



NORTHWESTERN UNIVERSITY

Electrical Engineering and Computer Science Department

Technical Report
NWU-EECS-09-22
December 18, 2009

Pitfalls for Testbed Evaluations of Internet Systems

David Choffnes and Fabian Bustamante

Abstract

The design, implementation and evaluation of large-scale systems relies on global views on network topologies and performance. In this article, we identify several issues with extending results from limited platforms to *Internet wide* perspectives. Specifically, we try to quantify the level of inaccuracy and incompleteness of testbed results when applied to the context of a large-scale peer-to-peer (P2P) system. Based on our results, we emphasize the importance of measurements in the appropriate environment when evaluating Internet-scale systems.

Keywords: Internet-scale systems, peer-to-peer, evaluation.

Pitfalls for Testbed Evaluations of Internet Systems

David Choffnes and Fabián E. Bustamante

Department of Electrical Engineering & Computer Science, Northwestern University
{drchoffnes,fabianb}@eecs.northwestern.edu

ABSTRACT

The design, implementation and evaluation of large-scale systems relies on global views on network topologies and performance. In this article, we identify several issues with extending results from limited platforms to *Internet wide* perspectives. Specifically, we try to quantify the level of inaccuracy and incompleteness of testbed results when applied to the context of a large-scale peer-to-peer (P2P) system. Based on our results, we emphasize the importance of measurements in the appropriate environment when evaluating Internet-scale systems.

1. INTRODUCTION

Global views of network topologies and performance are essential for informing protocol design and building systems and services at Internet-scale. A number of research efforts [9, 11] have focused on acquiring such views from research platforms such as PlanetLab [13] and using public topology data such as RouteViews [15] in an attempt to build a comprehensive network atlas. Such information has been successfully used to advise a variety of important applications addressing IP reachability, prefix hijacking and routing anomalies.

Recent studies, however, suggest that such testbed results for distributed systems do not always extend to the targeted deployment. For example, Ledlie et al [7] and Agarwal et al. [1] show that network positioning systems perform much worse “in the wild” than in PlanetLab deployments.

In this article, we identify several limitations with this model taking a look at the accuracy and completeness of testbed results when applied to the context of a large-scale peer-to-peer (P2P) system. To inform this study, we use a unique collection of traces gathered from a deployment containing hundreds of thousands of users located at the network edge.

We focus our analysis on the following three issues that affect the validity of any study extending testbed results to an Internet scale. First, we find that large and significant portions of the Internet topology used by P2P systems are invisible to research testbeds, limiting the effectiveness of testbed-based inference of Internet paths and relationships between autonomous systems (ASes). Next, we show that inferred properties of these topologies (latencies and bandwidth capacities) are inaccurate. Finally, we discuss how these issues prevent accurate evaluations of performance for distributed systems running on these topologies.

Based on our results, we emphasize the importance of representative network views when evaluating Internet-scale systems. There is a number of approaches to achieve this goal, including the use of application-level and network-level traces from the edge of the network (e.g., via Ark¹ and Ono datasets). We also encourage the design and deployment of additional edge-based monitoring services, either built into existing distributed services or provided independently with the proper incentives for Internet-scale adoption.

2. EDGE SYSTEM TRACES

The Internet is growing in ways that make increasingly difficult to attain a global view of the network. Large swathes of the network cannot be probed directly from our research testbeds and a number of valuable measurement techniques have side effects that render them impractical.

For our study, we address this issue using network- and application-level traces from tens of thousands of BitTorrent users worldwide running the Ono plugin [3]. In particular, our software passively records the data transferred over each host’s connection and performs active ping and traceroute measurements to a subset of these connections.

Our installed user base covers 204 countries, 53,000 routable prefixes and more than 7,000 ASes. While the amount of data collected per unit time varies according to the online user population, each day we record between 2.5 and 3.5 million traceroutes, tens of millions of latency measurements and more than 100 million per-connection transfer-rate samples.

The design, implementation and maintenance of the associated data-collection service is the subject of an ongoing systems project. As part of this work, we are making this dataset available to researchers through our EdgeScope project.²

In the following sections we use this dataset to explore several key pitfalls of testbed-based evaluations for Internet scale systems. We begin by exploring the completeness of the view from such testbeds.

3. GENERALIZING NETWORK VIEWS

A number of studies explicitly or implicitly rely on network topologies for estimating Internet resiliency, inferring

¹<http://www.caida.org/projects/ark/>

²<http://www.aqualab.cs.northwestern.edu/projects/EdgeScope.html>

Internet paths and estimating cross-network traffic costs, to name a few. While several research efforts have successfully extracted detailed topologies from public vantage points, it is well known that these are incomplete Internet views.

To explore a lower bound for missing topology information, we used AS-level paths and AS relationships inferred from traceroute data gathered from hundreds of thousands of users located at the edge of the network [2]. As Fig. 1 shows, the number of vantage points available from edge systems far outnumbers those from public views, particularly for lower tiers of the Internet hierarchy where most of the ASes reside.

Our dataset includes probes from nearly 1 million source IP addresses to more than 84 million unique destination IP addresses, all of which represent active users of the BitTorrent P2P system. By comparison, the BitProbes study [6] used a few hundred sources from the PlanetLab testbed to measure P2P hosts comprising 500,000 destination IPs.

Missing links. Besides providing a large number of vantage points, our dataset also discovers links invisible to the public view. In total, our approach identified 20% additional links missing from the public view, the vast majority of which were located below Tier-1 ASes in the Internet hierarchy. Not surprisingly, the number of missing links discovered increases with the tier number location for these links. Thus, when evaluating the interaction between network topologies and Internet systems at the edge (often located in lower Internet tiers), testbed-based topologies are less likely to include many relevant links.

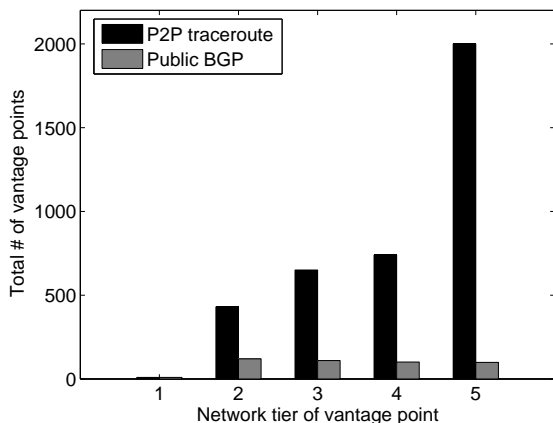


Figure 1: Number of vantage point ASes corresponding to inferred network tiers.

In addition to the locations of links in the Internet hierarchy, it is useful to understand what kinds of AS relationships are included in these missing links. Table 1 categorizes links into Tier-1, customer-provider or peering links and shows the missing links as a fraction of existing links in the public view, for each category. Note that there is a large number of additional peering links (44%) and, more surprisingly, a significant fraction of new customer-provider links (12%).

Tier-1	Customer-Provider	Peering
3.14%	12.86%	40.99%

Table 1: Percent of links missing from public views, but found from edge systems, for major categories of AS relationships.

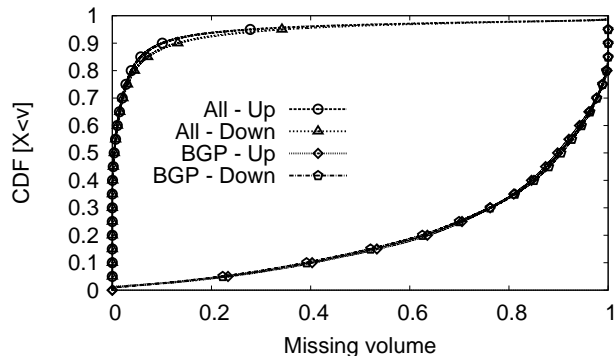


Figure 2: CDF of the portion of each host's traffic volume that could not be mapped to a path based on both public views and traceroutes between a subset of P2P users.

From these results it is clear that our community may need to revisit analysis that is based on public-view topologies (e.g., when looking at traffic cost or Internet resilience).

Impact of missing paths. To better understand how this missing information affects studies of Internet-scale systems, we investigate the impact of missing links using three weeks of connection data from P2P users. In particular, we would like to know how much of these users' traffic volumes can be mapped to AS-level paths – an essential step for evaluating P2P traffic costs and locality.

We begin by determining the volume of data transferred over each connection for each host, then we map each connection to a source/destination AS pair using the Team Cymru service [18]. We use the set of paths from public views and P2P traceroutes [2] and, finally, for each host we determine the portion of its traffic volume that could not be mapped to *any* AS path in our dataset.

Figure 2 uses a cumulative distribution function (CDF) to plot these unmapped traffic volumes using only BGP data (labeled *BGP*) and the entire dataset (labeled *All*). The figure shows that when using *All* path information, we cannot locate complete path information for 84% of hosts; fortunately, the median portion of traffic for which we cannot locate an AS path is only 6.7%. Of the hosts in our dataset 16% use connections for which we have path information for only half of their traffic volumes and 3% use connections for which we have no path information at all. When using only BGP information the median volume of unaccounted traffic is nearly 90%.

One implication of this result is that any Internet wide study from a testbed environment cannot accurately characterize path properties for traffic volumes from the majority of P2P users. Even though the additional links discovered

by Chen et al. [2] cannot identify every single paths carrying traffic volumes for P2P systems, a partial view of these paths based on traceroutes to randomly selected peers allows us to map nearly all of the flows.

4. GENERALIZING MEASUREMENTS

While the previous section showed that large portions of the network are invisible to current testbeds, in this section we show how properties of topology links measured from testbed vantage points do not extend to those measured from the edge of the network. We begin by focusing on estimating distances between Internet hosts, which is essential to a variety of network performance optimizations including server selection, central leader election and connection biasing. We close the section by examining Internet-wide bandwidth capacities as measured by BitTorrent throughput from users at the edge of the network.

There is a large body of research addressing the issue of how to measure, calculate and encode Internet distances in terms of round-trip latencies [5, 10, 16, 17]. Generally, these solutions rely on methods to predict latencies between arbitrary hosts without requiring the N^2 number of measurements that provide ground-truth information. Previous work has identified the following key properties that impact network positioning performance: the structure of the latency space, the rate of triangle-inequality violations (TIVs) in this latency space and last-mile delays. We now show how these key properties are significantly different when measured exclusively from edge systems compared to those measured from testbed environments.

We base our results on 2 billion latency samples gathered from edge systems during June 10–25th, 2008.³ Unlike studies that use PlanetLab hosts to measure latencies or infer them based on latencies between DNS servers, this dataset consists exclusively of directly measured latencies between edge systems. It is also an order of magnitude larger than the set used by Agarwal et al. [1] to evaluate server selection in gaming systems.

Latencies. To begin, Figure 3 compares the average latencies seen by hosts using the Ono plugin (labeled P2P) to those seen from three related projects: the RON testbed (MIT), PlanetLab (PL) and Ledlie et al.’s study (PL-to-P2P). The graph shows that latencies from edge systems are generally much larger than those from MIT King [5] and PlanetLab (PL). In fact, the median latency in our dataset is twice as large as reported by the study by Ledlie et al. [7], which used PlanetLab nodes to probe Vuze P2P users (PL-to-P2P).⁴

Triangle-Inequality Violations. TIVs in the Internet delay space occur when the latency between hosts A and




Figure 3: CDFs of latencies from different measurement platforms (semilog scale). Our measurement study exclusively between peers in Vuze (labeled P2P) exhibits double the median latency “in the wild” (labeled PL-to-P2P).

B is larger than the sum of the latency from A to C and C to B ($A \neq B \neq C$). This is caused by factors such as network topology and routing policies (see, for example, [8, 17]). Wang et al. [19] demonstrate that TIVs can significantly reduce the accuracy of network positioning systems.

We performed a TIV analysis on our dataset and found that over 13% of the triangles had TIVs (affecting over 99.5% of the source/destination pairs). Lumezanu et al. [8] study the dynamics of TIVs and demonstrate that using the minimum RTTs, as done in this study, is likely to underestimate the rate of TIVs. Thus our results can be considered a lower bound for TIVs in a large-scale P2P environment.

Compared to TIV rates reported in an analysis of datasets from Tang and Crovella [17], TIV rates in the P2P environment we studied are between 100% and 400% higher, and the number of source/destination pairs experiencing TIVs in our dataset (nearly 100%) is significantly greater than the 83% reported by Ledlie et al. [7]. These patterns for TIVs and their severity hints at the challenges in accounting for TIVs in coordinate systems.

Last-mile effects. It is well known that last-mile links often have poorer quality than the well provisioned links in transit networks. The problem is particularly acute in typical network edge settings. However, most of today’s network positioning systems either ignore or naively account for this effect.

We analyze last-mile effects by dividing the traceroute-based IP-level path between hosts into quartiles and determining the portion of the end-to-end latency contained in each quartile. If the latency were evenly distributed among IP hops along a path, each quartile would contain 25% of the end-to-end latency. In contrast, the first quartile (which is very likely to contain the entire first mile) accounts for disproportionately large fractions of the total end-to-end latency. For instance, when looking at the median values, the first quartile alone captures 80% of the end-to-

³For more details about this dataset and our methodology, see [4].

⁴We found that P2P traffic did not significantly impact latencies; when our measurement hosts were not transferring data their latencies were smaller than those in the complete dataset, but the difference in median latencies was less than 10%.

end latency. The middle two quartiles, in contrast, each account for only 8%. Also note that the first quartile (and a significant fraction of the last quartile) has a large number of values close to and larger than 1. This demonstrates the variance in latencies along these first and last miles, where measurements to individual hops along the path can yield latencies that are close to or larger than the total end-to-end latency (as measured by probes to the last hop). In fact, more than 10% of the first quartile samples have a ratio greater than 1. While Vivaldi uses “height” to account for (first- and) last-mile links [5], this analysis suggests that a single parameter is insufficient due to the large and variable latencies in a large-scale P2P environment.

Bandwidth capacities. Bandwidth capacities are an important factor in the design of distributed systems, from making encoding decisions in video streaming to informing peer selection in P2P systems. While there are many proposed techniques for estimating capacities, these techniques are not amenable to widespread studies due to limitations on measurement traffic volumes and the need for compliant endpoints. Further, previous work has cast doubts on their accuracy [14].

We take an alternative approach to estimating capacities based on passively monitoring BitTorrent throughput. BitTorrent generally attempts to maximize bandwidth capacities to minimize time to completion for downloading files, so we expect observed throughputs to be proportional to a host’s bandwidth capacity. This approach alleviates the issues of compliant endpoints and generating measurement traffic; however, this environment can be affected by ISP interference (e.g., traffic shaping) and user-specified limits on the maximum throughput consumed by a P2P application. While accounting for ISP interference is the topic of ongoing work, we have the necessary data to account for user-specified limits and filter out these cases.

After this filtering step, we use the maximum upstream and downstream transfer rates seen by each host during a three-week period in April, 2009. As such, our results are a lower bound for each host’s bandwidth capacity.

Figure 4 depicts a CDF of maximum upstream and downstream throughput seen for each host in our study. First, we note the rarity of step-like functions in the CDFs, which would occur if BitTorrent were, as commonly believed, most often saturating the full bandwidth capacity. Thus, while BitTorrent attempts to saturate each user’s downstream bandwidth capacity, in practice it does not always do so.

We also find that the median upstream rate is 54 KB/s while the median for downstream rates is 102 KB/s. Interestingly, this indicates that although asymmetric bandwidth allocation – often with about an order of magnitude larger downstream rates – is common in the Internet, the transfer rates achieved by P2P systems are indeed limited by the peers’ upstream capacities.

It is important to note that these CDFs do not imply that the ratio of upstream to downstream capacities is nearly




Figure 4: CDF of transfer rates for all users, where the median is only 50 to 100 KB/s. This suggests that the BitTorrent system is dominated by mid-to-low-capacity hosts.

one for most hosts – some of the above samples contain only upstream transfer rates. For those hosts where we can measure both upstream and downstream throughputs, we find that the median ratio is 0.32 and the 90th percentile ratio is 0.77. This is in line with the asymmetric bandwidth allocations typical of DSL and cable Internet technologies being used by the majority of our vantage points.

We now compare these lower-bound estimates of capacities with those measured from PlanetLab in 2006 as reported by Isdal et al. [6]. We expect that bandwidth capacities have increased since then, so BitTorrent-based capacities should be larger. The authors report that 70% of hosts have an upload capacity between 350Kbps and 1Mbps. However, we surprisingly find that only 45% of hosts in our study achieve such transfer rates. In fact, 40% of hosts in our study achieve *less than* 350 Kbps maximum upstream rates. This suggests that even if the testbed-based bandwidth capacity measurements were accurate, they are insufficient for predicting *achieved* transfer rates in a P2P system. Although Isdal et al. were unable to directly measure or estimate downstream rates at the edge of the network, Fig. 4 shows that they closely track upstream rates until after the 30th percentile, where downstream rates significantly exceed upstream ones.

Finally, we analyze the maximum throughput achieved by hosts grouped by country in Fig. 5. We find that hosts in Germany, Romania and Sweden achieve the highest transfer rates while those in India, the Philippines and Brazil achieve the lowest. This is in line with results from independent bandwidth tests from Speedtest.net, indicating that maximum transfer rates measured from P2P users, when grouped by location, are in fact predictive of the bandwidth capacity *rankings*.

5. INFERRING SYSTEM PERFORMANCE

While measurements from the edge of the network help us better understand network topologies, delay behavior and bandwidth capacity distributions, they also are essential to

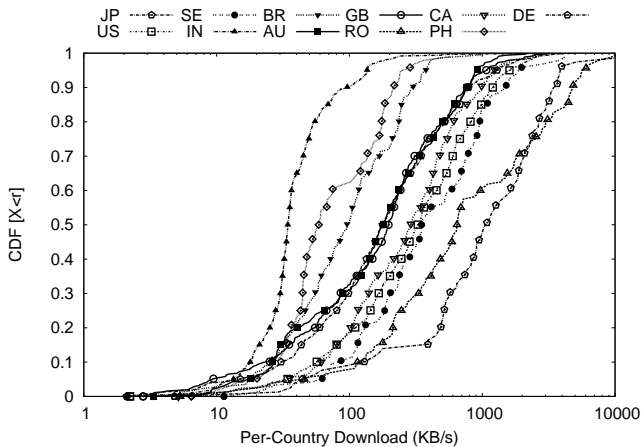


Figure 5: Per-country throughput CDFs, showing that Germany, Romania and Sweden have the highest average capacities while India, the Philippines and Brazil have the lowest.

designing, evaluating and optimizing distributed systems that run in this environment. We now show how more accurate views of the edge of the network affect system performance when compared to evaluations conducted from testbed environments.

Network positioning. We begin with network positioning systems and determine how the latency space measured in the previous section affects accuracy for a variety of positioning systems including GNP [10], Vivaldi [5], Meridian [20] and CRP [16].

The Vivaldi and CRP systems are implemented in our measurement platform, so their values represent true, “in the wild” performance. For evaluating GNP performance, we use the authors’ simulation implementation. The results are based on three runs of the simulation, each using a randomly chosen set of 15 landmarks, 464 targets and an 8-dimensional coordinate space. We also simulate Meridian using settings proportional to those in the original evaluation, with 379 randomly selected Meridian nodes, 100 target nodes, 16 nodes per ring and 9 rings per node. Our results are based on four simulation runs, each of which performs 25,000 latency queries.

We begin our analysis by evaluating the accuracy of GNP and of the Vuze Vivaldi implementations in terms of errors in predicted latency. Meridian and CRP are omitted here because they do not provide quantitative latency predictions. Figure 6(a) presents the cumulative distribution function (CDF) of errors on a semilog scale, where each point represents the absolute value of the *average* error from one measurement host. We find that GNP has lower measurement error (median is 59.8 ms) than the original Vivaldi implementation (labeled V1, median error is \approx 150 ms), partially due to GNP’s use of fixed, dedicated landmarks. Somewhat surprisingly, Ledlie et al.’s Vivaldi implementation (labeled V2) has slightly larger errors in

latency (median error is \approx 165 ms) than GNP and V1; however, we show in the next paragraph that its relative error is in fact smaller.

Relative error, the difference between the expected and measured latency, is a better measure of accuracy for network positioning systems. To compute relative errors, we first calculate the absolute value of the relative error between Vivaldi’s estimated latency and the ping latency for each sample, then find the average of these errors for each client running our software. Fig. 6(b) plots a CDF of these values; each point represents the average relative error for a particular client. For Vivaldi V1, the median relative error for each node is approximately 74%, whereas the same for V2 is 55% – both significantly higher than the 26% median relative error reported in studies based on PlanetLab nodes [7]. Interestingly, the median error for Vivaldi V2 is approximately the same as for GNP, indicating that decentralized coordinates do not significantly hurt relative performance. Finally, because Meridian and CRP do not predict distances, Fig. 6(b) plots the relative error for the closest peers they found. Meridian finds the closest peer approximately 20% of the time while CRP can locate the closest peer more than 70% of the time.

Network costs in P2P file sharing. Large traffic volumes generated by P2P file-sharing systems have generated a great deal of publicity as network providers attempt to reduce their costs by blocking, shaping or otherwise interfering with P2P connections. Given the popularity of these systems, a number of research efforts have investigated this issue by designing systems to reduce cross-network traffic [3,21] and evaluate the potential for P2P traffic locality [12].

Most previous work in this area relies on limited deployments and/or simulation results to estimate network costs of P2P systems. We now show how measurements from a large-scale, live deployment – combined with more complete AS topology information – provides a different view of the costs incurred by these systems.

A number of studies estimate the costs of P2P traffic as proportional to the number of AS hops along paths to different hosts. In this context, traffic is considered “no-cost” (also referred to as “local”) if it stays entirely in the same AS. We now refine this metric to include all paths for which no hop contains a customer-provider relationship (or vice-versa); i.e. our definition of “no-cost” includes traffic that remains in the origin AS or traverses peering and sibling links.

Figure 7(a) presents a CDF of the portion of each P2P user’s traffic that is “no-cost.” Our results from 130,000 source IPs and 12 million destination IPs indicates that the vast majority of hosts naturally generate at least some no-cost traffic. This result contradicts those from Piatek et al. [12], who use inferred testbed-based results and a single deployed vantage point to question the effectiveness of reducing ISP costs in P2P systems. In fact, we find that the majority of traffic volumes are no-cost for a significant

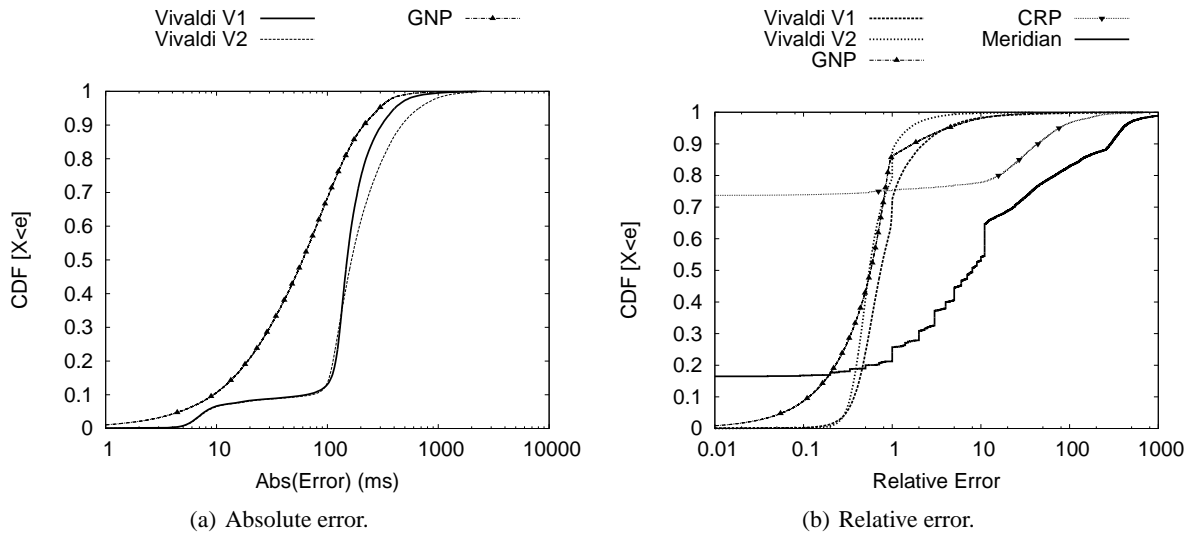


Figure 6: Absolute value of errors between estimated and measured latencies, in milliseconds (right), and absolute value of *relative* errors between estimated and measured latencies (left).

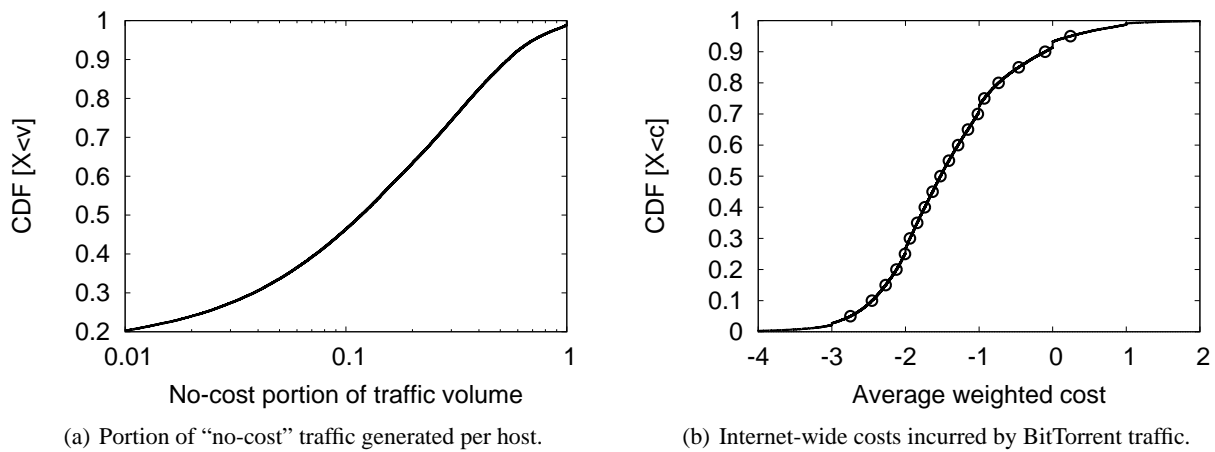


Figure 7: CDF of portion of "no-cost" traffic generated per host (right), and estimated Internet-wide costs incurred by BitTorrent traffic (left). The vast majority of hosts generate at least some no-cost traffic while the majority of traffic volumes are no-cost for 12% of hosts. Further, our results show that P2P traffic has a net effect of generating significant revenue for provider ISPs.

fraction (12%) of hosts.

It is important to note that these results may be biased by the fact that the measured hosts are using Ono to preferentially use no-cost peer connections. We do not believe that this effect strongly affects our results: while Ono has been installed nearly 800,000 times, it is still a small minority (less than half of one percent) of the total number of BitTorrent users worldwide. Thus, these results represent neither an upper nor a lower bound for the portion of no-cost traffic that P2P systems can successfully use.

Finally, Fig. 7(b) plots the average cost per byte for each user, based on the net costs of P2P traffic according to the traffic volumes per path and AS relationships along each path. Specifically, the cost of a path is the sum of the cost of each AS hop, where a hop between customer and provider is assigned a cost of 1, provider to customer a cost of -1 and zero otherwise (sibling and peer AS hops). We then determine, for each host, the portion of all traffic volume generated by each of its connections and multiply this by the cost of the path. Each point in Fig. 7(b) represents the sum of these values for each host.

As the figure shows, the vast majority of hosts generate flows with a net effect of generating revenue (i.e., negative costs) for ISPs. While this result is in agreement with commonly held notions that P2P traffic has generated revenue for ISPs, we believe that we are the first to attempt to quantify this effect. We leave a study of which ISPs are benefiting from this (and by how much) as part of our future work.

6. CONCLUSION

This article discussed potential issues with extending results from limited platforms to Internet wide perspectives. In particular, we showed that testbed-based views of Internet paths are limited, the properties of these paths do not extend to the edge of the network and these inaccuracies have a significant impact on inferred system-wide performance for services running at the edge. These results make a strong case for research in new evaluation strategies for Internet-scale systems, both through edge-systems traces (as those available via our CloudScope project) and new evaluation platforms.

7. REFERENCES

- [1] AGARWAL, S., AND LORCH, J. R. Matchmaking for online games and other latency-sensitive P2P systems. In *Proc. of ACM SIGCOMM* (2009).
- [2] CHEN, K., CHOFFNES, D., POTHARAJU, R., CHEN, Y., BUSTAMANTE, F., AND ZHAO, Y. Where the sidewalk ends: Extending the Internet AS graph using traceroutes from P2P users. In *Proc. of ACM CoNEXT* (2009).
- [3] CHOFFNES, D. R., AND BUSTAMANTE, F. E. Taming the torrent: A practical approach to reducing cross-ISP traffic in peer-to-peer systems. In *Proc. of ACM SIGCOMM* (2008).
- [4] CHOFFNES, D. R., SANCHEZ, M. A., AND BUSTAMANTE, F. E. Network positioning from the edge: An empirical study of the effectiveness of network positioning in p2p systems. In *Proc. of IEEE INFOCOM Mini Conference* (2010).
- [5] DABEK, COX, KAASHOEK, AND MORRIS, R. Vivaldi: A decentralized network coordinate system. In *Proc. of ACM SIGCOMM* (2004).
- [6] ISDAL, T., PIATEK, M., KRISHNAMURTHY, A., AND ANDERSON, T. Leveraging BitTorrent for end host measurements. In *Proc. of PAM* (2007).
- [7] LEDLIE, J., GARDNER, P., AND SELTZER, M. Network coordinates in the wild. In *Proc. of USENIX NSDI* (Apr. 2007).
- [8] LUMEZANU, C., BADEN, R., SPRING, N., AND BHATTACHARJEE, B. Triangle inequality variations in the internet. In *Proc. of IMC* (2009).
- [9] MADHYASTHA, H. V., ISDAL, T., MICHAEL PIATEK, DIXON, C., ANDERSON, T., KRISHNAMURTHY, A., AND UN VENKATARAMANI, A. iPlane: an information plane for distributed systems. In *Proc. of USENIX OSDI* (2006).
- [10] NG, T., AND ZHANG, H. Predicting Internet network distance with coordinates-based approaches. In *Proc. of IEEE INFOCOM* (2002).
- [11] OLIVEIRA, R., PEI, D., WILLINGER, W., ZHAN, B., AND ZHANG, L. Quantifying the completeness of the observed Internet AS-level structure. Tech. Rep. 080026, UCLA, September 2008.
- [12] PIATEK, M., MADHYASTHA, H. V., JOHN, J. P., KRISHNAMURTHY, A., AND ANDERSON, T. Pitfalls for ISP-friendly P2P design. In *Proc. of HotNets* (2009).
- [13] PLANETLAB. Planetlab. <http://www.planet-lab.org/>.
- [14] PRASAD, R. S., MURRAY, M., DOVROLIS, C., AND CLAFFY, K. Bandwidth estimation: Metrics, measurement techniques, and tools. *IEEE Network* 17 (2003), 27–35.
- [15] ROUTEVIEW. <http://www.routeviews.org/>.
- [16] SU, A.-J., CHOFFNES, D., BUSTAMANTE, F. E., AND KUZMANOVIC, A. Relative network positioning via CDN redirections. In *Proc. of ICDCS* (2008).
- [17] TANG, L., AND CROVELLA, M. Virtual landmarks for the internet. In *Proc. of IMC* (2003).
- [18] TEAM CYMRU. The Team Cymru IP to ASN lookup page. <http://www.cymru.com/BGP/asnlookup.html>.
- [19] WANG, G., ZHANG, B., AND NG, T. S. E. Towards network triangle inequality violation aware distributed systems. In *Proc. of IMC* (2007).
- [20] WONG, B., SLIVKINS, A., AND SIRER, E. Meridian: A lightweight network location service without virtual coordinates. In *Proc. of ACM SIGCOMM* (2005).
- [21] XIE, H., YANG, R., KRISHNAMURTHY, A., LIU, Y., AND SILBERSCHATZ, A. P4P: Provider portal for (P2P) applications. In *Proc. of ACM SIGCOMM* (2008).