Delivering Services on Today's Internet—Fall 2013, 694a

Ethan Katz-Bassett

November 17, 2013

Today, large web providers such as Google and Netflix account for significant portions of the Internet's traffic. In this course, we will investigate how these providers use (and abuse) the Internet's protocols to deliver their services.[1] Our goals include:

- Extend the simplified view of the Internet covered in standard breadth classes to a richer view of how it is used today.
- Understand how the protocols and networks fit together to deliver services. Much of the complexity and many of the problems on the Internet arise in these interactions.
- Understand how the protocols are–and are not–good fits for the big players on today's Internet. The protocols were designed in a very different setting, and providers today go to great lengths to achieve their goals in the face of limited protocols.
- Identify open problems and investigate potential approaches, both from the literature and in our own research for class.

1 Basic Information

- Place and time: KAP145, Wed 3:30pm-6:20pm
- Instructor: Ethan Katz-Bassett
- email: ethan.kb/usc.edu
- phone: –
- office: SAL 236
- office hours: By appointment.
- Course homepage: `http://www-bcf.usc.edu/~katzbass/teaching/csci694/2013FA/index.html`
- Prerequisites: CS551 of permission of the instructor. This class is appropriate for graduate students or advanced undergraduates with previous classwork in networking. Students from non-systems/networking areas are welcome.
- Textbooks and course materials: No required textbook. We will read research papers, whitepapers from industry, and other similar documents. We will generally read the equivalent of 3-4 research papers a week.

2 Grading and Coursework

There are no exams in this class. The course grade will be determined based on:

Written paper responses and class presentations/discussion (50%): Students are expected to write responses to 2-4 papers a week. Each week, one or two students will give a presentation on the papers and lead the discussions in each class. Other students are expected to participate in the discussion.

[1] In focusing on the Internet's role in these services, we will, for the most part, not cover issues that predominantly pertain to a service's clients or servers but not the network between them.

For two weeks, instead of the standard course structure, we will instead conduct a Shadow Program Committee for NSDI 2014. We will read submitted papers and go through the reviewing process, ultimately arriving at a shadow conference program. This is an opportunity to learn about the peer-review process and gain experience as a reviewer. It will expose us to cutting-edge papers related to the course. Hopefully, the papers will include some from industry, as NSDI is having an Operational Systems track this year.

A research project, including 6 pg writeup and 20 minute presentation (50%): The semester-long project is an open-ended research project on one of the covered topics. The instructor will provide some possible project topics, or (with instructor approval) you can work on a project of your own devising. Projects should be done in groups of two (or get instructor approval for a different size).

3 Reading List and Class Schedule

Week 1 8/28: Course Overview and Introduction

Week 2 9/4: Content Delivery Infrastructures (Presented by Brandon and Zahaib)

P1 Erik Nygren, Ramesh K. Sitaraman, and Jennifer Sun. The Akamai network: A platform for high-performance Internet applications. *ACM SIGOPS Operating Systems Review*, 2010.

P2 Matt Calder, Xu Fan, Zi Hu, Ramesh Govindan, John Heidemann, and Ethan Katz-Bassett. Mapping the expansion of Google's serving infrastructure. In *IMC*, 2013.

P3 Mukarram Bin Tariq, Kaushik Bhandankar, Vytautas Valancius, Amgad Zeitoun, Nick Feamster, and Mostafa Ammar. Answering "what-if" deployment and configuration questions with WISE: Techniques and deployment experience. *IEEE/ACM TON*, 2013.

Supplemental:

- Bernhard Ager, Wolfgang Mühlbauer, Georgios Smaragdakis, and Steve Uhlig. Web content cartography. In *IMC*, 2011.

Week 3 9/11: DNS and Content Routing (Presented by Himanshu and Pramit)

P4 Zhuoqing Morley Mao, Charles D. Cranor, Fred Douglis, Michael Rabinovich, Oliver Spatscheck, and Jia Wang. A precise and efficient evaluation of the proximity between web clients and their local DNS servers. In *USENIX Annual Technical Conference*, 2002.

P5 John S. Otto, Mario A. Sánchez, John P. Rula, and Fabián E. Bustamante. Content delivery and the natural evolution of DNS. In *IMC*, 2012.

P6 Rupa Krishnan, Harsha V. Madhyastha, Sridhar Srinivasan, Sushant Jain, Arvind Krishnamurthy, Thomas Anderson, and Jie Gao. Moving beyond end-to-end path information to optimize CDN performance. In *IMC*, 2009.

Week 4 9/18: Interdomain Routing: Peering and Internet Exchange Points (IXPs) (Presented by Himanshu and Pramit)

Background (you are expected to read and understand these, but you do not need to write responses):

- (short) Matthew Caesar and Jennifer Rexford. BGP routing policies in ISP networks. *IEEE Network Magazine, special issue on Interdomain Routing*, 2005.
- Peyman Faratin, David Clark, Steven Bauer, William Lehr, Patrick Gilmore, and Arthur Berger. The growing complexity of Internet interconnection. *COMMUNICATIONS & STRATEGIES*, 2008.
- (short) William B. Norton. The art of peering: The peering playbook. Technical report, DrPeering International, 2010.

Papers to write responses to:

P7 Wolfgang Mühlbauer, Steve Uhlig, Bingjie Fu, Mickael Meulle, and Olaf Maennel. In search for an appropriate granularity to model routing policies. In *SIGCOMM*, 2007.

P8 Yaping Zhu, Benjamin Helsley, Jennifer Rexford, Aspi Siganporia, and Sridhar Srinivasan. LatLong: Diagnosing wide-area latency changes for CDNs. *IEEE Transactions on Network and Service Management*, 2012.

P9 (short, *NOT YET PUBLISHED*) Nikolaos Chatzis, Georgios Smaragdakis, Anja Feldmann, and Walter Willinger. There is more to IXPs than meets the eye. *SIGCOMM CCR*, 2013.

P10 (short, *UNDER SUBMISSION*) Arpit Gupta, Muhammad Shahbaz, Laurent Vanbever, Hyojoon Kim, Russ Clark, Nick Feamster, Jennifer Rexford, and Scott Shenker. SDX: A software defined Internet exchange. Under submission, 2013.

Supplemental:

- (short, *UNDER SUBMISSION*) Phillipa Gill, Michael Schapira, and Sharon Goldberg. A survey of interdomain routing policies. Under submission, 2013.
- Neil Spring, Ratul Mahajan, and Thomas Anderson. Quantifying the causes of path inflation. In *SIGCOMM*, 2003.
- (short) William B. Norton. The art of peering: The IX playbook. Technical report, DrPeering International, 2010.
- (very short) William B. Norton. Why not peer? the top 10 reasons not to peer. Technical report, DrPeering International.

Week 5 9/25: Route Redistribution and Connecting Multiple Routing Instances (Presented by Nitish and Matt)

P11 Nikola Gvozdiev, Brad Karp, and Mark Handley. LOUP: The principles and practice of intra-domain route dissemination. In *NSDI*, 2013.

P12 Franck Le, Geoffrey G. Xie, Dan Pei, Jia Wang, and Hui Zhang. Shedding light on the glue logic of the Internet routing architecture. In *SIGCOMM*, 2008.

P13 (short) Nick Feamster, Hari Balakrishnan, Jennifer Rexford, Aman Shaikh, and Jacobus van der Merwe. The case for separating routing from routers. In *FDNA Workshop*, 2004.

Supplemental:

- Franck Le, Geoffrey G. Xie, and Hui Zhang. Understanding route redistribution. In *ICNP*, 2007.
- Franck Le, Geoffrey G. Xie, and Hui Zhang. Instability free routing: Beyond one protocol instance. In *CoNEXT*, 2008.
- Franck Le, Geoffrey G. Xie, and Hui Zhang. Theory and new primitives for safely connecting routing protocol instances. In *SIGCOMM*, 2010.
- Matthew Caesar, Donald Caldwell, Nick Feamster, Jennifer Rexford, Aman Shaikh, and Jacobus van der Merwe. Design and implementation of a routing control platform. In *NSDI*, 2005.

Week 6 10/2: Transport I: Split and Multiplexed TCP (Presented by Zahaib and Brandon)

P14 (short) Abhinav Pathak, Y. Angela Wang, Cheng Huang, Albert Greenberg, Y. Charlie Hu, Randy Kern, Jin Li, and Keith W. Ross. Measuring and evaluating TCP splitting for cloud services. In *PAM*, 2010.

P15 (short) Yingying Chen, Sourabh Jain, Vijay Kumar Adhikari, and Zhi-Li Zhang. Characterizing roles of front-end servers in end-to-end performance of dynamic content distribution. In *IMC*, 2011.

P16 (short) SPDY: An experimental protocol for a faster web. Technical report, Google Inc.

P17 Xiao Sophia Wang, Aruna Balasubramanian, Arvind Krishnamurthy, and David Wetherall. How speedy is SPDY? Under submission.

Supplemental

- Tobias Flach, Nandita Dukkipati, Andreas Terzis, Barath Raghavan, Neal Cardwell, Yuchung Cheng, Ankur Jain, Shuai Hao, Ethan Katz-Bassett, and Ramesh Govindan. Reducing web latency: the virtue of gentle aggression. In *SIGCOMM*, 2013.

Week 7 10/9: Transport II: Beyond TCP (Presented by Nitish and Matt)

P18 Jim Roskind. QUIC: Design document and specification rational. Technical report, Google Inc., 2013.

P19 Bryan Ford. Structured streams: a new transport abstraction. In *SIGCOMM*, 2007.

P20 Michael F. Nowlany Nabin Tiwari Michael Nowlan, Nabin Tiwari, Janardhan Iyengar, Syed Obaid Amin, and Bryan Ford. Fitting square pegs through round pipes: Unordered delivery wire-compatible with TCP and TLS. In *NSDI*, 2012.

Week 8 10/16: NSDI Shadow PC meeting 1

Week 9 10/23: **No class: Internet Measurement Conference**

Week 10 10/30: NSDI Shadow PC meeting 2

Week 11 11/6: Traffic Engineering (Presented by Haonan and Xiyue)

P21 *Bernard Fortz, Jennifer Rexford, and Mikkel Thorup. Traffic engineering with traditional IP routing protocols.* IEEE Communication Magazine*, 2002.*

P22 Nick Feamster, Jay Borkenhagen, and Jennifer Rexford. Guidelines for interdomain traffic engineering. *SIGCOMM CCR*, 2003.

P23 Anwar Elwalid, Cheng Jin, Steven Low, and Indra Widjaja. MATE: MPLS adaptive traffic engineering. In *INFOCOM*, 2001.

P24 Sushant Jain, Alok Kumar, Subhasree Mandal, Joon Ong, Leon Poutievski, Arjun Singh, Subbaiah Venkata, Jim Wanderer, Junlan Zhou, Min Zhu, Jonathan Zolla, Urs Hölzle, Stephen Stuart, and Amin Vahdat. B4: Experience with a globally deployed software defined WAN. In *SIGCOMM*, 2013.

P25 (very short) Martín Casado, Teemu Koponen, Scott Shenker, and Amin Tootoonchian. Fabric: A retrospective on evolving SDN. In *HotSDN*, 2012.

Supplemental

- Murali Kodialam and T. V. Lakshman. Minimum interference routing with applications to MPLS traffic engineering. In *INFOCOM*, 2000.
- B. Quoitin, Steve Uhlig, C. Pelsser, L. Swinnen, and Olivier Bonaventure. Interdomain traffic engineering with BGP. *IEEE Communications Magazine*, 2003.

Week 12 11/13: Video Delivery (Presented by Haonan)

P26 (tech report version, NOT the shorter INFOCOM version) Vijay Kumar Adhikari, Sourabh Jain, Yingying Chen, and Zhi-Li Zhang. Vivisecting youtube: An active measurement study. Technical report, University of Minnesota, 2011.

P27 (short) Monia Ghobadi, Yuchung Cheng, Ankur Jain, and Matt Mathis. Trickle: Rate limiting YouTube video streaming. In *USENIX ATC*, 2012.

P28 Xi Liu, Florin Dobrian, Henry Milner, Junchen Jiang, Vyas Sekar, Ion Stoica, and Hui Zhang. A case for a coordinated internet video control plane. In *SIGCOMM*, 2012.

Supplemental

- (Minlan will present this paper) Minlan Yu, Wenjie Jiang, Haoyuan Li, and Ion Stoica. Tradeoffs in CDN designs for throughput oriented traffic. In *CoNEXT*, 2012.
- Vijay Kumar Adhikari, Sourabh Jain, and Zhi-Li Zhang. YouTube traffic dynamics and its interplay with a tier-1 ISP: An ISP perspective. In *IMC*, 2010.
- (short) Vijay Kumar Adhikari, Yang Guo, Fang Hao, Volker Hilt, , and Zhi-Li Zhang. A tale of three CDNs: An active measurement study of Hulu and its CDNs. In *IEEE Global Internet Symposium*, 2012.
- (short) Vijay Kumar Adhikari, Yang Guo, Fang Hao, Matteo Varvello, Volker Hilt, Moritz Steiner, and Zhi-Li Zhang. Unreeling Netflix: Understanding and improving multi-CDN movie delivery. In *INFOCOM*, 2012.

Week 13 11/20: Looking Across Boundaries for New Solutions (Presented by Xiyue and Matt)

P29 Vytautas Valancius, Bharath Ravi, Nick Feamster, and Alex C. Snoeren. Quantifying the benefits of joint content and network routing. In *SIGMETRICS*, 2013.

P30 Ingmar Poese, Benjamin Frank, Georgios Smaragdakis, Steve Uhlig, Anja Feldmann, and Bruce Maggs. Enabling content-aware traffic engineering. *SIGCOMM CCR*, 2012.

P31 Ratul Mahajan, David Wetherall, and Thomas Anderson. Mutually controlled routing with independent ISPs. In *NSDI*, 2007.

Supplemental

- Ingmar Poese, Benjamin Frank, Bernhard Ager, Georgios Smaragdakis, and Anja Feldmann. Improving content delivery using provider-aided distance information. In *IMC*, 2010.
- Benjamin Frank, Ingmar Poese, Yin Lin, Georgios Smaragdakis, Anja Feldmann, Bruce Maggs, Jannis Rake, Steve Uhlig, and Rick Weber. Pushing CDN-ISP collaboration to the limit. *ACM CCR*, 2013.
- Ratul Mahajan, David Wetherall, and Tom Anderson. Negotiation-based routing between neighboring ISPs. In *NSDI*, 2005.

Week 14 11/27: **No class: Thanksgiving**

Week 15 12/4: Class presentations

Week 16 12/14: **Project report due**

Week 17 12/16: **Review another group's project report**

Statement for Students with Disabilities

Any student requesting academic accommodations based on a disability is required to register with Disability Services and Programs (DSP) each semester. A letter of verification for approved accommodations can be obtained from DSP. Please be sure the letter is delivered to me as early in the semester as possible. DSP is located in STU 301 and is open 8:30 a.m.5:00 p.m., Monday through Friday. The phone number for DSP is (213) 740-0776.

Statement on Academic Integrity

USC seeks to maintain an optimal learning environment. General principles of academic honesty include the concept of respect for the intellectual property of others, the expectation that individual work will be submitted unless otherwise allowed by an instructor, and the obligations both to protect one's own academic work from misuse by others as well as to avoid using another's work as one's own. All students are expected to understand and abide by these principles. Scampus, the Student Guidebook, contains the Student Conduct Code in Section 11.00, while the recommended sanctions are located in Appendix A: `http://www.usc.edu/dept/publications/SCAMPUS/gov/` Students will be referred to the Office of Student Judicial Affairs and Community Standards for further review, should there be any suspicion of academic dishonesty. The Review process can be found at: `http://www.usc.edu/student-affairs/SJACS/`

Emergency Preparedness/Course Continuity in a Crisis

In case of a declared emergency if travel to campus is not feasible, USC executive leadership will announce an electronic way for instructors to teach students in their residence halls or homes using a combination of Blackboard, teleconferencing, and other technologies.