Internet Measurement: Past, Present and Future

Tom Anderson, Mark Crovella, Christophe Diot

November 1, 2004

Abstract

This report summarizes observations about the past, present and future of research in the field of Internet measurement. Our goal is to broadly survey the range of past results in this area. We argue that Internet measurement research has resulted in a wide range of fundamental results, and these results have had significant impact on the broad domain of networking research. We show that measurement research has had a widespread impact on industry as well. We conclude with some observations future directions for research in this area.

1 Introduction

We are currently experiencing a period of rapid progress in the field of Internet measurement research. The past few years have seen more papers at the SIGCOMM conference on Internet measurement than any other topic, the creation of an entirely new annual conference, IMC, dedicated to network measurement, and frequent reliance on the results of measurement papers to motivate and evaluate broader network research.

On the other hand, research in Internet measurement can sometimes seem like the old adage about mountain climbing: we measure something because its there. There are a blindingly large number of properties that can be measured [71], and for each of those properties, there are likely to be multiple papers discussing different ways of measuring each one.

In this paper, we wanted to take a step back and ask: how can we tell whether this work is getting us anywhere? What have been the benefits of network measurement research, particularly outside the network measurement community? Put another way: suppose we were completely successful and able to understand literally everything about the Internet. Would that lead to solutions to the various problems facing the Internet today – would it lead to a more secure, more reliable, more manageable, and more cost-efficient Internet?

We believe the answer is a resounding yes, and we present concrete examples to back this up. In essence, there's been a sea change in how network research is evaluated, from the qualitative to the quantitative, driven by network measurement research. Even as recently as five years ago, much of network research was evaluated on very simple toy topologies and workloads, an approach that would be considered unacceptable today because of our much deeper understanding of network behavior. One might think that since network protocols should be designed to work on any possible topology and any possible workload and any possible failure pattern, that it doesn't or shouldn't matter what topologies, workloads and failures happen in practice. However, it is also crucial that network systems be designed to work well in the common case; by understanding the structure and behavior and failure properties of networks, we can design systems to take advantage of those properties, dramatically improving cost-performance and reliability, and making it much more likely that protocols work as expected when they are deployed.

Characterizing the common case (and common distributions) for Intenet properties has engaged the energies of the network measurement community for the past decade, and probably will consume us for the next decade. In addition, since we not only want our systems to work well on today's network but also on future networks, it is important to understand the root causes of network measurements, to be able to predict how those properties will evolve over time.

Finally, infrastructures for network measurement raise the intriguing possibility of redesigning network protocols and systems to take advantage of detailed knowledge of the current state and usage of the network. Although this is still in its early stages, such an "information plane" could alter how we approach network protocol design. In this, industry has been leading the way in deploying systems that adapt in real-time to measured data.

The rest of this paper is structured as follows. In the next section, we discuss past progress in network research, outlining its impact on our conceptual understanding of the Internet and how that has changed how we design protocols. We then survey some of the progress being made in industry, in taking the results of measurement research and applying them in practice. We conclude with several emerging trends in measurement infrastructures, and how these might impact network measurement research in the future.

2 An Overview of Research Results

2.1 Three Categories of Internet Measurement

The field of Internet measurement is still in its early stages, and new directions arise frequently. However one can broadly categorize results in the field so far into three categories:

1. Internet Structure and Infrastructure. This covers work in Internet topology; understanding the interconnection patterns of routers and

links; understanding the structure of the Interdomain system; bandwidth measurement; delay and loss measurement; and understanding the rates and nature of change in all these properties.

2. Traffic and Network Monitoring. This covers work in statistical properties of network traffic; understanding and predicting where traffic enters the network and where it exits; predicting network traffic for capacity planning and network engineering; and identifying unusual traffic patterns for network management and network security.
3. Application Properties. This covers work in characterizing the demand placed on the Internet by its principal applications; describing how Web, streaming media, file transfers, DNS, Email, and other applications make use of the network; understanding how application properties dictate traffic properties; understanding how varying nature of network service affects user-perceived performance.

While these three categories do not encompass all work in Internet measurement, they cover the majority of the work and provide a convenient framework for discussion.

2.2 A Look at Some Results

Taking each of these areas in turn, we can review some recent results and chart their impact. It is not our intent here to exhaustively survey the field, but rather to show sample results that are illustrative of the area (hence we only provide a subset of relevant citations).

2.2.1 Internet Infrastructure

- **Result: Router and AS level topology characterization.** The Internet has grown in a distributed manner as the result of countless local structural changes. While the lack of central coordination in Internet organization has been an important factor in its success, it means that the global structure of the network is largely unknown.

  Beginning in 1999, measurements of the interconnection patterns of routers, as well as that of autonomous systems (ASs), have uncovered surprising properties [32, 70]. The variability of the degree of routers and ASes is much higher than expected. Most routers or ASes have only a few connections to others; but some are very highly connected. Traditional models of random graphs do not capture this phenomenon.

  Additionally, these networks were found to have the "small world" property: low average path length combined with high local clustering.

  Impact: The structure of the network strongly affects the network's performance, fault tolerance, and the design of protocols and applications. Accordingly, these results have changed the way that researchers build models for network simulation; it has drawn attention to the potential for "weak points" in the Internet's connectivity; and it has stimulated the search for underlying explanations for the driving factors that determine router-level connectivity [70, 47].

As measurements of Internet structure have led to a richer understanding of small world topologies, we are starting to see those in the design of distributed systems. An example is the design of Symphony, a second-generation DHT that is more robust and easier to configure than e.g. Chord, since it only needs to preserve the small world property to provide log scale lookup and repair [50].

More broadly, the surprising nature of Internet connectivity has stimulated work to better understand so-called "complex networks." This has grown to be a subfield of statistical mechanics, primarily pursued in the physics community, but with contributions from researchers in computer science, statistics, probability, and graph theory. The immense amount of research stimulated in just a few short years is clear from the book "Evolution and Structure of the Internet," [61] which includes over 300 references, the vast majority of which are from the past five years. The field has even spawned a book aimed at general audiences [10].

- **Result: Characterization of BGP interdomain system.** Since 1998, researchers at the University of Oregon, and more recently elsewhere, have been monitoring the state of the interdomain (BGP) routing system [53, 63]. The interdomain routing system is the fundamental glue that allows independent service providers to cooperatively provide global interconnectivity to customers.

  Impact: Measurements of the size of BGP routing tables have been crucial in driving the architectural evolution of the Internet. Internet architectural work in the late 80's and early 90's was driven by concerns about exhaustion of the 32-bit IPv4 address space, and resulted in the adoption of 128 bit addresses in IPv6 as the next standard for Internet addressing. However, more recent measurements of the BGP system have demonstrated that IPv4 address space exhaustion is a long way off, perhaps 20 years or more [41].

  Furthermore, given the central role of the interdomain system in the Internet, its stability is of crucial concern. BGP measurements have been central in understanding the sources of instability in the Internet and will undoubtedly drive improvements in stability [62].

- **Result: Development of methods to infer network hidden properties.** A fundamental design principle of the Internet has to strive for simplicity of core elements, while pushing complexity into endsystems. This has often been called the "stupid network," and seems to have been an enabler of the Internet's rapid growth. However this design principle has led to an absence of instrumentation in the network's core. Furthermore, competitive pressures have discouraged providers from sharing the meager information that is available with researchers.

An example is the lack of accessible per-link and per-router statistics. The research community cannot even state with confidence where typical bottlenecks occur in the Internet, how long congestion events last, or the primary determiners of network congestion.

The result has been an intensive effort to develop methods capable of estimation network-internal characteristics using measurements taken only at endpoints. Such methods include estimating the capacity of Internet bottlenecks [18, 42] and identifying where in the network delays and losses are taking place [16, 75, 49].

Knowledge of the physical characteristics of the network, while useful, is not sufficient in and of itself. Other research has targeted inference of routing policy – how routes are chosen through the physical topology – at both the router and AS-level [36, 69] Other work has studied whether it is possible to measure workloads, such as the amount of traffic entering each point in the network that is destined for each egress point. Direct measurement of the flows between all sources and destinations (the "traffic matrix") is impractical, so methods to infer this information have received considerable attention [52, 78].

Impact: These developments have contributed immensely to our understanding of the properties of networks. Traffic matrix estimation is an area where each new method is immediately tested in the network. This is a typical example of close flow of research results from laboratories to operations.

Furthermore, these problems have stimulated the development of new statistical methods and have strengthened network research by drawing in many researchers from statistics.

- **Result: Development of Internet distance estimation methods.** In many planetary-scale applications there is opportunity for optimization based on knowledge of point-to-point delays in the underlying network. However, active measurement of network delays can be difficult, time consuming, and can add to network load. In response a number of methods have been proposed for estimation of network delays based on reduced or incomplete measurements [35, 20, 38].

  A promising method that has received considerable attention is to assign *coordinates* to nodes [56, 48]. The idea is to assign coordinates in such a manner that the associated distance approximates network delay (round-trip propagation and transmission time). Large scale measurements have shown the method to be surprisingly accurate and scalable [73].

  Impact: These methods offer the potential to enable a scalable network distance estimation service, which has been a goal since the early days of the World Wide Web [55]. Example applications that can benefit from such knowledge include content delivery networks, peer-to-peer networks, multiuser games, overlay routing networks, and applications employing dynamic server selection.

- **Result: Development of Geolocation methods.** The Internet is often credited with the "death of distance" for its remarkable ability to unify global information systems and support wide-area collaboration. However this very strength is sometimes a liability; at times it can be important to obtain some knowledge of the true physical location of Internet systems or users. An analagous need has driven the introduction of geolocation capability for enhanced 911 services in cellular telephony.

  To address this need, promising methods have been developed for internet based geolocation based on various kinds of network measurement. Methods based on network structure [57], finding nearby landmarks [57, 79], or direct triangulation [37] have been proposed, and the state of the art suggests that geolocation to the level of a metropolitan area is often feasible.

  Impact: Geolocation is a basic capability that enables a wide range of applications. Network security can be improved if it is possible to identify the physical location of network abusers. As the Internet takes on a strategic role in society, homeland security is enhanced when cyberattackers can be located. Geolocation is also important for content localization, targeted advertizing, and server selection.

- **Result: Characterization of Internet route efficiency.** In principle, the algorithms used in Internet routing seek to find short, reliable paths between network endpoints. However this principle is complicated and sometimes frustrated by the nature of business relationships between service providers, and by interdomain routing protocols that hide the information needed to find better routes. Measurements of the actual paths taken in the Internet show that routing in practice can often be sub-optimal, in both efficiency and reliability [65, 74, 8, 69].

  Impact: These observations have motivated a large amount of work in systems to take advantage of this inefficiency. One approach is to route traffic through a performance-oriented overlay; in these systems, application-level routing is used to find routes that are better than those provided by the IP routing layer [9, 64]. Another approach is so-called smart routing through a more intelligent selection of the next hop AS; several commercial products have been developed in this space [4, 2], and there is recent evidence that smart routing achieves most of the benefits of overlay routing [7].

2.2.2 Traffic and Network Monitoring

- **Result: Models for network traffic on short timescales.** Performance evaluation is a key tool in the design of routers and other network elements. Performance evaluation techniques rely on accurate traffic models, particularly on short timescales (ranging from micoseconds to minutes). Starting with a landmark paper in 1993 [46], researchers have come to realize that network traffic has unusual and surprising characteristics that are not well captured by traditional models. This is referred to as traffic *self-similarity*, and has significant implications for the performance of routers and networks as a whole [29].

  Impact: Recognition of the phenomenon of self-similarity has revolutionized traffic modeling and network performance evaluation [60, 12]. It is now generally understood that any simulation or analytic study intended to apply to Internet traffic must incorporate self-similar models, or otherwise model the high burstiness attributable to self-similar traffic.

- **Result: Models for network traffic on longer timescales (hours to years).** The past few years have seen the beginning of systematic large-scale network measurements by ISPs such as Sprint and AT&T [34, 22]. These measurements have identified the crucial features of network traffic over long timescales, separating long-range (multi-year) trends from predictable variation such as daily and weekly periodicities [58].

  Impact: The vast majority of network engineering today is done by trial and error adjustments to network configurations. Models for how traffic varies on timescales from hours to years open the potential for scientifically grounded capacity planning and configuration of networks.

- **Result: Development of new algorithms for high speed network monitoring.** Although traces of network traffic are crucial for addressing a wide range of issues (both short- and long-timescale), technology trends make it increasingly difficult to collect traffic in the network core. The dramatically increasing link rates in ISP networks make capturing all packets on a link impractical.

  In response a wide variety of techniques have been developed to allow routers and other network elements to sample or estimate metrics of interest rather than relying on full packet capture [28, 31, 30, 27].

  Impact: These approaches address an acute need and have been quickly embraced by network operators. A number of proposals are moving quickly toward standardization [26, 81] and we will likely soon see supported implementations in routers.

- **Result: Discovery of the prevalence of long-tailed distributions (the "elephants and mice" phenomenon) in network workloads.** While it has long been understood that all aspects of network workloads show variability, the extremely skewed nature of many workloads – in which the vast majority of demands are small, but most work is concentrated in a few large demands – has only recently been well understood. This phenomenon has been termed "mice and elephants." It's been documented with respect to the sizes of network flows, the sizes of objects transferred over the network, the different rates at which flows send, the distribution of traffic across network ingress and egress points, and other related aspects of network workloads [23].

Impact: The mice and elephants phenomenon now informs almost all work on traffic engineering and congestion control. For example, the fact that generally most bytes are contained in a small number of long flows has suggested new routing algorithms [33, 66] and has made available performance improvements arising from differential treatment of long and short flows [51, 39, 25, 40].

2.2.3 Internet Applications Research

- **Result: Applicability of Zipf's Law to Internet Applications.** Zipf's Law is a phenomenon appearing in popularity distributions. When one examines the relative popularity of a set of objects, such as a set of Web pages, one often finds a highly skewed distribution (usually taking a power-law shape). Thus most object requests are concentrated in a small subset of the available objects. This phenomenon has been widely noted in requests for Web pages [5].

  Impact: The concentration of requests implied by Zipf's law means that caching can be much less effective than otherwise hoped when applied to Web workloads [13, 77]. This has reduced reliance on caching as a solution to Web latency and driven trends toward more proactive schemes such as content delivery networks.

  Furthermore, theoretical work shows that the locality that is present in Web workloads is effectively removed by caching [14]. This has led to the conclusion that there are practical limits on the ability of caching for Zipf-like workloads. Prior to these results, there were substantial efforts to build deep hierarchies of proxy caches, but they were in essence proven to be infeasible as a result of measurement.

- **Result: Discovery of the structure of the Web graph.** The Web can be considered to be a directed graph, in which pages are vertices and hyperlinks are directed edges. Measurements of the structure of the Web graph have identified the presence of key nodes. Within all the pages touching on a given topic, some pages will have high outdegree ("hubs") and some will have high indegree ("authorities") [43, 44]. This structure of hubs and authorities arises as an emergent property through the distributed action of the set of all Web page authors.

  Impact: Mining the structure of the Web has become essential to the performance of the best search engines. This is because the construction of hubs and authorities in fact encodes user knowledge about the relative value of web pages touching on a given topic. This provides a measure of relevance, which is a crucial metric for a useful search engine. By mining this structure (using, for example, Google's PageRank algorithm [15]), search engines are able to select the "right" pages to return first from the set of all pages matching a given query string (which is often an immense set).

- **Result: Heavy-tailed distribution of Web object sizes and Traffic Self-similarity.** Although, as already noted, traffic self-similarity was first noted in 1993, the immediate explanation for *why* traffic showed this surprising property was not immediately clear.

  Theoretical arguments subsequently established that self-similar traffic could arise due to the presence of many sources, each transferring for time periods with lengths drawn from a heavy-tailed distribution (a highly skewed distribution with a power-law tail shape) [76].

  At this juncture, measurements of the Web played a role in tracing the underlying causes of self-similarity. By 1996, the Web was the most popular application on the Internet, and measurements of Web object sizes showed that they too exhibited a heavy-tailed distribution [24]. The use of the network to transport Web objects was therefore shown to be sufficient to explain the presence of self-similarity in the resulting traffic [59].

  Impact: The realization of the connection between traffic self-similarity (a traffic property) and Web object sizes (an application and filesystem property) drew into focus the importance of application demands as a determiner of network properties, and provided significant insight into the underlying causes of self-similarity in network traffic.

  Furthermore, these results led to a natural approach for workload generation for network simulation. Previously, although the importance of self-similarity to network performance evaluation was understood, the proper way to introduce self-similar traffic into simulation settings was not clear. The understanding of the connection between object transfer sizes and traffic properties meant that *source-level* models could easily be used to generate self-similar traffic. Such source-level models became the standard for accurate traffic generation [11].

- **Result: Measurements of Internet Worm Propagation.** As worms have become increasingly common on the Internet, attention has focused on methods to counter them.

  Initial methods focused on patching system vulnerabilities before worms could exploit them. This is exactly analagous to vaccination in the case of infectious diseases, and it was natural to use models developed to characterize epidemics to evaluate the effectiveness of such methods.

Based on a combination of modeling and network measurement, it has become clear that worms can be capable of extremely rapid spread [72]. This has been borne out, with recent worms appearing to spread with unprecedented speed [67].

Impact: These measurement results have refocused the community away from solutions that require human response (such as patching systems, interrupting network access, or reconfiguring firewalls). The research community is now focused on automated worm detection and containment [54, 80].

3 Observations

The results and impact reviewed here span a wide range of topics and directions, but they have some common elements. Good network measurement studies (and there are many more than are reported here) have a direct impact on practice, and that's a good way to tell whether the measurement study is worth pursuing – will it provide insight into how systems should be built?

These results are not just about the need to collect ever-larger quantities of data, nor are they a form of "stamp collecting." They are about squeezing insight from data. The insights gained form the foundation for system design – whether it be the design of network infrastructure or the applications that it supports.

This foundational aspect of Internet measurement can be seen in the research impact that such studies have. The Citeseer database compiles a list of the most-cited papers published in each year. For more than ten years, (going back to 1993), Internet measurement papers have been among the 20 most-cited published each year [21].

As a result, and not surprisingly, Internet measurement is experiencing an explosion of research interest. The Internet Measurement Workshop (held in 2001 and 2002) needed to convert to conference format in 2003 to accommodate the demand for attendance. Submissions to IMW/IMC have gone from 93 in 2002, to 103 in 2003, to 157 in 2004. Likewise, submissions to the Passive and Active Measurement Workshop (PAM) hit an all-time high of over 200 in 2004 (for a two-day workshop!).

3.1 Activities at Industry Research Labs

ISPs and vendors are quite active in expanding our understanding of the Internet through measurement-based research and in putting the results of that improved understanding into practice. ISPs in particular have leveraged their ability to collect and analyze data about the operation of their own networks. Two tier-1 ISP's have been at the forefront of this effort: AT&T and Sprint.

AT&T has been the pioneer in measurement based research, work that has been driving the evolution of the AT&T IP technology and backbone. We can mention three areas in which research results have been integrated into network design and management, yielding a more cost-efficient and reliable network:

- Gigascope is a passive monitoring infrastructure designed to track traffic evolution and anomalies in the AT&T backbone. Gigascope is deployed on passive taps in the backbone and can be queried by operations to provide various kind of traffic statistics.
- Research on inference of traffic matrices has led to the deployment of a tool that is currently used to evaluate the evolution of the traffic in the backbone and track changes.
- Tools and models issued from research projects are used to control the stability of routing through the analysis of routing information.

Sprint started a couple of years after AT&T. Research at Sprint has contributed to the Sprintlink engineering and operation the following (non exhaustive) list of tools and techniques:

- Weights in ISIS are now allocated using a tool and a heuristic developed by the labs.
- The matching of IP routes to fiber network is realized using techniques developed at the lab in order to maximize the redundancy and the robustness of the network.
- BGP configuration is now assisted by research-developed tools.
- Traffic matrices are also computed using tools designed at Sprintlabs. These tools are very similar to those designed by AT&T and benefited from parallel and cooperative research works.
- Traffic forecasting is also based on traffic analysis.

It should be noticed that most of these tools have replaced techniques relying on human expertise and intuition. The net result is more robust and stable networks.

Despite having no research lab, Cisco is funding a large number of measurement-based research projects in universities. These research results directly impact products:

- Traffic measurement tools, such as netflow, are being continuously improved based on researcher input. For example, netflow sampling strategies are currently being improved by research labs in Europe.
- In order to implement PSAMP (proposed by AT&T) to measure delay at various point of a network, application of hash functions have been added to the forwarding path.
- Results of studies led at Sprint on fast convergence of IGP routing protocols have been implemented on Cisco routers. More recently, AT&T helped Cisco modify their BGP implementation.

Following the lead of Sprint and AT&T, numerous ISPs are providing traces to univerities. Abilene has a very sophisticated measurement infrastructure. Geant, the european research backbone is following in the steps of Abilene. Global Crossing, Akamai, and France Telecom are regular publishers in measurement related conferences and workshops. Recently, Intel has begun providing support to two measurement related projects: Planetlab [3] and CoMo [17]. These two projects aid the wide dissemination of measurement based research at the overlay level (planetlab) and in the form of passive monitoring (CoMo). Intel is also making traces available to academic researchers.

This is a limited but illustrative set of example of the strong involvement of industry in measurement driven network research. We could certainly mention other research activities led by industry or promoted and funded by industry.

In addition to pursuing measurement based research themselves, industry labs are also playing an important role in the diffusion of research and in the building of a community, through

- Funding. Sprint has been funding $1M of university research for the past 6 years; AT&T's investment is of a similar scope.
- Making traces available. Even though the data are often sensitive, both Sprint and AT&T have made traces available to numerous visiting students and faculty.
- Publishing their results.

3.2 Research Testbeds

The increased interest in Internet measurement means that new testbed-oriented projects have sprung up to address measurement needs. Of particular note are:

1. Planetlab. Planetlab is a globally distributed open platform for experimenting with wide-area services. As such, a large number of Internet measurement projects have begun to make use of Planetlab.

   From its inception in March 2002, Planetlab has grown incredibly rapidly; it currently comprises 392 nodes at 164 sites. Planetlab has generated considerable energy and its software base has become quite sophisticated. The key to this has been its open development and use model.
2. EMUlab/Netbed. Netbed is a shared environment for distributed systems research. It provides integrated access to simulated, emulated, and wide-area network testbeds. The simulated and emulated testbeds are hosted at U. Utah; the wide-area testbed includes 50-60 nodes geographically distributed across approximately 30 sites.
3. Scriptroute. Scriptroute is a software suite that enables active network measurement by application-level programs. As such, it provides a safe

way for administrators to allow machines to be used for active network measurement (without granting root priviledges to users). Released in fall 2002, Scriptroute is now widely used for active Internet measurement projects. Scriptroute runs on Planetlab hosts, but is not restricted to those hosts.

4. Archival storage. Skitter and RouteViews have shown the value of continual measurement of network properties, both to evaluate new measurement techniques and to show how the network is evolving over time. New efforts have been launched to measure a wider spectrum of network properties, continuously over a period of years, to help undergird models of the Internet's evolution.
5. End host measurement systems. Several recent efforts have focused on enlisting end hosts in the task of passively or actively monitoring Internet behavior, enabling more complete and real-time measurements than is possible even with a widely distributed system such as Planetlab; of course, a key challenge is collecting and distributing the data to where it is needed. Possible benefits include automated response to viruses and worms, more rapid diagnosis of network problems, and various forms of network optimization. Examples of such systems include Neti@home [19], Intel's network oracle project, DIMES [68], and DShield [1].
6. Wireless testbeds. Wireless networking is increasingly relying on measurements of deployed systems to guide designs and drive protocols; several researchers argue that purely simulation-based models have led the field astray [45] and are working to fix those problems [6].

These projects suggest that a new stage of Internet measurement research is beginning: one in which large-scale, cross-institutional cooperation is developing, to address problems with dramatically increased size and scope. It is tempting to compare this state to the early days of high energy physics research.

4 Conclusions

Internet measurement is a relatively young field. Although the first significant results may be found as far back as 1993 or before, the emergence of a coherent research community probably dates to the first Internet Measurement Workshop in 2001. For such a young field, the results obtained already are very encouraging, but there are still many aspects of the Internet's structure, workload, and applications that are only poorly understood. Given the accelerating pace of work in the field, the increasing number of researchers working in the area, and the emergence of large-scale cooperative testbeds for Internet measurement, we are confident that over the next ten years we will gain a much deeper understanding of these issues, leading in turn to a much more secure, reliable, and efficient set of Internet protocols.

References

[1] Distributed intrusion detection system. http://www.dshield.org.

[2] Internap, inc. http://www.internap.com.

[3] Planetlab. http://www.planet-lab.org.

[4] Routescience, inc. http://www.routescience.com.

[5] Lada A Adamic and Bernado A. Huberman. Zipf's Law and the Internet. *Glottometrics*, 3:143–150, 2002.

[6] Daniel Aguayo, John Bicket, Sanjit Biswas, Glenn Judd, and Robert Morris. Link-level measurements from an 802.11b mesh network. In *Proceedings of ACM SIGCOMM*, 2004.

[7] Aditya Akella, Jeffrey Pang, Bruce Maggs, Srinivasan Seshan, and Anees Shaikh. A comparison of overlay routing and multihoming route control. In *Proceedings of ACM SIGCOMM*, pages 93–106. ACM Press, 2004.

[8] David G. Andersen, Hari Balakrishnan, M. Frans Kaashoek, and Robert Morris. Resilient overlay networks. In *Proceedings of Symposium on Operating Systems Principles (SOSP)*, pages 131–145, 2001.

[9] David G. Andersen, Hari Balakrishnan, M. Frans Kaashoek, and Robert Morris. Resilient overlay networks. In *Proceedings of the 18th ACM SOSP*, October 2001.

[10] Albert-Laszlo Barabasi. *Linked: The New Science of Networks*. Perseus Publishing, Cambridge, MA, 2002.

[11] Paul Barford and Mark E. Crovella. Generating representative Web workloads for network and server performance evaluation. In *Proceedings of Performance '98/SIGMETRICS '98*, pages 151–160, July 1998.

[12] Paul Barford and Sally Floyd. The self-similarity and long range dependence in networks web site. Available at http://www.cs.bu.edu/pub/barford/ss_lrd.html.

[13] Azer Bestavros, Robert L. Carter, Mark E. Crovella, Carlos R. Cunha, Abdelsalam Heddaya, and Sulaiman A. Mirdad. Application-level document caching in the internet. In *Proceedings of the Second International Workshop on Services in Distributed and Networked Environments (SDNE'95)*, June 1995.

[14] Lee Breslau, Pei Cao, Li Fan, Graham Phillips, and Scott Shenker. Web caching and zipf-like distributions: Evidence and implications. In *Proceedings of INFOCOM '99*, pages 126–134, 1999.

[15] Sergey Brin and Lawrence Page. The anatomy of a large-scale hypertextual web search engine. In *Proceedings of World-Wide Web '98 (WWW7)*, April 1998.

[16] Tian Bu, Nick Duffield, Francesco Lo Presti, and Don Towsley. Network tomography on general topologies. In *Proceedings of ACM SIGMETRICS*, pages 21–30, Marina Del Rey, CA, June 2002.

[17] Intel Research Cambridge. Continuous monitoring (como). http://www.cambridge.intel-research.net/como/.

[18] Robert L. Carter and Mark E. Crovella. Measuring bottleneck link speed in packet switched networks. *Performance Evaluation*, 27&28:297–318, 1996.

[19] Jr. Charles Robert Simpson and George F. Riley. NETI@home: A distributed approach to collecting end-to-end network performance measurements. In *The Fifth Annual Passive and Active Measurement Workshop (PAM)*, April 2004.

[20] Y. Chen, K. Lim, R. H. Katz, and C. Overton. On the stability of network distance estimation. In *Proceedings of ACM SIGMETRICS Practical Aspects of Performance Analysis Workshop (PAPA 2002), in ACM SIGMETRICS Performance Evaluation Review*, September 2002.

[21] Most cited source documents in the CiteSeer.IST database as of july 2004. available at http://citeseer.ist.psu.edu/source.html.

[22] Chuck Cranor, Theodore Johnson, and Oliver Spatscheck. Gigascope: a stream database for network applications. In *Proceedings of ACM SIGMOD*, June 2003.

[23] Mark E. Crovella. Performance evaluation with heavy tailed distributions. In *Lecture Notes in Computer Science 1786*, pages 1–9, March 2000.

[24] Mark E. Crovella and Azer Bestavros. Self-similarity in World Wide Web traffic: Evidence and possible causes. *IEEE/ACM Transactions on Networking*, 5(6):835–846, December 1997.

[25] Mark E. Crovella, Robert Frangioso, and Mor Harchol-Balter. Connection scheduling in Web servers. In *1999 USENIX Symposium on Internet Technologies and Systems (USITS '99)*, 1999.

[26] N.G. Duffield, A. Greenberg, M. Grossglauser, J. Rexford, D. Chiou, B. Claise, P. Marimuthu, and G. Sadasivan. A framework for packet selection and reporting. Internet Draft, October 2003.

[27] Nick Duffield and Carsten Lund. Predicting Resource Usage and Estimation Accuracy in an IP Flow Measurement Collection Infrastructure. In *Proceedings of the ACM SIGMETRICS Internet Measurement Conference*, November 2003.

[28] Nick Duffield, Carsten Lund, and Mikkel Thorup. Estimating Flow Distributions from Sampled Flow Statistics. In *Proceedings of ACM SIGCOMM*, August 2003.

[29] A. Erramilli, O. Narayan, and W. Willinger. Experimental queueing analysis with long-range dependent packet traffic. *IEEE/ACM Transactions on Networking*, 4(2):209–223, April 1996.

[30] Cristian Estan, Stefan Savage, and George Varghese. Automated measurement of high volume traffic clusters. In *Proceedings of the ACM SIGMETRICS Internet Measurement Workshop*, pages 177–178, Marseilles, November 2002.

[31] Cristian Estan and George Varghese. New directions in traffic measurement and accounting. In *Proceedings of ACM SIGCOMM*, Pittsburgh, PA, August 2002.

[32] Michalis Faloutsos, Petros Faloutsos, and Christos Faloutsos. On power-law relationships of the internet topology. In *Proceedings of SIGCOMM '99*, 1999.

[33] Anja Feldmann, Jennifer Rexford, and Ramon Caceres. Efficient policies for carrying web traffic over flow-switched networks. *IEEE/ACM Transactions on Networking*, December 1998.

[34] C. Fraleigh, C. Diot, S. Moon, P. Owezarski, K. Papagiannaki, and F. Tobagi. Design and deployment of a passive monitoring infrastructure. In *Proceedings of the Passive and Active Monitoring Workshop*, April 2001.

[35] Paul Francis, Sugih Jamin, Vern Paxson, Lixia Zhang, Daniel F. Gryniewicz, and Yixin Jin. An architecture for a global Internet host distance estimation service. In *IEEE INFOCOM 1999*, pages 210–217, New York, NY, March 1999. IEEE.

[36] Lixin Gao. On inferring autonomous system relationships in the Internet. *IEEE/ACM Transactions on Networking*, 9(6):733–745, 2001.

[37] Bamba Gueye, Artur Ziviani, Mark Crovella, and Serge Fdida. Constraint-based geolocation of Internet hosts. In *Proceedings of the ACM/SIGCOMM Internet Measurement Conference 2004*, October 2004.

[38] Krishna P. Gummadi, Stefan Saroiu, and Steven D. Gribble. King: Estimating latency between arbitrary Internet end hosts. In *Proceedings of 2nd ACM Sigcomm Internet Measurement Workshop 2002*, November 2002.

[39] Liang Guo and Ibrahim Matta. The War between Mice and Elephants. In *Proceedings of ICNP'2001: The 9th IEEE International Conference on Network Protocols*, Riverside, CA, November 2001.

[40] Mor Harchol-Balter, Bianca Schroeder, Nikhil Bansal, and Mukesh Agrawal. Size-based scheduling to improve web performance. *ACM Transactions on Computer Systems*, 21(2):207–233, May 2003.

[41] Geoff Huston. IPv4 - how long have we got? ISP Column, available at www.potaroo.net, July 2003.

[42] Manish Jain and Constantinos Dovrolis. End-to-end available bandwidth: Measurement methodology, dynamics, and relation with TCP throughput. In *Proceedings of ACM SIGCOMM*, pages 295–308, Pittsburgh, PA, August 2002.

[43] Jon M. Kleinberg. Authoritative sources in a hyperlinked environment. In *Proceedings of ACM-SIAM Symposium on Discrete Algorithms*, pages 668–677, January 1998.

[44] Jon M. Kleinberg. Hubs, authorities, and communities. *ACM Computing Surveys*, 31(4), December 1999.

[45] David Kotz, Calvin Newport, and Chip Elliott. The mistaken axioms of wireless-network research. Technical Report TR2003-467, Dept. of Computer Science, Dartmouth College, July 2003.

[46] W. Leland, M. Taqqu, W. Willinger, and D. Wilson. On the self-similar nature of ethernet traffic. In *Proceedings of SIGCOMM '93*, pages 183–193, September 1993.

[47] Lun Li, David Alderson, Walter Willinger, and John Doyle. A first-principles approach to understanding the Internet's router-level topology. In *Proceedings of SIGCOMM 2004*, pages 3–14. ACM Press, 2004.

[48] Hyuk Lim, Jennifer C. Hou, and Chong-Ho Choi. Constructing Internet coordinate system based on delay measurement. In *Proceedings of the ACM/SIGCOMM Internet Measurement Conference (IMC-03)*, 2003.

[49] Ratul Mahajan, Neil Spring, David Wetherall, and Thomas Anderson. User-level Internet path diagnosis. In *Proceedings of the nineteenth ACM Symposium on Operating Systems Principles (SOSP)*, pages 106–119. ACM Press, 2003.

[50] Gurmeet Singh Manku, Mayank Bawa, and Prabhakar Raghavan. Symphony: Distributed hashing in a small world. In *Proceedings of USITS 2003*, March 2003.

[51] Ibrahim Matta and Liang Guo. Differentiated Predictive Fair Service for TCP Flows. In *Proceedings of ICNP'2000: The 8th IEEE International Conference on Network Protocols*, Osaka, Japan, October 2000.

[52] Alberto Medina, Nina Taft, Kave Salamatian, Supratik Bhattacharyya, and Christophe Diot. Traffic matrix estimation: Existing techniques and new directions. In *Proceedings of ACM SIGCOMM*, pages 161–174, Pittsburgh, PA, August 2002.

[53] David Meyer. Routeviews project page. available at www.routeviews.org, 2004.

[54] David Moore, Colleen Shannon, Geoffrey M. Voelker, and Stefan Savage. Internet quarantine: Requirements for containing self-propagating code. In *Proceedings of INFOCOM*, 2003.

[55] Keith Moore, Jason Cox, and Stan Green. Sonar - a network proximity service. Internet Draft, available at http://www.netlib.org/utk/projects/sonar/draft-moore-sonar-01.txt, 1996.

[56] E. Ng and H. Zhang. Predicting Internet network distance with coordinates-based approaches. In *Proceedings of Infocom*, 2002.

[57] Venkata N. Padmanabhan and Lakshminarayanan Subramanian. An Investigation of Geographic Mapping Techniques for Internet Hosts. In *Proceedings of ACM SIGCOMM*, August 2001.

[58] Konstantina Papagiannaki, Nina Taft, Zhi-Li Zhang, and Christophe Diot. Long-Term Forecasting of Internet Backbone Traffic: Observations and Initial Models. In *Proceedings of IEEE Infocom*, March 2003.

[59] Kihong Park, Gi Tae Kim, and Mark E. Crovella. On the relationship between file sizes, transport protocols, and self-similar network traffic. In *Proceedings of the Fourth International Conference on Network Protocols (ICNP'96)*, pages 171–180, October 1996.

[60] Kihong Park and Walter Willinger, editors. *Self-Similar Network Traffic and Performance Evaluation*. Wiley / Wiley Interscience, New York, 1999.

[61] Romualdo Pastor-Satorras and Allesandro Vespignani. *Evolution and Structure of the Internet*. Cambridge University Press, Cambridge, UK, 2004.

[62] Jennifer Rexford, Jia Wang, Zhen Xiao, and Yin Zhang. BGP routing stability of popular destinations. In *Proceedings of the ACM SIGMETRICS Internet Measurement Workshop*, pages 197–202, Marseilles, November 2002.

[63] RIPE. Ripe ncc routing information service. available at www.ripe.net/ris, 2004.

[64] Stefan Savage, Tom Anderson, Amit Aggarwal, David Becker, Neal Cardwell, Andy Collins, Eric Hoffman, John Snell, Amin Vahdat, Geoff Voelker, and John Zahorjan. Detour: a case for informed internet routing and transport. *IEEE Micro*, 19(1):50–59, January 1999.

[65] Stefan Savage, Andy Collins, Eric Hoffman, John Snell, and Thomas Anderson. The end-to-end effects of Internet path selection. In *Proceedings of ACM SIGCOMM*, pages 289–299. ACM Press, 1999.

[66] Anees Shaikh, Jennifer Rexford, and Kang Shin. Load-sensitive routing of long-lived IP flows. In *Proceedings of ACM SIGCOMM '99*, pages 215–226, September 1999.

[67] Colleen Shannon and David Moore. The spread of the witty worm. Available at http://www.caida.org/analysis/security/witty/, March 2004.

[68] Yuval Shavitt. The dimes projecct. http://www.netdimes.org.

[69] N. Spring, R. Mahajan, and T. Anderson. Quantifying the causes of path inflation. In *Proceedings of ACM SIGCOMM*, 2003.

[70] N. Spring, R. Mahajan, D. Wetherall, and T. Anderson. Measuring ISP topologies with rocketfuel. *IEEE/ACM Transactions on Networking*, 12(1):2–16, Feb 2004.

[71] N. Spring, D. Wetherall, and T. Anderson. Reverse-engineering the internet. In *Proc. of the 2nd Workshop on Hot Topics in Networks*, Nov 2003.

[72] Stuart Staniford, Vern Paxson, and Nicholas Weaver. How to 0wn the internet in your spare time. In *Proceedings of the 11th USENIX Security Symposium (Security '02)*, 2002.

[73] Liying Tang and Mark Crovella. Virtual landmarks for the Internet. In *Proceedings of the ACM/SIGCOMM Internet Measurement Conference 2003*, October 2003.

[74] H. Tangmunarunkit, R. Govindan, D. Estrin, and S. Shenker. The impact of routing policy on Internet paths. In *Proceedings of IEEE Infocom*, April 2001.

[75] Y. Tsang, M. J. Coates, and R. Nowak. Network delay tomography. *IEEE Trans. Signal Processing, special issue on Signal Processing in Networking*, 51(8):2125–2136, Aug 2003.

[76] Walter Willinger, Murad S. Taqqu, Robert Sherman, and Daniel V. Wilson. Self-similarity through high-variability: Statistical analysis of Ethernet LAN traffic at the source level. In *Proceedings of SIGCOMM '95*, pages 100–113, Boston, MA, 1995.

[77] Alec Wolman, M. Voelker, Nitin Sharma, Neal Cardwell, Anna Karlin, and Henry M. Levy. On the scale and performance of cooperative Web proxy caching. In *Proceedings of the seventeenth ACM Symposium on Operating Systems Principles (SOSP)*, pages 16–31. ACM Press, 1999.

[78] Yin Zhang, Matthew Roughan, Carsten Lund, and David Donoho. An Information-Theoretic Approach to Traffic Matrix Estimation. In *Proceedings of ACM SIGCOMM*, August 2003.

[79] Artur Ziviani, Serge Fdida, Jose de Rezende, and Otto Carlor Duarte. Toward a Measurement-based Geographic Location Service. In *Proceedings of the Passive and Active Measurement Workshop*, Juan Les Pins, France, April 2004.

[80] C. Zou, W. Gong, and D. Towsley. Worm propagation modeling and analysis under dynamic quarantine defense. In *Proceedings of the Workshop on Rapid Malcode (WORM)*, 2003.

[81] T. Zseby, M. Molina, F. Raspall, and N.G. Duffield. Sampling and filtering techniques for ip packet selection. Internet Draft, October 2003.