ADITYA AKELLA

aditya@cs.cmu.edu
http://www.cs.cmu.edu/∼aditya
Phone: 412-818-3779
Fax: 412-268-5576 (Attn: Aditya Akella)

Computer Science Department
Carnegie Mellon University
5000 Forbes Avenue
Pittsburgh, PA 15232

Education

PhD in Computer Science — May 2005 (expected)
Carnegie Mellon University, Pittsburgh, PA
Dissertation: "An Integrated Approach to Optimizing Internet Performance"
Advisor: Prof. Srinivasan Seshan

Bachelor of Technology in Computer Science and Engineering — May 2000
Indian Institute of Technology (IIT), Madras, India

Honors and Awards

IBM PhD Fellowship	2003-04 & 2004-05
Graduated third among all IIT Madras undergraduates	2000
Institute Merit Prize, IIT Madras	1996-97
19th rank, All India Joint Entrance Examination for the IITs	1996
Gold medals, Regional Mathematics Olympiad, India	1993-94 & 1994-95
National Talent Search Examination (NTSE) Scholarship, India	1993-94

Research Interests

Computer Systems and Internetworking

PhD Dissertation

An Integrated Approach to Optimizing Internet Performance

In my thesis research, I adopted a systematic approach to understand how to optimize the performance of well-connected Internet end-points, such as universities, large enterprises and data centers. I showed that constrained bottlenecks inside and between carrier networks in the Internet could limit the performance of such end-points. I observed that a clever end point-based strategy, called Multihoming Route Control, can help end-points route around these bottlenecks, and thereby improve performance. Furthermore, I showed that the Internet's topology, routing and trends in its growth may essentially worsen the wide-area congestion in the future. To this end, I proposed changes to the Internet's topology in order to guarantee good future performance.

Research Experience

Apart from my thesis research outlined above, I have worked on a variety of other projects, both at CMU and elsewhere. Some of these are outlined next.

Carnegie Mellon University — 2000–Present
Graduate Research Assistant with Prof. Srinivasan Seshan

DNS Characterization: With Jeffrey Pang and James Hendricks I conducted a measurement analysis of the deployment, usage and availability characteristics of servers in the Domain Name System. We

also explored the effectiveness of employing DNS as a mechanism for network-level redirection and control of end-clients.

Representative Internet Measurements: With Dr. Anees Shaikh, I worked on developing a framework for evaluating the representativeness of a given Internet measurement testbed, and the measurement observations drawn from the testbed.

DDoS Attacks on ISP Networks: With Ashwin Bharambe, I worked on extending algorithms developed by the data streaming community to compute signatures of large-scale DDoS attacks on ISP networks, and to efficiently and accurately identify them.

False Sharing in Congestion Management Systems: As part of the Congestion Manager project, I studied the impact of false sharing of congestion state among flows in a shared congestion management system. I developed fast and accurate algorithms to detect false sharing and respond to it.

Exploring Congestion Control: I worked with Prof. Scott Shenker and Prof. Ion Stoica on evaluating the performance of various linear congestion control algorithms under modern settings of efficient end-point loss recovery schemes and AQM-enhanced routers.

Realistic Selfish Routing: With Shuchi Chawla, I developed new, more realistic theoretical models for selfish routing. These models incorporated multihoming route control, long-lived TCP flows and new social objective functions (such as maximizing the fair-share bandwidth).

IBM T. J. Watson Research Center Summer 2003
Intern

Multihoming Route Control Implementation: I worked with Dr. Anees Shaikh on the design and implementation of simple, practical route control strategies for optimizing web performance in multihomed enterprise settings. I also identified best common route control practices for multihomed networks. The solutions I developed will soon be incorporated into IBM's zSeries mainframe servers.

ICSI Center for Internet Research (ICIR) Summer 2001
Intern

Selfish Congestion Control: I worked with Prof. Richard Karp, Dr. Scott Shenker and Prof. Christos Papadimitriou on analyzing the impact of selfish TCP congestion control behavior on the operation of the Internet. We showed the surprising result that the traditional setting of end-points employing TCP Reno and routers employing FIFO-Droptail buffering is robust to selfish congestion control behavior. However, recent advances in AQM schemes and TCP loss recovery mechanisms have made the network vulnerable to selfish behavior.

IBM India Research Lab Summer 1999
Intern

Capacity Auctions: I worked with Dr. Alok Agarwal and S. Ananthan on new models of auctions for e-commerce establishments. We analyzed the hardness of approximation of a class of combinatorial auctions called "capacity" auctions. We also developed efficient heuristics for such auctions.

Teaching Experience

CMU Course 15744, Graduate Computer Networks Spring 2004

Prof. Hui Zhang and I designed this course. I outlined the reading list for the class, designed the problems sets and exams, gave several lectures and held office hours.

CMU Course 15441, Undergraduate Computer Networks Spring 2002

As a teaching assistant for Prof. Peter Steenkiste and Prof. Mor Harchol-Balter, I helped with grading and designing exams. I also designed a month-long development project for the class, held office hours and gave a lecture.

IIT Madras Course CS110, Fundamentals of Programming Spring 2000

As a teaching assistant for Prof. C. Pandu Rangan and Prof. Hema Murthy, I helped with grading, held office hours and coordinated lab sessions.

Mentoring Experience

In the academic year 2002-03, I mentored Arvind Kannan's undergraduate senior thesis titled "Scaling Properties of the Internet Graph". For the academic year 2004-05, I am mentoring Sameer Moidu's senior thesis work on the the effectiveness of 802.11b rate adaptation algorithms. I advised a younger graduate student, Jeffrey Pang, during the first year of his graduate studies.

Conference Publications

(1) Jeffrey Pang, James Hendricks, Aditya Akella, Bruce Maggs, Roberto De Prisco and Srinivasan Seshan. "Availability, Usage, and Deployment Characteristics of the Domain Name System", *Internet Measurement Conference (IMC)*, October 2004.

(2) Jeffrey Pang, Aditya Akella, Anees Shaikh, Balachander Krishnamurthy and Srinivasan Seshan. "On the Responsiveness of DNS-based Network Control", *Internet Measurement Conference (IMC)*, October 2004.

(3) Aditya Akella, Jeffrey Pang, Anees Shaikh, Bruce Maggs and Srinivasan Seshan. "A Comparison of Overlay Routing and Multihoming Route Control", *ACM SIGCOMM*, August 2004.

(4) Aditya Akella, Anees Shaikh and Srinivasan Seshan. "Multihoming Performance Benefits: An Experimental Evaluation of Practical Enterprise Strategies", *USENIX Annual Technical Conference*, June 2004.

(5) Aditya Akella, Hari Balakrishnan and Srinivasan Seshan. "The Impact of False Sharing on Shared Congestion Management", *IEEE International Conference on Network Protocols (ICNP)*, November 2003.

(6) Aditya Akella, Srinivasan Seshan and Anees Shaikh. "An Empirical Evaluation of Wide-Area Internet Bottlenecks", *Internet Measurement Conference (IMC)*, October 2003.

(7) Aditya Akella, Bruce Maggs, Srinivasan Seshan, Anees Shaikh and Ramesh Sitaraman. "A Measurement-Based Analysis of Multihoming", *ACM SIGCOMM*, August 2003.

(8) Aditya Akella, Shuchi Chawla, Arvind Kannan and Srinivasan Seshan. "Scaling properties of the Internet Graph", *ACM Principles of Distributed Computing (PODC)*, July 2003.

(9) Aditya Akella, Richard Karp, Christos Papadimitriou, Srinivasan Seshan and Scott Shenker. "Selfish Behavior and Stability of the Internet: A Game-Theoretic Analysis of TCP", *ACM SIGCOMM*, August 2002.

Journal Articles

(10) Aditya Akella, Shuchi Chawla, Arvind Kannan and Srinivasan Seshan, "On the Scaling of Congestion in the Internet Graph", *ACM SIGCOMM Computer Communication Review, Vol. 34, No.3, Special Issue on Science of Network Design*, pp 43–55, July 2004.

Selected Technical Reports

(11) Aditya Akella, Ashwin Bharambe, Suman Nath and Srinivasan Seshan. "Multi-Modal Network Protocols: Adapting to Highly Variable Operating Conditions", CMU Technical Report CMU-CS-02-170.

(12) Aditya Akella, Shuchi Chawla and Srinivasan Seshan. "Mechanisms for Internet Routing: A Study", CMU Technical Report CMU-CS-02-163.

(13) Aditya Akella, Srinivasan Seshan, Scott Shenker and Ion Stoica. "Exploring Congestion Control", CMU Technical Report CMU-CS-02-139.

Selected Short Papers and Posters

(14) Aditya Akella. "Understanding the Impact of Route Control Product Deployment", Invited paper, *Workshop on Internet Routing Evolution and Design*, October 2003.

(15) Aditya Akella, Srinivasan Seshan and Anees Shaikh. "Toward Representative Internet Measurements", *3rd New York Metro Area Networking Workshop*, September 2003.

(16) Aditya Akella, Srinivasan Seshan and Anees Shaikh, "An Empirical Evaluation of Wide-Area Internet Bottlenecks", *ACM SIGMETRICS*, June 2003.

(17) Aditya Akella, Ashwin Bharambe, Mike Reiter and Srinivasan Seshan. "Detecting DDoS Attacks on ISP Networks", *SIGMOD Workshop on Management and Processing of Data Streams*, June 2003.

(18) Aditya Akella, Rajesh Balan, Nikhil Bansal and Srinivasan Seshan. "Multi-modal Network Protocols", Poster at *ACM SIGCOMM*, August 2001.

Selected Invited Talks

"Routing Strategies for Improving End-to-End Performance and Reliability", *UC Berkeley*, April 2004.

"A Comparison of Overlay Routing and Multihoming Route Control", *IBM T.J. Watson Research Center*, March 2004.

"End-to-End Routing: Today and Tomorrow", *Workshop on Internet Routing Evolution and Design*, October 2003.

Service and Activities

External reviewer for SIGCOMM 2002–04, SIGMETRICS 2003, INFOCOM 2002–05, USENIX 2004, MMCN 2003, IMC 2003 & 2004 and IEEE Transactions on Networking.
Tutorial organization volunteer for SIGCOMM 2002 held at Pittsburgh, PA.
Member of ACM, SIGCOMM, IEEE and USENIX.

References

Prof. Srinivasan Seshan
Computer Science Department
Carnegie Mellon University
E-mail: srini+@cs.cmu.edu
Phone: +1 (412) 268-8734

Prof. Bruce Maggs
Computer Science Department
Carnegie Mellon University
E-mail: bmm+@cs.cmu.edu
Phone: +1 (412) 268-7654

Prof. Hui Zhang
Computer Science Department
Carnegie Mellon University
E-mail: hzhang+@cs.cmu.edu
Phone: +1 (412) 268-8945

Prof. Richard Karp
Computer Science Department
University of California, Berkeley
E-mail: karp@cs.berkeley.edu
Phone: +1 (510) 642-5799

Prof. Scott Shenker
Computer Science Department
University of California, Berkeley
E-mail: shenker@eecs.berkeley.edu
Phone: +1 (510) 666-2880

Selected Paper Abstracts

Selfish Behavior and Stability of the Internet: A Game-Theoretic Analysis of TCP
Aditya Akella, Richard Karp, Srinivasan Seshan, Scott Shenker and Christos Papadimitriou
ACM SIGCOMM 2002

For years, the conventional wisdom has been that the continued stability of the Internet depends on the widespread deployment of "socially responsible" congestion control. In this paper, we seek to answer the following fundamental question: If network end-points behaved in a selfish manner, would the stability of the Internet be endangered?

We evaluate the impact of greedy end-point behavior through a game-theoretic analysis of TCP. In this "TCP Game" each flow attempts to maximize the throughput it achieves by modifying its congestion control behavior. We use a combination of analysis and simulation to determine the Nash Equilibrium of this game. Our question then reduces to whether the network operates efficiently at these Nash equilibria.

Our findings are twofold. First, in more traditional environments – where end-points use TCP Reno-style loss recovery and routers use drop-tail queues – the Nash Equilibria are reasonably efficient. However, when endpoints use more recent variations of TCP (e.g., SACK) and routers employ either RED or drop-tail queues, the Nash equilibria are very inefficient. This suggests that the Internet of the past could remain stable in the face of greedy end-user behavior, but the Internet of today is vulnerable to such behavior. Second, we find that restoring the efficiency of the Nash equilibria in these settings does not require heavy-weight packet scheduling techniques (e.g., Fair Queuing) but instead can be done with a very simple stateless mechanism based on CHOKe.

Scaling Properties of the Internet Graph
Aditya Akella, Shuchi Chawla, Arvind Kannan and Srinivasan Seshan
ACM PODC 2003

As the Internet grows in size, it becomes crucial to understand how the speeds of links in the network must improve in order to sustain the pressure of new end-nodes being added each day. Although the speeds of links in the core and at the edges improve roughly according to Moore's law, this improvement alone might not be enough. Indeed, the structure of the Internet graph and routing in the network might necessitate much faster improvements in the speeds of key links in the network.

In this paper, using a combination of analysis and extensive simulations, we show that the worst congestion in the Internet's AS-level graph in fact scales poorly with the network size ($n^{1+\Omega(1)}$, where n is the number of nodes), when shortest-path routing is used between ASes. We also show, somewhat surprisingly, that policy-based routing does not exacerbate the maximum congestion when compared to shortest-path routing.

Our results show that it is crucial to identify ways to alleviate this congestion to avoid some links from being perpetually congested. To this end, we show that the congestion scaling properties of the Internet graph can be improved dramatically by introducing moderate amounts of redundancy in the graph (e.g., adding multiple parallel edges between pairs of adjacent nodes).

A Measurement-Based Analysis of Multihoming
Aditya Akella, Bruce Maggs, Srinivasan Seshan, Ramesh Sitaraman and Anees Shaikh
ACM SIGCOMM 2003

Multihoming has traditionally been employed by stub networks to enhance the reliability of their network connectivity. With the advent of commercial "intelligent route control" products, stubs now leverage multihoming to improve performance. Although multihoming is widely used for reliability and, increasingly for performance, not much is known about the tangible

benefits that multihoming can offer, or how these benefits can be fully exploited. In this paper, we aim to quantify the extent to which multihomed networks can leverage performance and reliability benefits from connections to multiple providers. We use data collected from servers belonging to the Akamai content distribution network to evaluate performance benefits from two distinct perspectives of multihoming: high-volume content-providers which transmit large volumes of data to many distributed clients, and enterprises which primarily receive data from the network. In both cases, we find that multihoming can improve performance significantly and that not choosing the right set of providers could result in a performance penalty as high as 40%. We also find evidence of diminishing returns in performance when more than four providers are considered for multihoming. In addition, using a large collection of measurements, we provide an analysis of the reliability benefits of multihoming. Finally, we provide guidelines on how multihomed networks can choose ISPs, and discuss initial practical strategies of using multiple upstream connections to achieve optimal performance benefits.

An Empirical Evaluation of Wide-Area Internet Bottlenecks
Aditya Akella, Srinivasan Seshan and Anees Shaikh
IMC 2003

Conventional wisdom has been that the performance limitations in the current Internet lie at the edges of the network – *i.e* last mile connectivity to users, or access links of stub ASes. As these links are upgraded, however, it is important to consider where new bottlenecks and hotspots are likely to arise. In this paper, we address this question through an investigation of *non-access* bottlenecks. These are links within carrier ISPs or between neighboring carriers that could *potentially* constrain the bandwidth available to long-lived TCP flows. Through an extensive measurement study, we discover, classify, and characterize bottleneck links (primarily in the U.S.) in terms of their location, latency, and available capacity.

We find that about 50% of the Internet paths explored have a non-access bottleneck with available capacity less than 50 Mbps, many of which limit the performance of well-connected nodes on the Internet today. Surprisingly, the bottlenecks identified are roughly equally split between intra-ISP links and peering links between ISPs. Also, we find that low-latency links, both intra-ISP and peering, have a significant likelihood of constraining available bandwidth. Finally, we discuss the implications of our findings on related issues such as choosing an access provider and optimizing routes through the network. We believe that these results could be valuable in guiding the design of future network services, such as overlay routing, in terms of which links or paths to avoid (and how to avoid them) in order to improve performance.

The Impact of False Sharing on Shared Congestion Management
Aditya Akella, Hari Balakrishnan and Srinivasan Seshan
ICNP 2003

Several recent proposals for sharing congestion information across concurrent flows between end-systems overlook an important problem: two or more flows sharing congestion state may in fact not share the same bottleneck. In this paper, we categorize the origins of this *false sharing* into two distinct cases: (i) networks with QoS enhancements such as differentiated services, where a flow classifier segregates flows into different queues, and (ii) networks with path diversity where different flows to the same destination address are routed differently. We evaluate the impact of false sharing on flow performance and investigate how false sharing can be detected by a sender. We discuss how a sender must respond upon detecting false sharing. Our results show that persistent overload can be avoided with window-based congestion control even for extreme false sharing, but higher bandwidth flows run at a slower rate. We find that delay and reordering statistics can be used to develop robust detectors of false sharing and are superior to those based on loss patterns. We also find that it is markedly easier to detect and react to false sharing than it is to start by isolating flows and merge their congestion state afterward.

Multihoming Performance Benefits: An Experimental Evaluation of Practical Enterprise Strategies
Aditya Akella, Srinivasan Seshan and Anees Shaikh
USENIX 2004

Multihoming is increasingly being employed by large enterprises and data centers as a mechanism to extract good performance from their provider connections. Today, multihomed end-networks can employ a variety of commercial *route control* products to optimize performance over multiple ISP links. However, little is known about the mechanisms employed by such products and their relative trade-offs.

In this paper, we propose and evaluate a wide range practical schemes that could go into the design of a route control device and analyze their trade-offs. We implement the proposed schemes on a Linux-based Web proxy and perform a trace-based emulation of their relative performance benefits. We show that both passive and active monitoring based techniques are equally effective and could improve Web performance by about 25% when compared to using a single provider. Another key observation is that the conventional practice of employing historical measurement samples to monitor and predict ISP performance could, in fact, result in sub-optimal performance.

A Comparison of Overlay Routing and Multihoming Route Control
Aditya Akella, Jeffrey Pang, Bruce Maggs, Srinivasan Seshan and Anees Shaikh
ACM SIGCOMM 2004

The limitations of BGP routing in the Internet are often blamed for poor end-to-end performance and prolonged connectivity interruptions. Recent work advocates using overlays to effectively bypass BGP's path selection in order to improve performance and fault tolerance. In this paper, we explore the possibility that intelligent control of BGP routes, coupled with ISP multihoming, can provide competitive end-to-end performance and reliability. Using extensive measurements of paths between nodes in a large content distribution network, we compare the relative benefits of overlay routing and multihoming route control in terms of round-trip latency, TCP connection throughput, and path availability. We observe that the performance achieved by route control together with multihoming to three ISPs (3-multihoming), is within 3–12% of overlay routing employed in conjunction 3-multihoming, in terms of both end-to-end RTT and throughput. We also show that while multihoming cannot offer the nearly perfect resilience of overlays, it can eliminate almost all failures experienced by a singly-homed end-network. Our results demonstrate that, by leveraging the capability of multihoming route control, it is not necessary to circumvent BGP routing to extract good wide-area performance and availability from the existing routing system.

Availability, Usage, and Deployment Characteristics of the Domain Name System
Jeffrey Pang, James Hendricks, Aditya Akella, Bruce Maggs, Roberto De Prisco and Srinivasan Seshan
IMC 2004

The Domain Name System (DNS) is a critical part of the Internet's infrastructure, and is one of the few examples of a robust, highly-scalable, and operational distributed system. Although a few studies have been devoted to characterizing its properties, such as its workload and the stability of the top-level servers, many key components of DNS have not yet been examined. Based on large-scale measurements taken from servers in a large content distribution network, we present a detailed study of key characteristics of the DNS infrastructure, such as load distribution, availability, and deployment patterns of DNS servers. Our analysis includes both local DNS servers and servers in the authoritative hierarchy. We find that (1) the vast majority of users use a small fraction of deployed name servers, (2) the availability of most name servers is high, and (3) there exists a larger degree of diversity in local DNS server deployment and usage than for authoritative servers. Furthermore, we use our DNS measurements to draw conclusions about federated infrastructures in general. We evaluate and discuss the impact of federated deployment models on future systems, such as Distributed Hash Tables.