimates were between 80 and 120 Mbps.

Temporal Analysis Since the measurements from Set 1 and Set 2 were done almost 5 months apart, the obvious question to ask is whether the data suggests significant changes in the PlanetLab infrastructure during this period. Note that we already know of two significant changes - the PlanetLab software version was upgraded to version 3 and the pathrate tool was upgraded to a new version. To answer the above question, we computed the common source-destination node pairs between the two sets and analyzed the bandwidth measurements. We found 128 common nodes in the two sets and 3,409 common source-destination node pair measurements.

The summary statistics of the measured capacity for these node pairs common to both measurement sets are given in Table 5. There are some interesting differences between the two sets, which could be caused by infrastructure changes, measurement errors

Table 6. nd-to-end capacity distribution of common node pairs in Sets 1 and 2.

Capacity (C)	Set 1		Set 2	
	Number of paths	Percentage (%)	Number of paths	Percentage (%)
$C < 20$ Mbps	1041	30.54	909	26.66
$20 \text{ Mbps} \leq C < 50$ Mbps	1491	43.74	103	3.02
$50 \text{ Mbps} \leq C < 80$ Mbps	105	3.08	180	5.28
$80 \text{ Mbps} \leq C < 120$ Mbps	587	17.22	2205	64.68
$120 \text{ Mbps} \leq C < 200$ Mbps	37	1.09	12	0.35
$200 \text{ Mbps} \leq C < 500$ Mbps	86	2.52	0	0.00
$500 \text{ Mbps} \leq C$	62	1.82	0	0.00

or both. The average capacity between the measured node-pairs increased to 71 Mbps from 55 Mbps, as did the minimum measured bandwidth, implying an upgrade of the infrastructure on average. An interesting point to note is that the maximum capacity between any node pair decreased significantly from 1 Gbps to 152 Mbps. This could have been due to stricter bandwidth limits imposed on PlanetLab nodes. In the first set, the capacity between the nodes planetlab1.cse.nd.edu and planetlab2.cs.umd.edu were measured to be 1 Gbps, which in the second set is now close to 100 Mbps. We were unable to determine whether this is due to an infrastructure change, imposed bandwidth limit or measurement error. Diagnosing the causes for these measurement changes is future work.

While the stated goal of this work was not to verify accuracy of the pathrate tool (this has been done by other researchers in earlier work), we mentioned earlier that in some of the measurements in Set 1, the low estimate of bandwidth reported by pathrate were found to be negative and the authors of pathrate rectified this in the subsequent release. Although the negative values do not affect any of our presented results as we use the high estimate of the bandwidth, it is interesting to note that with the new version of pathrate, of the 3,409 measurements, no negative low estimates were observed in Set 2, while there were 93 negative measurements in Set 1.

The capacity distribution of the 3,409 common node pairs is given in Table 6. The biggest changes are in the paths with capacity between 20 Mbps and 50 Mbps and those with capacity between 80 Mbps and 120 Mbps. From the data presented it seems that significant number of paths were upgraded from the first band (20~50 Mbps) to the second band (80~120 Mbps) in the time between our measurements.

As mentioned earlier, given two nodes A and B in this common set, capacity measurements in both directions, i.e., source destination node pairs (A,B) and (B,A) may not both be available. Of the 3,409 source-destination node pairs common to sets 1 and 2, 661 node pairs (i.e., 1,322 measurements) had bandwidth measurements in both directions and hence the asymmetry metric could be computed for these. The asymmetry factor distribution is tabulated in Table 7. Again, the second set of experiments show a significantly reduced asymmetry than the first set.

Table 7. Asymmetry factor distribution of common node pairs in Sets 1 and 2.

Asymmetry factor (α)	Set 1		Set 2	
	Number of pairs	Percentage (%)	Number of pairs	Percentage (%)
$\alpha < 0.01$	65	9.83	145	21.94
$0.01 \leq \alpha < 0.05$	167	25.26	299	45.23
$0.05 \leq \alpha < 0.1$	57	8.62	86	13.01
$0.1 \leq \alpha < 0.2$	64	9.68	70	10.59
$0.2 \leq \alpha < 0.5$	83	12.56	48	7.26
$0.5 \leq \alpha$	225	34.04	13	1.97

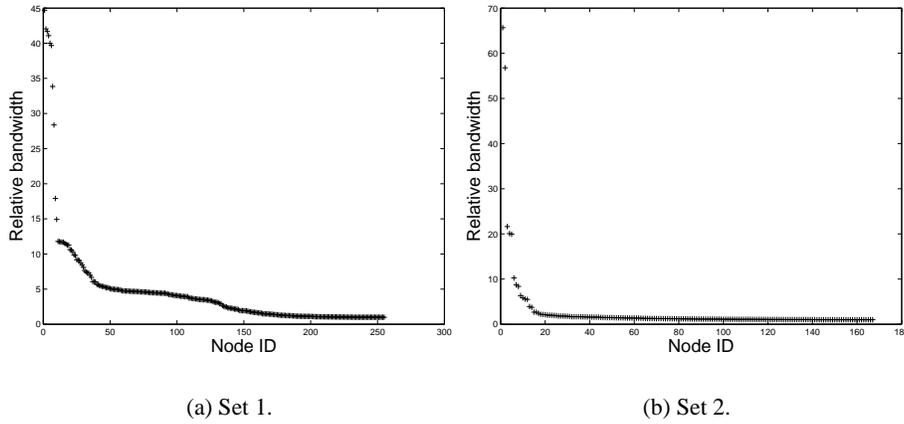


Fig. 4. Bandwidth/delay correlation.

Correlation Study We now study the correlation between bandwidth and latency. Before we report the result of this study, we explain the motivation of attempting to relate the delay with bandwidth. As mentioned in Section 1, our ultimate goal is to gain insights into potential correlation that will enable scalable bandwidth estimation. For example, to find a node whose path from a given node has the largest capacity, instead of performing bandwidth estimates to all the nodes, can we do the probing to just a small number of nodes (five for instance)? Since measuring latency can be done with less probing overhead with quick turnaround time than measuring bandwidth, there already exist tools that perform scalable network distance estimation [26]. With these tools available and latency values easily in hand, if there is any relationship or trend between latency and bandwidth, we can scalably estimate network bandwidth without excessive bandwidth probing. That is the main motivation of this trend analysis. Note also that using capacity, instead of available bandwidth, is more appropriate as the values of available bandwidth vary with time, and unless the measurement of bandwidth and latency are done at the same period, the analysis could be meaningless.

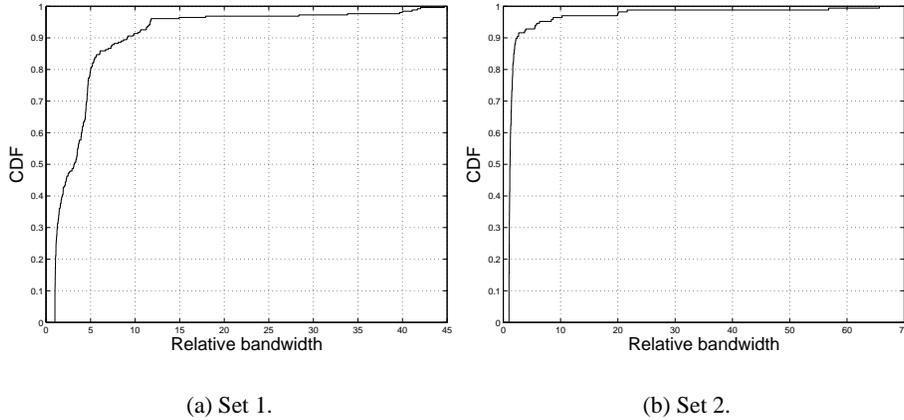


Fig. 5. Cumulative distribution function of bandwidth/delay correlation.

For the latency measurement, we initially used the all pair ping data.⁴ However, due to some missing ping data, there was little overlap between the available ping data and collected bandwidth measurement. Hence we also used the RTT measurements from pathrate. The resulting trends from ping and RTT measures from pathrate are quite similar, and hence we only present the results based upon the pathrate RTT latency.

We used two metrics for studying the bandwidth and latency correlation. The first metric, called relative bandwidth correlation metric, captures the ratio of maximum bandwidth path and bandwidth to the closest node. For a given host ($node_i$), using the latency data, we find the host that has the minimum latency ($node_{minLat}$). Similarly, using the capacity measurements, we find the host that provides the maximum bandwidth ($node_{maxBW}$). The relative bandwidth correlation penalty metric for $node_i$ is then defined as the ratio of the maximum bandwidth ($BW_{i,maxBW}$) and the bandwidth from $node_i$ to $node_{minLat}$ ($BW_{i,minLat}$). This metric takes values greater than or equal to 1.0. The closer this metric is to 1.0, we consider the correlation between latency and bandwidth to be stronger. We plot this metric and its cumulative distribution in Figures 4 and 5.

We see from the figure that latency and bandwidth are surprisingly quite correlated, especially in the second set. On further analysis with the CDF for the first set, we notice that for about 40% of the nodes, the bandwidth to the closest node is roughly 40% smaller than the actual maximum bandwidth. Our preliminary investigations reveal one of the primary reasons for this behavior is the imposed bandwidth limit. In some cases, the capacities to even the nearby nodes are quite less than the maximum bandwidth as they have bandwidth limits. For instance, in the Set 1, in lots of cases the node with highest capacity is planetlab1.ls.fi.upm.es which did not have any imposed bandwidth limit.

⁴ We obtain this from http://www.pdos.lcs.mit.edu/~strib/pl_app/.

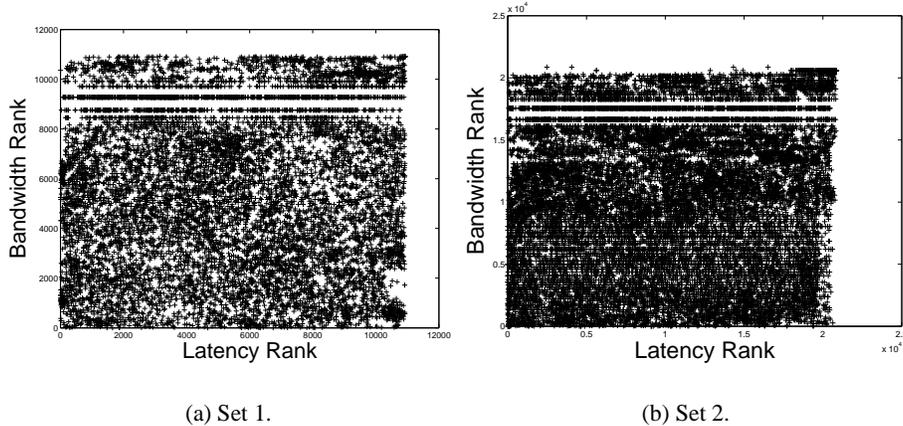


Fig. 6. Bandwidth/delay correlation using rank correlation.

For the second set on the other hand, we see that about 90% of the closest nodes has the capacity that is nearly equal to the maximum bandwidth. We must remember however, that the roundtrip time measurement was performed by pathrate itself. Tools such as Netvigator [26] that perform network distance estimation typically uses traceroute or ping, and those tools may return different values. We can further test the correlation between bandwidth and latency using these tools.

We also used the Spearman rank correlation coefficient [13], which is commonly used instead of the Pearson’s correlation coefficient, when the two sets of data come from different distributions. In Figure 6, we rank all the node pairs based on both bandwidth and latency. We plot the bandwidth rank versus the latency rank. If there is good correlation, we expect the points to be clustered along the $y = x$ diagonal line. However, using the Spearman coefficient we did not find any such correlation. The value of Spearman coefficient for sets 1 and 2 are 0.027 and 0.138 respectively based on the data presented in these figures. We believe that the degree of rank-order correlation is higher between closeby nodes and it decreases as distance between nodes increases. Hence, the low value of Spearman coefficient might be due to its computation over all the possible nodes. We plan to evaluate other rank-order correlation coefficients in this context.

3 Summary and Conclusions

In this paper, we presented a large scale measurement study of end-to-end capacity of paths between PlanetLab nodes. Our contributions are two-fold in that in addition to presenting the analysis of the data we collected from two sets of experiments, we also described the issues with the deployment of a number of bandwidth measurement tools on the PlanetLab platform. The measurement work is ongoing, but even with the data we have collected so far, there are a number of interesting conclusions one can draw.

Foremost, we verified our intuition that network paths connecting PlanetLab nodes are highly heterogeneous in the capacity values and in planning PlanetLab experiments, one needs to take this into account. The capacity of paths may have an order of magnitude difference even when they are sourced from the same node and similarly for the same receiver. Paths between two nodes do not necessarily show capacity symmetry. According to PlanetLab policy, bandwidth limits on outgoing traffic have been implemented. But we observed violations of the policy, which could have been due to the inaccuracy of the tool itself and we are investigating this further. In attempting to draw insights for scalable bandwidth estimation, we studied the correlation between latency and bandwidth of a path. Our preliminary results are promising and we plan to investigate this further.

One of our future work includes modifying the SProbe tool to keep attempting measurements until a valid estimate is made. We can then measure the bandwidth to nodes outside PlanetLab as SProbe does not require user access to destination hosts. We also plan to periodically measure the all pair bandwidth between planet-lab hosts and make it available to public. Although there exists a running service that measures bandwidth between PlanetLab nodes,⁵ it uses Iperf [25] that measures achievable TCP throughput, which is not necessarily raw capacity or available bandwidth.

References

1. A. Akella, S. Seshan, and A. Shaikh, "An empirical evaluation of wide-area Internet bottlenecks," in *Proceedings of the ACM IMC 2003*, Miami, FL, October 2003.
2. S. Banerjee, M. Pias, and T. Griffin, "The interdomain connectivity of planetlab nodes," in *Proceedings of the PAM 2004*, Sophia-Antipolis, France, April 2004.
3. R. L. Carter and M. E. Crovella, "Server selection using dynamic path characterization in wide-area networks," in *Proceedings of the IEEE INFOCOM'97*, Kobe, Japan, April 1997.
4. C. Dovrolis, P. Ramanathan, and D. Moore, "What do packet dispersion techniques measure?" in *Proceedings of the IEEE INFOCOM 2001*, Anchorage, AK, April 2001.
5. A. B. Downey, "Using pathchar to estimate Internet link characteristics," in *Proceedings of the ACM SIGCOMM'99*, Cambridge, MA, September 1999.
6. N. Hu, L. Li, Z. M. Mao, P. Steenkiste, and J. Wang, "Locating internet bottlenecks: Algorithms, measurements, and implications," in *Proceedings of the ACM SIGCOMM 2004*, Portland, OR, August 2004.
7. N. Hu and P. Steenkiste, "Evaluation and characterization of available bandwidth probing techniques," *IEEE J. Select. Areas Commun.*, vol. 21, no. 6, August 2003.
8. M. Jain and C. Dovrolis, "Pathload: A measurement tool for end-to-end available bandwidth," in *Proceedings of the PAM 2002*, Fort Collins, CO, March 2002.
9. ———, "Ten fallacies and pitfalls on end-to-end available bandwidth estimation," in *Proceedings of the ACM IMC 2004*, Taormina, Italy, October 2004.
10. V. Jacobson. pathchar: A tool to infer characteristics of internet paths. [Online]. Available: <ftp://ftp.ee.lbl.gov/pathchar>
11. R. Kapoor, L.-J. Chen, L. Lao, M. Gerla, and M. Y. Sanadidi, "CapProbe: A simple and accurate capacity estimation technique," in *Proceedings of the ACM SIGCOMM 2004*, Portland, OR, August 2004.
12. S. Katti, D. Kitabi, C. Blake, E. Kohler, and J. Strauss, "MultiQ: Automated detection of multiple bottleneck capacities along a path," in *Proceedings of the ACM IMC 2004*, Taormina, Italy, October 2004.

⁵ <http://www.planet-lab.org/logs/iperf/>

13. E. L. Lehmann and H. J. M. D' Abrera, *Nonparametrics: Statistical Methods Based on Ranks*, rev. ed. Prentice-Hall, 1998.
14. B. A. Mah. pchar: A tool for measuring internet path characteristics. [Online]. Available: <http://www.kitchenlab.org/www/bmah/Software/pchar>
15. F. Montesino-Pouzols, "Comparative analysis of active bandwidth estimation tools," in *Proceedings of the PAM 2004*, Sophia-Antipolis, France, April 2004.
16. J. Navratil and R. L. Cottrell, "ABwE: A practical approach to available bandwidth estimation," in *Proceedings of the PAM 2003*, La Jolla, CA, April 2003.
17. T. S. E. Ng and H. Zhang, "Predicting Internet network distance with coordinates-based approaches," in *Proceedings of the IEEE INFOCOM 2002*, New York, NY, June 2002.
18. R. Prasad, C. Dovrolis, M. Murray, and kc claffy, "Bandwidth estimation: Metrics, measurement techniques, and tools," *IEEE Network*, vol. 17, no. 6, pp. 27–35, November/December 2003.
19. V. Ribeiro, R. Riedi, R. Baraniuk, J. Navratil, and L. Cottrell, "pathChirp: Efficient available bandwidth estimation for network paths," in *Proceedings of the PAM 2003*, La Jolla, CA, April 2003.
20. V. J. Ribeiro, R. H. Riedi, and R. G. Baraniuk, "Locating available bandwidth bottlenecks," *IEEE Internet Comput.*, vol. 8, no. 6, pp. 34–41, September/October 2004.
21. S. Saroiu. SProbe: A fast tool for measuring bottleneck bandwidth in uncooperative environments. [Online]. Available: <http://sprobe.cs.washington.edu/>
22. S. Saroiu, P. K. Gummadi, and S. D. Gribble, "SProbe: Another tool for measuring bottleneck bandwidth," in *Work-in-Progress Report at the USITS 2001*, San Francisco, CA, March 2001.
23. N. Spring, D. Wetherall, and T. Anderson. Scriptroute: A facility for distributed internet debugging and measurement. [Online]. Available: <http://www.scriptroute.org>
24. J. Strauss, D. Katabi, and F. Kaashoek, "A measurement study of available bandwidth estimation tools," in *Proceedings of the ACM IMC 2003*, Miami, FL, October 2003.
25. A. Tirumala, F. Qin, J. Dugan, J. Ferguson, and K. Gibbs. Iperf: The TCP/UDP bandwidth measurement tool. [Online]. Available: <http://dast.nlanr.net/Projects/Iperf/>
26. Z. Xu, P. Sharma, S.-J. Lee, and S. Banerjee, "Netvigator: Scalable network proximity estimation," HP Laboratories, Tech. Rep., 2005.