

Measuring the Internet by Coordinating Distributed Vantage Points *

Ethan Katz-Bassett
Department of Computer Science
University of Washington, Seattle
ethan@cs.washington.edu

ABSTRACT

Network operators and researchers rely on traceroute to measure Internet paths, monitor performance, and localize faults. This approach is not ideal, as traceroute, though extremely useful, is inherently limited in the data it can provide. We have designed practical techniques that dynamically coordinate distributed vantage points to yield substantially better measurements for the types of data currently provided by traceroute. To provide a richer view than is available with existing tools, our measurement techniques take advantage of topological relationships vantage points have to each other and to destinations. We propose to empirically demonstrate that these techniques work on the Internet without special support and often provide ideal information in scenarios in which people now depend on traceroute’s limited information. We also describe applications we plan to build that utilize our novel measurement techniques.

1. BACKGROUND AND MOTIVATION

As we depend on the Internet more and more, it is critically important that it functions well. However, Internet communication depends on complex interactions between multiple protocols and multiple autonomous domains, making it difficult to guarantee high performance or understand the causes of problems. With limited visibility or control outside their local network, network operators and researchers make measurements to monitor the Internet and attempt to understand the wide area. Network operators rely on the traceroute tool to ensure the proper functioning of their networks and to troubleshoot problems. Researchers use traceroute measurements to study properties of the Internet and to build useful systems. Traceroute, which measures the router-level path from the source to a specified destination, is the “number one go-to tool” for troubleshooting problems [13]. Operators and researchers use traceroute to determine paths [10, 8], locations of failures [10, 14], laten-

cies [11, 9], and (via DNS names) geographic locations of routers [10, 9].

Although traceroute is an extremely useful measurement tool, it does not provide complete information, and the limits on what it provides restrict the accuracy of inferences made with it. As networks and operators have become more sophisticated, few problems encountered are simple to diagnose and resolve [13], and existing diagnostic techniques no longer suffice given the more complicated problems that dominate. The majority of Internet communication is bidirectional, and most paths are asymmetric [2], yet traceroute measures only the forward path, forcing systems into the unrealistic assumption of path symmetry [8, 15]. One cannot differentiate forward and reverse path failures using traceroute, limiting our ability to diagnose Internet anomalies [14, 15]. Traceroute measures the round-trip latency to each hop, but many systems want link or one-way latencies [9, 8], which are difficult to estimate because of path asymmetry. Many IP addresses do not have DNS names with geographic meaning [3], forcing systems to rely on limited data [8, 9]. Ideally, traceroute would give (1) the forward and reverse path, (2) the link causing any failure, (3) the latency of every link, and (4) each hop’s geographic location.

2. TECHNIQUES

We propose techniques for these four types of data, in some cases approaching the ideal. We focus on building real systems that work on today’s Internet. Our reverse traceroute system provides the same basic information as traceroute – IP-address-level hops along the path, plus round-trip delay to each – but along the reverse path from the destination back to the source. First, we employ the IP timestamp and record route options to identify hops along the reverse path. Second, we combine the view of multiple vantage points to infer information unavailable from any single one. Third, we use limited source spoofing – spoofing from one vantage point as another – to use the vantage point best positioned to make the measurement. Because the spoofed source address is always one of our hosts, we avoid issues of concealment that can arise with spoofing, and this controlled spoofing allows us to overcome many of the limitations inherent in using IP options. Using our current PlanetLab-based deployment, we find that, in the median (mean) case, our reverse traceroute technique reveals 87% (83%) of the routers and 100% (94%) of the points-of-presence (PoPs) on a traceroute issued from the destination. We are building a tool that allows users to use our system and serve as vantage points for others’ mea-

*This extended abstract describes the author’s work-in-progress thesis research.

measurements, increasing the coverage, and a number of operators are supporting our work as beta testers [4].

In ongoing work, we are using our reverse traceroute system to provide better failure, latency, and geolocation information. We propose to use ideas from our reverse traceroute to isolate failures to either the forward or reverse path and provide detailed information about the location of the failures. We propose to calculate individual link latencies by using traceroute and our reverse traceroute to constrain latencies of links observed on round-trip paths. Expanding on our earlier topology-based geolocation technique [5], we propose to estimate the geographic locations of arbitrary Internet hosts by using our reverse traceroute and link latency techniques, along with other topological measurements, to constrain the hosts' feasible locations.

3. APPLICATIONS

Path asymmetry is “the number one plague of traceroute” [13]; our reverse traceroute system's ability to overcome this challenge and measure reverse paths will enable us to conduct novel studies of Internet routing and to deploy real systems that aid operators in troubleshooting actual problems.

First, we propose to conduct the first Internet-wide study of the routing policies used by multihomed edge networks. Without a way to measure paths from most edge networks, previous work studied either routing at the core [12] or from a small number of networks [7]. Recent route control products enable multihomed networks to optimize performance by dynamically selecting providers, and we will study the behaviors networks exhibit in practice.

Second, we built *Hubble*, a system to identify Internet reachability problems and locate their likely sources [6]. *Hubble* vantage points coordinate to cover the entire Internet, a much larger scale than existing systems and studies [10, 1, 14]. *Hubble* uses novel failure localization techniques to characterize the nature of many failures. For example, in many cases, a multihomed AS is reachable through one provider, but probes through another terminate; using spoofed packets, we isolated the direction of failure in 84% of cases and found all problems to be exclusively because the provider was not forwarding traffic to the destination.

Finally, we propose to work with collaborators at a major content provider to use our reverse traceroute and link latency techniques to help them troubleshoot poor client performance. They found that 22% of client prefixes experience more than 50ms latency over the minimum latency to the prefix's geographical region. Our collaborators want a way to point to an AS as the cause of inflation, but are hindered by the lack of information about reverse paths back to their servers from clients.

4. SUMMARY

Network operators and researchers rely on traceroute for path, failure location, latency, and geographic information. Using practical techniques that dynamically coordinate probes sent from multiple vantage points, we can provide substantially better information than is obtainable using traceroutes. Our proposed techniques will allow us to build applications and conduct studies with dramatically better accu-

racy and coverage than existing tools provide. The primary contributions we hope to make are:

- The design and implementation of practical techniques to provide reverse path, failure localization, link latency, and geographic location information.
- A demonstration of the fact that applications of Internet measurements can benefit from using our techniques, including:
 - The first detailed study of routing policies in place at multihomed edge networks.
 - *Hubble*, a real-time system to monitor and classify reachability problems and black holes across the entire Internet.
 - A system to aid a content provider in diagnosing the cause of poor client performance.

5. ACKNOWLEDGMENTS

This research is supported by grants from Google, Cisco, and the National Science Foundation.

6. REFERENCES

- [1] D. Andersen, H. Balakrishnan, F. Kaashoek, and R. Morris. Resilient overlay networks. In *SOSP*, 2001.
- [2] Y. He, M. Faloutsos, S. Krishnamurthy, and B. Huffaker. On routing asymmetry in the Internet. In *Autonomic Networks Symposium in Globecom*, 2005.
- [3] iPlane. <http://iplane.cs.washington.edu>.
- [4] E. Katz-Bassett. Practical reverse traceroute. In *NANOG 45*, 2009. http://www.nanog.org/meetings/nanog45/presentations/Tuesday/Katz_reversetraceroute_N45.pdf.
- [5] E. Katz-Bassett, J. P. John, A. Krishnamurthy, D. Wetherall, T. Anderson, and Y. Chawathe. Towards IP geolocation using delay and topology measurements. In *IMC*, 2006.
- [6] E. Katz-Bassett, H. V. Madhyastha, J. P. John, A. Krishnamurthy, D. Wetherall, and T. Anderson. Studying black holes in the Internet with Hubble. In *NSDI*, 2008.
- [7] S. Lee, Z.-L. Zhang, and S. Nelakuditi. Exploiting AS hierarchy for scalable route selection in multi-homed stub networks. In *IMC*, 2004.
- [8] H. V. Madhyastha, T. Isdal, M. Piatek, C. Dixon, T. Anderson, A. Krishnamurthy, and A. Venkataramani. iPlane: An information plane for distributed services. In *OSDI*, 2006.
- [9] R. Mahajan, M. Zhang, L. Poole, and V. Pai. Uncovering performance differences among backbone ISPs with Netdiff. In *NSDI*, 2008.
- [10] V. Paxson. End-to-end routing behavior in the Internet. *IEEE/ACM TON*, 5(5):601–615, 1997.
- [11] S. Savage, A. Collins, E. Hoffman, J. Snell, and T. Anderson. The end-to-end effects of Internet path selection. In *SIGCOMM*, 1999.
- [12] N. Spring, R. Mahajan, and T. Anderson. Quantifying the causes of path inflation. In *SIGCOMM*, 2003.
- [13] R. A. Steenberg. A practical guide to (correctly) troubleshooting with traceroute. In *NANOG 45*, 2009. http://www.nanog.org/meetings/nanog45/presentations/Sunday/RAS_traceroute_N45.pdf.
- [14] M. Zhang, C. Zhang, V. Pai, L. Peterson, and R. Wang. PlanetSeer: Internet path failure monitoring and characterization in wide-area services. In *OSDI*, 2004.
- [15] Z. Zhang, Y. Zhang, Y. C. Hu, Z. M. Mao, and R. Bush. iSpy: detecting IP prefix hijacking on my own. In *SIGCOMM*, 2008.