

CMSC 818S: Fall 2005 Reading List

Neil Spring

October 24, 2005

This reading list is subject to change. Networking has three large and interesting conferences/workshops in the fall: SIGCOMM in August, IMC in October, and HotNets in November. These proceedings are likely to have new and discussion-sparking material.

Some papers appear more than once—they fit into the theme for more than one “week.” There are more than 15 weeks—we won’t cover it all, we’ll pick and chose what’s most interesting or most worth talking about. There’s no designed order except at the beginning—I’ll have you sign up for a “week” and order the weeks that have been claimed, then we’ll go through topics as the time takes us.

Unit 0: First Class – Neil

No readings. However, be prepared to start reading a textbook to brush up on vocabulary.

Unit 1: Big Picture – Neil

I reuse the “Reverse Engineering” title often. This is a big-picture survey or position paper that argues that we (researchers) could, if we tried, measure everything of interest on every link of the Internet every day. The “Looking over the fence” report describes networking research at a high level, and puts network measurement into context: what we know and don’t, what the roadblocks are, and why it matters.

Readings

- [1] Neil Spring, David Wetherall, and Thomas Anderson. Reverse-engineering the Internet. In *Proceedings of the ACM Workshop on Hot Topics in Networks (HotNets)*, pages 3–8. Cambridge, MA, November 2003. URL <http://www.cs.umd.edu/~nspring/papers/reverse.pdf>.
- [2] Computer Science and Telecommunications Board, National Research Council. *Looking Over the Fence at Networks: A Neighbor’s View of Networking Research*. The National Academies Press, 2001. Document is pages 13–26, URL http://books.nap.edu/html/looking_over_the_fence/report.pdf.

Unit 2: Big Picture II – Neil

The “crisis conditions” paper motivates measuring and understanding the network by the expectation that we will see more, more severe attacks upon the communication network. The failures imply open research topics, the successes (in news distribution, for example) imply faulty motivation for some research. The story is an interesting one.

The “knowledge plane” is a (position paper) attempt to address open issues in network design—primarily that networks should manage themselves—by combining measurement work with artificial intelligence techniques.

The “better models” paper is helpful in guiding researchers looking to evaluate proposals convincingly. I believe the “IAB concerns” writeup is similar in spirit. The Internet Architecture Board is a part of the IETF (Internet Engineering Task Force, which is the standards body involved in Internet protocols) that manages high-level decisions, much like a board of directors of a company.

Readings

- [1] Computer Science and Telecommunications Board, National Research Council. *The Internet Under Crisis Conditions: Learning from September 11*. The National Academies Press, 2003. URL <http://books.nap.edu/books/0309087023/html/R1.html>.
- [2] David D. Clark, Craig Partridge, J. Christopher Ramming, and John T. Wroclawski. A knowledge plane for the Internet. In *Proceedings of the ACM SIGCOMM Conference on Applications, Technologies, Architectures, and*

Protocols for Computer Communication (SIGCOMM), pages 3–10. Karlsruhe, Germany, August 2003. Position paper, URL <http://www.acm.org/sigcomm/sigcomm2003/papers/p3-clark.pdf>.

- [3] Sally Floyd and Eddie Kohler. Internet research needs better models. In *Proceedings of the ACM Workshop on Hot Topics in Networks (HotNets)*, pages 29–34. Princeton, NJ, October 2002. URL <http://www.icir.org/models/hotnetsFinal.pdf>.
- [4] Ran Atkinson and Sally Floyd, editors. IAB concerns and recommendations regarding Internet research and evolution. Internet Engineering Task Force Request for Comments RFC-3869, August 2004. URL <http://www.icir.org/floyd/papers/rfc3869.ps>.

Unit 3: Fundamental Tools – Neil

This week will have little reading, in part because the basic tools we use often come from operators or from practice, or are in some other way not research. However, we will likely have some exercises in using a few basic tools.

Oops. I wanted [5] instead of nettimer.

Readings

- [1] Allen B. Downey. Using pathchar to estimate Internet link characteristics. In *Proceedings of the ACM SIGCOMM Conference on Applications, Technologies, Architectures, and Protocols for Computer Communication (SIGCOMM)*, pages 241–250. Cambridge, MA, September 1999. URL <http://allendowney.com/research/clink/downey.ps.gz>.
- [2] Manish Jain and Constantinos Dovrolis. End-to-end available bandwidth: measurement methodology, dynamics, and relation with TCP throughput. In *Proceedings of the ACM SIGCOMM Conference on Applications, Technologies, Architectures, and Protocols for Computer Communication (SIGCOMM)*, pages 295–308. Pittsburgh, PA, August 2002. URL <http://www.cc.gatech.edu/~jain/publications/sigcomm02.pdf>.
- [3] Kevin Lai and Mary Baker. Nettimer: A tool for measuring bottleneck link bandwidth. In *Proceedings of the USENIX Symposium on Internet Technologies and Systems (USITS)*, pages 123–134. San Francisco, CA, March 2001. URL http://mosquitonet.stanford.edu/~laik/projects/nettimer/publications/usits2001/a_tool_for_measuring_bottleneck_link_bandwidth.pdf.
- [4] Rick Jones. Netperf. <http://www.netperf.org/>.
- [5] Kevin Lai and Mary Baker. Measuring link bandwidths using a deterministic model of packet delay. In *Proceedings of the ACM SIGCOMM Conference on Applications, Technologies, Architectures, and Protocols for Computer Communication (SIGCOMM)*. Stockholm, Sweden, August 2000.

Unit 4: Neato – Dave – 9/29

Some tools are, I think, a bit more creative than others. They might not always work, but they make you think.

Read first two.

Readings

- [1] Stefan Savage. Sting: a TCP-based network measurement tool. In *Proceedings of the USENIX Symposium on Internet Technologies and Systems (USITS)*, pages 71–79. Boulder, CO, October 1999. URL <http://www.cs.ucsd.edu/~savage/papers/Usits99.pdf>.
- [2] Ningning Hu, Oliver Spatscheck, Jia Wang, and Peter Steenkiste. Locating Internet bottlenecks: Algorithms, measurements, and implications. In *Proceedings of the ACM SIGCOMM Conference on Applications, Technologies, Architectures, and Protocols for Computer Communication (SIGCOMM)*. Portland, OR, August 2004. URL <http://gs274.sp.cs.cmu.edu/www/papers/bottleneck.pdf>.
- [3] Krishna P. Gummadi, Stefan Saroiu, and Steven D. Gribble. King: Estimating latency between arbitrary Internet end hosts. In *Proceedings of the ACM SIGCOMM Internet Measurement Workshop (IMW)*, pages 5–18. Marseille, France, November 2002. URL www.cs.washington.edu/homes/gummadi/papers/king.pdf.

- [4] Ramesh Govindan and Vern Paxson. Estimating router ICMP generation delays. In *Proceedings of Passive & Active Measurement (PAM)*, pages 6–13. Fort Collins, CO, March 2002. URL <http://www.icir.org/vern/papers/fsd-pam-02.pdf>.

Unit 5: Mapping – Nick – 10/4

Network mapping is the process of using tools to determine the connectivity between routers, between autonomous systems, and sometimes between other abstract network concepts (links in a web graph, for example).

Read first two.

Readings

- [1] Neil Spring, Ratul Mahajan, and David Wetherall. Measuring ISP topologies with Rocketfuel. In *Proceedings of the ACM SIGCOMM Conference on Applications, Technologies, Architectures, and Protocols for Computer Communication (SIGCOMM)*, pages 133–146. Pittsburgh, PA, August 2002. URL <http://www.cs.umd.edu/~nspring/papers/sigcomm2002.pdf>.
- [2] Ramesh Govindan and Hongsuda Tangmunarunkit. Heuristics for Internet map discovery. In *Proceedings of the IEEE Joint Conference of the IEEE Computer and Communications Societies (INFOCOM)*, pages 1371–1380. Tel Aviv, Israel, March 2000. URL http://www.isi.edu/div7/publication_files/heuristics.pdf.
- [3] Benoit Donnet, Philippe Raoult, Timur Friedman, and Mark Crovella. Efficient algorithms for large-scale topology discovery. In *Proceedings of the ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Systems*. Banff, Canada, June 2005. URL <http://www.cs.bu.edu/faculty/crovella/paper-archive/sigm05-trathome.pdf>.

Unit 6: Mapping + Inference – Mike S. 10/6 and 10/11

Routing policy is as important as network topology in determining which available paths will be used. The research described this week focuses on determining how to understand, or at least estimate, low-level routing configuration (policy).

Read the first two.

Readings

- [1] Lixin Gao. On inferring autonomous system relationships in the Internet. *IEEE/ACM Transactions on Networking*, 9(6):733–745, December 2001. URL <http://www-unix.ecs.umass.edu/~lgao/ton.ps>.
- [2] Neil Spring, Ratul Mahajan, and Thomas Anderson. Quantifying the causes of path inflation. In *Proceedings of the ACM SIGCOMM Conference on Applications, Technologies, Architectures, and Protocols for Computer Communication (SIGCOMM)*, pages 113–124. Karlsruhe, Germany, August 2003. URL <http://www.cs.umd.edu/~nspring/papers/pathpolicy.pdf>.
- [3] Ratul Mahajan, Neil Spring, David Wetherall, and Thomas Anderson. Inferring link weights using end-to-end measurements. In *Proceedings of the ACM SIGCOMM Internet Measurement Workshop (IMW)*, pages 231–236. Marseille, France, November 2002. URL <http://www.cs.umd.edu/~nspring/papers/imw2002.pdf>.

Unit 7: When Mapping Goes Bad – (Neil) 10/11 and 10/13

I split this from mapping because it describes a somewhat new and interesting realization that we cannot properly sample topology. yet.

Readings

- [1] Michalis Faloutsos, Petros Faloutsos, and Christos Faloutsos. On power-law relationships of the Internet topology. In *Proceedings of the ACM SIGCOMM Conference on Applications, Technologies, Architectures, and Protocols for Computer Communication (SIGCOMM)*, pages 251–262. Cambridge, MA, September 1999. URL <http://www.cs.ucr.edu/~michalis/PAPERS/FFF.pdf>.
- [2] Anukool Lakhina, John Byers, Mark Crovella, and Peng Xie. Sampling biases in IP topology measurements. In *Proceedings of the IEEE Joint Conference of the IEEE Computer and Communications Societies (INFOCOM)*, pages 332–341. San Francisco, CA, April 2003. URL <http://cs-people.bu.edu/anukool/pubs/samplingbias-infocom03.pdf>.

Unit 8: Kernel stuff – Mike F. 10/18

Precise packet timing and precise probe transmission are important for the accuracy of active network measurements. Passive measurements rely on fast processing in the kernel of a small interpreted bytecode language, BPF. Finally, sometimes tools have to keep the kernel out of the way so that it does not send packets in response to measurement traffic.

Readings

- [1] Khaled Harfoush, Azer Bestavros, and John Byers. PeriScope: An active measurement API. In *Proceedings of Passive & Active Measurement (PAM)*, pages 230–241. Fort Collins, CO, March 2002. URL <http://www.cs.bu.edu/fac/byers/pubs/PAM2002.ps>.
- [2] Steven McCanne and Van Jacobson. The BSD packet filter: A new architecture for user-level packet capture. In *USENIX Winter Technical Conference*. January 1993. URL <http://www.tcpdump.org/papers/bpf-usenix93.pdf>.
- [3] Ravi Prasad, Manish Jain, and Constantinos Dovrolis. Effects of interrupt coalescence on network measurements. In *Proceedings of Passive & Active Measurement (PAM)*. Antibes Juan-les-Pins, France, April 2004. URL http://www.cc.gatech.edu/fac/Constantinos.Dovrolis/Papers/intcoal_pam04.pdf.
- [4] Stefan Savage. Sting: a TCP-based network measurement tool. In *Proceedings of the USENIX Symposium on Internet Technologies and Systems (USITS)*, pages 71–79. Boulder, CO, October 1999. URL <http://www.cs.ucsd.edu/~savage/papers/Usits99.pdf>.

Unit 9: Passive – Randy 10/20

Passive measurement is practically half of measurement work, but occupies only a small portion of my head. A problem with packet trace analysis is the privacy of those whose packets are inspected. Once their privacy is preserved (or the privacy of the institution embarrassed by their employees’ inappropriate web surfing), little more can be understood about the network than the prevalence of some properties.

Readings

- [1] Venkata N. Padmanabhan, Lili Qiu, and Helen J. Wang. Server-based inference of Internet performance. In *Proceedings of the IEEE Joint Conference of the IEEE Computer and Communications Societies (INFOCOM)*, pages 145–155. San Francisco, CA, April 2003. URL <http://www.research.microsoft.com/~padmanab/papers/infocom2003.pdf>.
- [2] Jun Xu, Jinliang. Fan, Mostafa Ammar, and Sue Moon. Prefix-preserving IP address anonymization: Measurement-based security evaluation and a new cryptography-based scheme. In *Proceedings of the IEEE International Conference on Network Protocols (ICNP)*. Paris, France, November 2002.
- [3] Ruoming Pang and Vern Paxson. A high-level programming environment for packet trace anonymization and transformation. In *Proceedings of the ACM SIGCOMM Conference on Applications, Technologies, Architectures, and Protocols for Computer Communication (SIGCOMM)*, pages 339–351. Karlsruhe, Germany, August 2003. URL <http://portal.acm.org/citation.cfm?id=863994>.

Unit 10: DNS – Gleneesha – 10/25

DNS measurements are not uncommon, and they can be important because they understand a large, decentralized system that has a central role and nice properties if done well. (there are more papers, some “lint”-like ones from UCLA, some early ones, CAIDA has some stuff, the “anycast” DNS servers can be interesting).

Readings

- [1] Jaeyeon Jung, Emil Sit, Hari Balakrishnan, and Robert Morris. DNS performance and the effectiveness of caching. *IEEE/ACM Transactions on Networking*, 10(5), October 2002. Originally in IMW 2001, URL <http://nms.lcs.mit.edu/papers/dns-ton2002.pdf>.
- [2] Krishna P. Gummadi, Stefan Saroiu, and Steven D. Gribble. King: Estimating latency between arbitrary Internet end hosts. In *Proceedings of the ACM SIGCOMM Internet Measurement Workshop (IMW)*, pages 5–18. Marseille, France, November 2002. URL www.cs.washington.edu/homes/gummadi/papers/king.pdf.

- [3] Vasileios Pappas, Zhiguo Xu, Songwu Lu, Daniel Massey, Andreas Terzis, and Lixia Zhang. Impact of configuration errors on DNS robustness. In *Proceedings of the ACM SIGCOMM Conference on Applications, Technologies, Architectures, and Protocols for Computer Communication (SIGCOMM)*, pages 319–330. Portland, OR, August 2004. URL http://www.cs.ucla.edu/~vpappas/p/vp_sigcomm04.pdf.

Unit 11: Faking it with coordinates – Skylar

Attaching a “coordinate” in a multidimensional space to each of many hosts allows the latency between hosts to be predicted with surprising accuracy (most of the time) as the distance between coordinates. How does this work? Why does it work? Could we do better? Does the approach generalize?

Readings

- [1] T. S. Eugene Ng and Hui Zhang. Predicting Internet network distance with coordinates-based approaches. In *Proceedings of the IEEE Joint Conference of the IEEE Computer and Communications Societies (INFOCOM)*, volume 1, pages 170–179. New York, NY, June 2002.
- [2] Russ Cox, Frank Dabek, Frans Kaashoek, Jinyang Li, and Robert Morris. Practical, distributed network coordinates. In *Proceedings of the ACM Workshop on Hot Topics in Networks (HotNets)*. Cambridge, MA, November 2003.
- [3] Eng Keong Lua, Timothy Griffin, Marcelo Pias, Han Zheng, and Jon Crowcroft. On the accuracy of embeddings for Internet coordinate systems. In *Proceedings of the ACM SIGCOMM Internet Measurement Conference (IMC)*. Berkeley, CA, October 2005.
- [4] Paul Francis, Sugih Jamin, Cheng Jin, Yixin Jin, Danny Raz, Yuval Shavitt, and Lixia Zhang. IDMaps: A global Internet host distance estimation service. *IEEE/ACM Transactions on Networking*, 9(5):525–540, October 2001.
- [5] Ningning Hu and Peter Steenkiste. Exploiting Internet route sharing for large scale available bandwidth estimation. In *Proceedings of the ACM SIGCOMM Internet Measurement Conference (IMC)*. Berkeley, CA, October 2005. URL <http://gs274.sp.cs.cmu.edu/www/papers/broute-imc05.pdf>.

Unit 12: Making stuff up – Justin

Measurement only goes so far. We cannot measure the money that changes hands for a link or how much traffic goes from one end of the network to another. But with enough other information, and a good model of how the economics work, maybe we can guess really well.

Readings

- [1] Yin Zhang, Matthew Roughan, Nick Duffield, and Albert Greenberg. Fast accurate computation of large-scale IP traffic matrices from link loads. In *Proceedings of the ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Systems*, pages 206–217. San Diego, CA, June 2003.
- [2] Lun Li, David Alderson, Walter Willinger, and John Doyle. A first-principles approach to understanding the Internet’s router-level topology. In *Proceedings of the ACM SIGCOMM Conference on Applications, Technologies, Architectures, and Protocols for Computer Communication (SIGCOMM)*, pages 3–14. Portland, OR, August 2004. URL <http://netlab.caltech.edu/pub/papers/topology-sigcomm04.pdf>.
- [3] Hyunseok Chang, Sugih Jamin, and Z. Morley Mao and Walter Willinger. An empirical approach to modeling inter-as traffic matrices. In *Proceedings of the ACM SIGCOMM Internet Measurement Conference (IMC)*. Berkeley, CA, October 2005. URL <http://topology.eecs.umich.edu/archive/imc05.pdf>.

Unit 13: Platforms

Several execution platforms enable measurement collection. Mine is Scriptroute, though there are others.

Readings

- [1] Neil Spring, David Wetherall, and Thomas Anderson. Scriptroute: A public Internet measurement facility. In *Proceedings of the USENIX Symposium on Internet Technologies and Systems (USITS)*, pages 225–238. Seattle, WA, March 2003. URL <http://www.cs.umd.edu/~nspring/papers/scriptroute.pdf>.
- [2] Fotis Georgatos, Florian Gruber, Daniel Karrenberg, Mark Santcroos, Ana Susanj, Henk Uijterwaal, and René Wilhelm. Providing active measurements as a regular service for ISP's. In *Proceedings of Passive & Active Measurement (PAM)*. March 2001.
- [3] Sunil Kalidindi and Matthew J. Zekauskas. Surveyor: An infrastructure for Internet performance measurements. In *INET'99*. June 1999.
- [4] Vern Paxson, Andrew Adams, and Matt Mathis. Experiences with NIMI. In *Proceedings of Passive & Active Measurement (PAM)*. April 2000. URL <http://www.icir.org/vern/papers/nimi-pam-00.ps.gz>.
- [5] Vern Paxson, Jamshid Mahdavi, Andres Adams, and Matt Mathis. An architecture for large-scale internet measurement. *IEEE Communications*, 36(8):48–54, August 1998. URL <ftp://ftp.ee.lbl.gov/papers/nimi-ieee-comm98.ps.gz>.
- [6] Charles Robert Simpson, Jr. and George F. Riley. NETI@home: A distributed approach to collecting end-to-end network performance measurements. In *Proceedings of Passive & Active Measurement (PAM)*. Antibes Juan-les-Pins, France, April 2004. URL <http://www.pam2004.org/papers/127.pdf>.

Unit 14: PlanetLab – Galileo

PlanetLab is a very large network testbed where users, like yourselves, can run code both to measure and to do other things.

Readings

- [1] Suman Banerjee, Timothy G. Griffin, and Marcelo Pias. The interdomain connectivity of PlanetLab nodes. In *Proceedings of Passive & Active Measurement (PAM)*, pages 73–82. Antibes Juan-les-Pins, France, April 2004.
- [2] Ming Zhang, Chi Zhang, Vivek Pai, Larry Peterson, and Randy Wang. PlanetSeer: Internet path failure monitoring and characterization in wide-area services. In *Proceedings of the Symposium on Operating Systems Design and Implementation (OSDI)*. San Francisco, CA, December 2004.
- [3] Larry Peterson, Thomas Anderson, David Culler, and Timothy Roscoe. A blueprint for introducing disruptive technology into the Internet. In *Proceedings of the ACM Workshop on Hot Topics in Networks (HotNets)*, pages 59–64. Princeton, NJ, October 2002. URL http://www.cs.princeton.edu/nsg/papers/planetlab_hotnets_02/hotnets.ps.
- [4] Ruoming Pang, Vinod Yegneswaran, Paul Barford, Vern Paxson, and Larry Peterson. Characteristics of Internet background radiation. In *Proceedings of the ACM SIGCOMM Internet Measurement Conference (IMC)*. Taormina, Sicily, Italy, October 2004. URL http://www.cs.princeton.edu/nsg/papers/radiation_imc_04/radiation.pdf.
- [5] Sung-Ju Lee, Puneet Sharma, Sujata Banerjee, Sujoy Basu, and Rodrigo Fonseca. Measuring bandwidth between PlanetLab nodes. In *Proceedings of Passive & Active Measurement (PAM)*. Boston, MA, March 2005.

Unit 15: errors – Aaron

A few papers describe common mistakes of tools and of interpretation.

Readings

- [1] Ramesh Govindan and Vern Paxson. Estimating router ICMP generation delays. In *Proceedings of Passive & Active Measurement (PAM)*, pages 6–13. Fort Collins, CO, March 2002. URL <http://www.icir.org/vern/papers/fsd-pam-02.pdf>.
- [2] Ravi S. Prasad, Constantinos Dovrolis, and Bruce A. Mah. The effect of layer-2 switches on pathchar-like tools. In *Proceedings of the ACM SIGCOMM Internet Measurement Workshop (IMW)*, pages 321–322. Marseille, France, November 2002.

- [3] Sue B. Moon, Paul Skelly, and Don Towsley. Estimation and removal of clock skew from network delay measurements. In *Proceedings of the IEEE Joint Conference of the IEEE Computer and Communications Societies (INFOCOM)*. New York, NY, March 1999.
- [4] Z. Morley Mao, Jennifer Rexford, Jia Wang, and Randy Katz. Towards an accurate AS-level traceroute tool. In *Proceedings of the ACM SIGCOMM Conference on Applications, Technologies, Architectures, and Protocols for Computer Communication (SIGCOMM)*, pages 365–378. Karlsruhe, Germany, August 2003.
- [5] Renata Teixeira and Jennifer Rexford. A measurement framework for pin-pointing routing changes. In *ACM SIGCOMM Network Troubleshooting Workshop*, pages 313–318. Portland, OR, August 2004. URL <http://www.cs.princeton.edu/~jrex/papers/omni.pdf>.

Unit 16: Vern – Rob

Vern Paxson is a pioneer of network measurement; this week will focus on his early (dissertation) work and on some of his recent measurement-related work.

Anything in <http://www.icir.org/vern/papers.html> is fair game, what follows are some of the more heavily cited ones.

Readings

- [1] Vern Paxson. Strategies for sound Internet measurement. In *Proceedings of the ACM SIGCOMM Internet Measurement Conference (IMC)*, pages 263–271. Taormina, Sicily, Italy, October 2004. URL <http://www.icir.org/vern/papers/meas-strategies-imc04.pdf>.
- [2] Vern Paxson. End-to-end routing behavior in the Internet. In *Proceedings of the ACM SIGCOMM Conference on Applications, Technologies, Architectures, and Protocols for Computer Communication (SIGCOMM)*, pages 25–38. Palo Alto, CA, August 1996.
- [3] Vern Paxson. On calibrating measurements of packet transit times. In *Proceedings of the ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Systems*. 1998.

Unit 17: Workload – Sandro

Passive measurements often determine the characteristics of the workload different applications see, and how those characteristics relate to protocols.

Readings

- [1] Stefan Saroiu, Steven D. Gribble, and Henry M. Levy. Measurement and analysis of spyware in a university environment. In *Proceedings of the ACM/USENIX Symposium on Networked Systems Design and Implementation (NSDI)*. San Francisco, CA, March 2004. URL <http://www.cs.washington.edu/homes/tzoozmpy/publications/nsdi/2004/>.
- [2] Alec Wolman, Geoffrey 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 ACM Symposium on Operating Systems Principles (SOSP)*, pages 16–31. Kiawah Island Resort, SC, December 1999. URL <http://www.cs.washington.edu/research/networking/websys/pubs/sosp99/sosp99.pdf>.
- [3] Hari Balakrishnan, Venkata N. Padmanabhan, Srinivasan Seshan, Mark Stemm, and Randy H. Katz. TCP behavior of a busy Internet server: Analysis and improvements. In *Proceedings of the IEEE Joint Conference of the IEEE Computer and Communications Societies (INFOCOM)*. San Francisco, CA, March 1998.

Unit 18: Big Data – Vijay

Few networking researchers have investigated how best to store network measurement data while making it available for further study using something more sophisticated than FTP. We could use some help.

I also really like the OMNI paper because it clears up misconceptions about Internet behavior.

I wouldn't mind reading papers outside networking to understand better how this sort of (real-time and historical) data is made available.

Readings

- [1] Mark Allman, Ethan Blanton, and Wesley M. Eddy. A scalable system for sharing Internet measurement. In *Proceedings of Passive & Active Measurement (PAM)*, pages 189–191. Fort Collins, CO, March 2002.
- [2] Sridhar Srinivasan and Ellen Zegura. Network measurement as a cooperative enterprise. In *International Workshop on Peer-to-Peer Systems (IPTPS)*, pages 166–177. Cambridge, MA, March 2002.
- [3] Renata Teixeira and Jennifer Rexford. A measurement framework for pin-pointing routing changes. In *ACM SIGCOMM Network Troubleshooting Workshop*, pages 313–318. Portland, OR, August 2004. URL <http://www.cs.princeton.edu/~jrex/papers/omni.pdf>.

Unit 19: Measurement Protocols – Amy

Although clever hacks to determine properties are nice, perhaps there are alternatives, using explicit protocols that would report performance measures of interest. Are proposed approaches narrow or elegant? Nice ways of doing the same old stuff, or motivated by an ability to do the otherwise impossible? If explicit, would routers lie?

Readings

- [1] Mathew J. Luckie, Anthony J. McGregor, and Hans-Werner Braun. Towards improving packet probing techniques. In *Proceedings of the ACM SIGCOMM Internet Measurement Workshop (IMW)*, pages 145–150. San Francisco, CA, November 2001.
- [2] Ratul Mahajan, Neil Spring, David Wetherall, and Thomas Anderson. User-level Internet path diagnosis. In *Proceedings of the ACM Symposium on Operating Systems Principles (SOSP)*, pages 106–119. Bolton Landing, NY, October 2003. URL <http://www.cs.washington.edu/research/networking/tulip/bits/sosp2003-tulip.pdf>.
- [3] Matthew Luckie and Tony McGregor. Path diagnosis with IPMP. In *ACM SIGCOMM Workshop on Network Troubleshooting: Research, Theory, and Operations Practice Meet Malfunctioning Reality (NetTs'04)*. Portland, OR, September 2004. URL <http://wand.cs.waikato.ac.nz/pubs/208/pdf/nts35-luckie1.pdf>.
- [4] Stanislav Shalunov, Benjamin Teitelbaum, Anatoly Karp, Jeff W. Boote, and Matthew J. Zekauskas. A one-way active measurement protocol (OWAMP). Internet Draft: draft-ietf-ippm-owdp-14.txt, November 2004. Work in progress.

Unit 20: Constantinos

Constantinos Dovrolis, and his student Manish, have worked extensively on available bandwidth estimation techniques. These are interesting tools that are hard to evaluate because they measure dynamically varying properties.

Readings

- [1] Manish Jain and Constantinos Dovrolis. End-to-end estimation of the available bandwidth variation range. In *Proceedings of the ACM SIGMETRICS International Conference on Measurement and Modeling of Computer Systems*. Banff, Canada, June 2005. URL <http://www.cc.gatech.edu/fac/Constantinos.Dovrolis/Papers/F105-jain.pdf>.
- [2] Manish Jain and Constantinos Dovrolis. End-to-end available bandwidth: measurement methodology, dynamics, and relation with TCP throughput. In *Proceedings of the ACM SIGCOMM Conference on Applications, Technologies, Architectures, and Protocols for Computer Communication (SIGCOMM)*, pages 295–308. Pittsburgh, PA, August 2002. URL <http://www.cc.gatech.edu/~jain/publications/sigcomm02.pdf>.
- [3] Ravi S. Prasad, Margaret Murray, Constantinos Dovrolis, and kc claffy. Bandwidth estimation: metrics, measurement techniques, and tools. *IEEE Network*, 17(6):27–35, November 2004.
- [4] Constantinos Dovrolis, Parameswaran Ramanathan, and David Moore. What do packet dispersion techniques measure? In *Proceedings of the IEEE Joint Conference of the IEEE Computer and Communications Societies (INFOCOM)*, pages 905–914. Anchorage, AK, April 2001.
- [5] Ravi Prasad, Manish Jain, and Constantinos Dovrolis. Effects of interrupt coalescence on network measurements. In *Proceedings of Passive & Active Measurement (PAM)*. Antibes Juan-les-Pins, France, April 2004. URL http://www.cc.gatech.edu/fac/Constantinos.Dovrolis/Papers/intcoal_pam04.pdf.

Unit 21: Tools

This “week” is a bag of tools; some of them are more interesting than others, I expect we would cover only a handful or run quickly through the insight and evaluation of some.

Readings

- [1] Dina Katabi, Issam Bazzi, and Xiaowei Yang. A passive approach for detecting shared bottlenecks. In *IEEE International Conference on Computer Communications and Networks (ICCCN)*. 2001.
- [2] Krishna P. Gummadi, Stefan Saroiu, and Steven D. Gribble. King: Estimating latency between arbitrary Internet end hosts. In *Proceedings of the ACM SIGCOMM Internet Measurement Workshop (IMW)*, pages 5–18. Marseille, France, November 2002. URL www.cs.washington.edu/homes/gummadi/papers/king.pdf.
- [3] Allen Downey. Evidence for long-tailed distributions in the Internet. In *Proceedings of the ACM SIGCOMM Internet Measurement Workshop (IMW)*, pages 229–241. San Francisco, CA, November 2001. URL <http://allendowney.com/research/longtail/downey01evidence.pdf>.
- [4] Vern Paxson and Sally Floyd. Wide-area traffic: The failure of poisson modeling. *IEEE/ACM Transactions on Networking*, 3(3):226–244, June 1995.
- [5] Attila Pásztor and Darryl Veitch. Active probing using packet quartets. In *Proceedings of the ACM SIGCOMM Internet Measurement Workshop (IMW)*, pages 293–305. Marseille, France, November 2002.
- [6] Stefan Saroiu, P. Krishna Gummadi, and Steven D. Gribble. Sprobe: A fast technique for measuring bottleneck bandwidth in uncooperative environments. In *In submission*. 2002. <http://sprobe.cs.washington.edu/sprobe.ps>.
- [7] Jitendra Padhye and Sally Floyd. Identifying the TCP behavior of Web servers. In *Proceedings of the ACM SIGCOMM Conference on Applications, Technologies, Architectures, and Protocols for Computer Communication (SIGCOMM)*, pages 287–298. San Diego, CA, August 2001.
- [8] Khaled Harfoush, Azer Bestavros, and John Byers. Measuring bottleneck bandwidth of targeted path segments. In *Proceedings of the IEEE Joint Conference of the IEEE Computer and Communications Societies (INFOCOM)*, pages 2079–2089. San Francisco, CA, April 2003.
- [9] Kostas G. Anagnostakis, Michael B. Greenwald, and Raphael S. Ryger. cing: Measuring network-internal delays using only existing infrastructure. In *Proceedings of the IEEE Joint Conference of the IEEE Computer and Communications Societies (INFOCOM)*, pages 2112–2121. San Francisco, CA, April 2003.
- [10] Jonathan Stone and Craig Partridge. When the CRC and TCP checksum disagree. In *Proceedings of the ACM SIGCOMM Conference on Applications, Technologies, Architectures, and Protocols for Computer Communication (SIGCOMM)*, pages 309–319. Stockholm, Sweden, August 2000.
- [11] Vinay Ribeiro, Rudolf Riedi, Richard Baraniuk, Jiri Navratil, and Les Cottrell. pathChirp: Efficient available bandwidth estimation for network paths. In *Proceedings of Passive & Active Measurement (PAM)*. La Jolla, CA, April 2003.
- [12] Steven M. Bellovin. A technique for counting NATted hosts. In *Proceedings of the ACM SIGCOMM Internet Measurement Workshop (IMW)*, pages 267–272. Marseille, France, November 2002.
- [13] Jon C. R. Bennett, Craig Partridge, and Nicholas Schectman. Packet reordering is not pathological network behavior. *IEEE/ACM Transactions on Networking*, 7(6):789–798, December 1999.