Toward Representative Internet Measurements

Aditya Akella, Srinivasan Seshan
Dept. of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

Anees Shaikh
Networking and Policy Technologies
IBM T.J. Watson Research Center
Hawthorne, NY 10532

INTRODUCTION

Efforts to characterize and understand the structure and behavior of the Internet have a long history in the network research community. Though our understanding of the network has evolved, our knowledge of its configuration, workload, and failure modes still is far from complete. Characterizing the operation of the current Internet infrastructure and its usage patterns is essential to improve the understanding of the Internet and shape its future.

Despite the numerous studies of Internet properties, there has been little work to understand *how* to measure the network in terms of the number and distribution of measurement sites. Research studies that aim to collect and analyze Internet data from end-hosts are typically limited to a relatively small number of probes, located primarily at universities or research institutions. Naturally, the authors of such studies express caveats to their results, and offer some justification of the limited perspective. Some examples of this reasoning in studies of Internet path properties include statements such as “routes between the 37 hosts are plausibly representative” [10], “we believe that measuring the paths between our sources and a wide variety of different ISPs would provide a representative view” [2], and “the testbed topology contains paths that traverse most of the ‘large’ AS’s in the Internet” [8]. Some recent work delves into this issue in the context of discovering the Internet router-level topology [5], concluding that having more than 1 or 2 vantage points does not improve discovery of nodes and links. The authors also point out, however, that it is impossible to claim that the measurement sites used in the study are representative, and that their conclusions are unavoidably sensitive to the choice of measurement sites.

While deploying and maintaining a large-scale measurement infrastructure is extremely challenging, it is clear that further guidance as to the representativeness of results obtained using existing measurement testbeds is a crucial missing element of current network measurement research. In this paper we argue that for measurements of traffic, paths, workloads, etc., to be representative, data should be collected from locations actually used by a large portion of Internet users (with an obvious trade-off between “large” and the degree of representativeness). For the common client-server communication pattern, this means that measurement infrastructure should be placed at, or near, places (i.e., stub networks, ingress points, egress points) that are sources or destinations for most of the traffic flowing in the Internet. For popular peer-to-peer applications, measurement probes should be placed at concentration points (i.e., close to busy dial-up access servers or broadband central offices and cable head-ends). Given enough resources, a more complete approach would be to place a measurement probe at every access point in the Internet. This would allow, for example, studies of general properties of the entire network, such as a survey of delays along all paths regardless of their usage. Our view is that though such studies might be valuable, it is more important and interesting, not to mention more practical, to characterize the network based on common usage.

Realizing a vision of well-placed infrastructure for collecting representative measurements is difficult for a number of reasons. A key challenge lies in how to determine which parts of the network are the most commonly used. The notion of commonly used paths is dictated largely by application deployment (e.g., locations of the most popular Web sites, file servers, mail servers, or game servers) and user populations. The question is how to determine which application servers are most popular and where the densest groups of users reside. Another challenge arises from the variety of objectives of Internet measurement studies. Requirements for deploying probe stations to measure path properties may be quite different from those used to characterize session interarrivals or durations for a particular Internet application.

In this paper we outline a methodology for determining a more representative set of Internet measurement locations, based on actual usage of the network. While we recognize that measuring different aspects of the network implies different notions of representativeness, we focus our attention on determining locations that enable meaningful measurements of network-layer properties. The main idea behind our technique is to leverage access records from a large network of distributed application servers to understand where most of the demand originates, and also where these client requests are directed. With this information, we describe how to arrive at a set of network measurement locations that captures a significant fraction of this traffic, and ways to reduce the number of locations to a manageable level by detecting and removing redundancies. We also provide an overview of our work on evaluating several current measurement testbeds in terms of their ability to capture views of the commonly used paths in the network.

PROPOSED METHODOLOGY

As discussed above, our focus is on developing a methodology to identify representative points in the Internet from which network-layer properties related to routing, path and link metrics, and topology, can be measured. Our proposed approach is based on our view that the key to representative measurements is to collect information from the most often used portions of the network. Given the client-server nature of most Internet applications, this can be achieved by focusing on parts of the network that carry traffic from the most active clients to the most popular servers. Hence, our approach is to determine where these significant clients and servers reside.

Approximating representative points

Determining popular servers is not necessarily hard. A number of companies publish lists of the top Web sites they expect (or have measured) clients to access. Examples include Alexa Internet [3] and Nielsen/NetRatings [9]. In addition, a recent study by Wills *et al.* [14], describes a measurement-based technique to estimate the relative popularity of Internet applications by querying local DNS servers used by active client populations.

Estimating the most active client prefixes, on the other hand, is more challenging. If we assume that the geographic distribution of Internet clients roughly follows the geographic population density then using census data and other heuristics, we could potentially enumerate areas where we expect the most active clients to reside. In order to characterize and measure network-layer properties, however, it is important to further classify active clients in terms of which ISPs they are connected to. A first cut, approximate method of achieving this is to again turn to published lists which rank ISPs in metro areas according to their subscriber base [6]. From the above two classifications, it is possible to obtain a rough classification of the most active clients base on both geography and upstream ISP.

Leveraging CDNs

Clearly, the above categorization of clients is approximate at best, and calls for a more informed, accurate method. We propose employing data from distributed application servers such as large content distribution networks (CDNs) to collect accurate information about active client clusters. In addition, data from a large CDN could also be used to cross-check data about the most popular servers, since such servers are typically customers of large CDNs.

Consider, for example the Akamai CDN, which consists of a large collection of servers and monitoring nodes distributed throughout the world and connected to a variety of upstream carrier ISPs [1]. Akamai deploys collections of servers in most major cities, with each cluster attached to a distinct upstream ISP serving the city. Internally, Akamai maintains statistics of total traffic received and served by each cluster both in terms of bytes and number of requests. Using this data and ranking server clusters in terms of the total number of request arrivals or the total numbers of bytes served to end-clients, we can obtain an exhaustive and accurate list of the most active client clusters that Akamai serves. Since CDNs like Akamai serve a significant fraction of bytes in the Internet [13] and since they serve objects from some of the most popular Web sites, this approach may be expected to yield a reasonably representative set of active client locations (again, in terms of geography and topology).

An important advantage of employing data from a large CDN is that we can also obtain the actual traffic volumes seen by the most active clients and content providers. From these two pieces of information, it is possible to construct an approximate matrix of the amount of traffic between a given client cluster and a Web server. This knowledge is useful in understanding how important a particular path is for measurement purposes, and how often it should be sampled relative to other paths to obtain, for example, estimates of typical loss rates and delays experience in the Internet.

Selecting measurement locations

Once we have information on the locations of the most active Web servers and clients, the next step is to formulate a set of guidelines for the placement of measurement vantage points in order to estimate network-layer properties and metrics. Here, the goal is to select measurement points which allow us to characterize paths between the set of popular Web servers and the active client clusters. Clearly, having a measurement node located at each active client location and each popular Web server would be ideal, but this is certainly impractical, as it requires a very large number of measurement nodes. To reduce the number of measurement nodes, we propose the a simple approach in which we first identify the smallest set of nodes (where a node is characterized by both its geographic location and the ISP it belongs in, e.g., AT&T in New York) which cover a significant fraction of the most important paths. Here, paths generally means end-to-end AS level paths between the set of frequently visited Web sites and active clients, in each direction (i.e., client → server and server → client). The number of measurement nodes could be reduced by exploiting ISP customer-provider relationships to remove redundancies. For example, suppose active client groups are attached to AT&T in New York and to a smaller regional provider which is customer of AT&T. In this case, choosing a measurement node attached to AT&T in New York may be enough to “cover” AS paths to or from both client clusters. However, if the customer of AT&T is multihomed to multiple providers, it may be necessary to choose distinct measurement points covering both sets of active clients.

In selecting measurement locations, it is important to consider the trade-off between the number of client clusters and Web servers covered with respect to the total weights of the covered paths. While it is very important to measure the most used paths (paths with the highest weight), it is also critical to ensure that a reasonable spectrum of paths, traversing a variety of geographic locations and ISPs, are also measured.

EVALUATING CURRENT MEASUREMENT TESTBEDS

In the absence of an available measurement infrastructure that provides more representative vantage points, it is important to understand the completeness of results obtained from existing measurement testbeds. An additional objective of this work is to use our proposed methodology to evaluate several currently available network measurement testbeds. We can analyze, for example, how well-placed these testbed nodes are in terms of their ability to probe important Internet paths taken by a large number of clients to reach popular destination servers. Such information is a useful first step in determining the suitability of a given measurement infrastructure for observing various Internet properties. Since our evaluation is currently ongoing, our discussion is limited to an overview of our approach.

Despite being used often as vehicles for network measurements, many existing network research testbeds are designed and deployed to serve broader research objectives. In our analysis we focus primarily on the placement of testbed nodes and do not evaluate their other features, such as special functionality or the degree to which they can be customized or extended. In our initial study, we are considering the following four widely-used research testbeds:

NIMI [11]: The National Measurement Internet Infrastructure (NIMI) is a government-supported measurement infrastructure currently consisting of 37 probes. The probes are located primarily at universities and industrial and government research labs, with five nodes in Europe.

RON [4]: Resilient Overlay Networks (RON) is an MIT project focused on overlay network services. The testbed consists of about 30 nodes, located in the U.S. and five international locations. The RON testbed is notable in that it also includes a small number of broadband-connected nodes.

Skitter [7]: The Skitter project is a CAIDA project that aims to map Internet paths and uncover topology changes. The testbed consists of about 30 nodes, with a relatively large number of international nodes (approximately 10).

PlanetLab [12]: PlanetLab is a large testbed of centrally managed nodes located in about 70 distinct locations (each PlanetLab site typically has two or more machines). The testbed is intended to enable the development, deployment, and evaluation of global-scale network services. The nodes are located primarily at educational institutions in the Americas, Europe, and Asia-Pacific regions, though there are a few nodes in commercial data centers and industry labs.

Though geographic information about these testbeds is readily available, understanding their topological representativeness requires more effort. We have begun experiments, primarily based on traceroute measurements from multiple points, to generate a list of the access providers used by each node. Using this data we can analyze the distribution of networks covered by a particular testbed in terms of their size and placement in the Internet hierarchy. For studies that require bidirectional network-layer measurements between clients and servers, paths between nodes in existing testbeds may not be very useful. Many research testbed nodes are at educational or research institutions, and hence they often use Abilene (i.e, Internet2) to reach other. For example, more than half of the paths between PlanetLab nodes passed through Abilene based on data available in October 2002. Therefore, measurements between these nodes would not represent typical network characteristics. By constructing AS-paths between testbed nodes we can understand the relative extent to which they are useful for bidirectional measurements.

SUMMARY AND ONGOING WORK

In this paper we have outlined our proposal for improving the representativeness of Internet measurement studies. We advocate an approach that focuses attention on the most heavily used paths, which can be determined by examining traffic data from distributed application servers. In particular, we are interested in leveraging data from large CDNs which provide a view of active client regions, as well as the most frequently visited Web content providers. By identifying these active client clusters and popular destinations, we can begin to determine the number and locations of probes necessary to “cover” the paths taken between them. Finally, we also described our ongoing effort to analyze and quantify the ability of a few available testbeds to measure these widely used paths.

This paper provides only an overview of our approach, as this work is still in its preliminary phases. We are currently working with a large CDN service provider to collect client request statistics that would allow us to identify the most active endpoints for Web-based content and applications. Using this data, we plan to apply and further refine our proposed methodology to arrive at a set of probe locations that would provide a representative view of typical Internet path properties. Using the same data, we are also working on characterizing the ability of existing testbeds to measure these representative paths.

REFERENCES

[1] Akamai Technologies, Inc. `http://www.akamai.com`.

[2] A. Akella, S. Seshan, and A. Shaikh. An empirical evaluation of wide-area internet bottlenecks. In *Proceedings of ACM SIGCOMM Internet Measurement Conference (IMC)*, Miami, FL, October 2003.

[3] Alexa Internet, Inc. Top english language sites. `http://www.alexa.com`.

[4] D. Andersen, H. Balakrishnan, M. Kaashoek, and R. Morris. Resilient Overlay Networks. In *Proceedings of the Symposium on Operating System Principles (SOSP)*, Banff, Canada, October 2001.

[5] P. Barford, A. Bestavros, J. Byers, and M. Crovella. On the marginal utility of network topology measurements. In *Proceedings of ACM SIGCOMM Internet Measurement Workshop (IMW)*, San Francisco, CA, November 2001.

[6] N. Christenson. Ranking Internet service providers by size. `http://www.jetcafe.org/~npc/isp/large.html`, July 2003.

[7] Cooperative Association for Internet Data Analysis. skitter. `http://www.caida.org/tools/measurement/skitter/`.

[8] N. Feamster, D. G. Andersen, H. Balakrishnan, and M. F. Kaashoek. Measuring the effects of internet path faults on reactive routing. In *Proceedings of ACM SIGMETRICS*, San Diego, CA, June 2003.

[9] NetRatings, Inc. Weekly top 10 parent companies. `http://www.netratings.com`.

[10] V. Paxson. End-to-end routing behaviour in the internet. In *Proceedings of the SIGCOMM ’96 Symposium on Communications Architectures and Protocols*, September 1996.

[11] V. Paxson, A. Adams, and M. Mathis. Experiences with NIMI. In *Proceedings of Passive and Active Measurement Workshop (PAM)*, Hamilton, New Zealand, April 2000.

[12] L. Peterson, T. Anderson, D. Culler, and T. Roscoe. A blueprint for introducing disruptive technology into the internet. In *Proceedings of ACM Workshop on Hot Topics in Networks (HOTNETS)*, Princeton, NJ, October 2002. see also `http://www.planet-lab.org`.

[13] S. Saroiu, K. P. Gummadi, R. J. Dunn, S. D. Gribble, and H. M. Levy. An analysis of Internet content delivery systems. In *Proceedings of the Symposium on Operating Systems Design and Implementation (OSDI)*, Boston, MA, December 2002.

[14] C. Wills, M. Mikhailov, and H. Shang. Inferring relative popularity of internet applications by actively querying dns caches. In *Proceedings of ACM SIGCOMM Internet Measurement Conference (IMC)*, Miami, FL, October 2003.