Fall 2020 EECS E6890

Delivering Modern Services on the Internet

Services such as Facebook and Netflix occupy much of our time and account for most Internet traffic. This course investigates how these providers use the Internet to deliver their services, focusing on hurdles they face because the Internet was designed in a very different setting than today's.

Course goals include answering:

- How are large Internet services architected and operated?
- How (and how well) do Internet protocols and networks fit together?
- What led to the Internet's design? Is it a good fit for today's services?
- What problems do services still face? What are potential solutions?

Time and place:	Wednesday 4:10pm-6:40pm on Zoom Live attendance required, but we will discuss timing to find a good option for as many people as possible.
Instructor:	Ethan Katz-Bassett
Prerequisites:	CSEE 4119 or equivalent, or permission of the instructor. Appropriate for grad students or advanced undergrads with previous classwork in networking. Students from non-systems/networking areas are welcome.
Format:	Reading, writing about, and discussing research papers. Short presentations. Projects in small groups.

For course page and draft syllabus, see:
http://www.columbia.edu/~ebk2141/DeliveringServices.html

Schedule

Course introduction and overview

Sep 9: Course introduction

1. [type:background] A Brief History of the Internet. Barry M. Leiner, Vinton G. Cerf, David D. Clark, Robert E. Kahn, Leonard Kleinrock, Daniel C. Lynch, Jon Postel, Larry G. Roberts, Stephen Wolff. Internet Society, 1997.
2. [type:background] (very short paper) How to Read a Paper. S. Keshav.

   Suggested supplemental paper:
   An evaluation of the ninth SOSP submissions -or- How (and how not) to write a good systems paper. Roy Levin and David D. Redell. ACM SIGOPS Operating Systems Review 17(3):35-40 (July, 1983).

Design of the Internet

Sep 16: Philosophy of the Internet's design

3. [type:design/type:design-philosophy] The design philosophy of the DARPA internet protocols. David D. Clark. In Proceedings of the ACM SIGCOMM Conference, pages 106–114. ACM, August 1988.
4. [type:design-philosophy] (short paper) End-to-end arguments in system design. J. H. Saltzer, D. P. Reed, and D. D. Clark. Proceedings of the 2nd International Conference on Distributed Computing Systems, pages 509–512, April 1981.

Sep 23: Evolution of the Internet's protocols and traffic

5. [type:design-philosophy] (short paper) Why The Internet Only Just Works. Mark Handley. BT Technology Journal, Vol 24, No 3, July 2006.
6. [type:measurement-internal] Internet inter-domain traffic. Craig Labovitz, Scott Iekel-Johnson, Danny McPherson, Jon Oberheide, and Farnam Jahanian. In Proceedings of the ACM SIGCOMM Conference, pages 75–86, New Delhi, India, August 2010. ACM.
7. [type:measurement-external] (short paper) Are We One Hop Away from a Better Internet? Y. Chiu, B. Schlinker, A. B. Radhakrishnan, E. Katz-Bassett, R. Govindan. ACM Internet Measurement Conference (IMC), 2015.

   Suggested supplemental paper:
   Tussle in cyberspace: Defining tomorrow's Internet. David D. Clark, John Wroclawski, Karen Sollins, and Robert Braden. In Proceedings of the ACM

SIGCOMM Conference, pages 347–356, Pittsburgh, PA, USA, August 2002. ACM.

Content delivery infrastructures: Bringing content close to users

Sep 30: The rise of content delivery networks

8. [type:background] (short paper) Globally distributed content delivery. J. Dilley, B. Maggs, J. Parikh, H. Prokop, R. Sitaraman, and B. Weihl, IEEE Internet Computing, September/October 2002, pp. 50-58.
9. [type:measurement-external] Mapping the expansion of Google's serving infrastructure. Matt Calder, Xu Fan, Zi Hu, Ramesh Govindan, John Heidemann, and Ethan Katz-Bassett. In IMC, 2013.
10. [type:measurement-internal] Leveraging Interconnections for Performance: The Serving Infrastructure of a Large CDN. Florian Wohlfart, Nikolaos Chatzis, Caglar Dabanoglu, Georg Carle, Walter Willinger. In SIGCOMM, 2018.

Routing traffic to users

Oct 7: Flexible routing at the BGP edge

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

11. [type:background] (short paper) Matthew Caesar and Jennifer Rexford. BGP routing policies in ISP networks. IEEE Network Magazine, special issue on Interdomain Routing, 2005.

Papers to write responses to:

12. [type:design] Engineering egress with Edge Fabric: Steering oceans of content to the world. Brandon Schlinker, Hyojeong Kim, Timothy Cui, Ethan Katz-Bassett, Harsha V. Madhyastha, Italo Cunha, James Quinn, Saif Hasan, Petr Lapukhov, and Hongyi Zeng. In SIGCOMM 2017.
13. [type:design] Taking the Edge off with Espresso: Scale, Reliability and Programmability for Global Internet Peering. KK Yap, Murtaza Motiwala, Jeremy Rahe, Steve Padgett, Matthew Holliman, Gary Baldus, Marcus Hines, TaeEun Kim, Ashok Narayanan, Ankur Jain, Victor Lin, Colin Rice, Brian Rogan, Arjun Singh, Bert Tanaka, Manish Verma, Puneet Sood, Mukarram Tariq, Matt Tierney, Dzevad Trumic, Vytautas Valancius, Calvin Ying, Mahesh Kallahalla, Bikash Koley, Amin Vahdat. In SIGCOMM 2017.

Suggested supplemental paper:

[type:background] Peyman Faratin, David Clark, Steven Bauer, William Lehr, Patrick Gilmore, and Arthur Berger. The growing complexity of Internet interconnection. COMMUNICATIONS & STRATEGIES, 2008.

Oct 14: Interconnection

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

14. [type:background] (short paper) William B. Norton. The art of peering: The peering playbook. Technical report and slides. DrPeering International, 2010.

Papers to write responses to:

15. [type:measurement-external] Inferring Persistent Interdomain Congestion. A. Dhamdhere, D. Clark, A. Gamero-Garrido, M. Luckie, R. Mok, G. Akiwate, K. Gogia, V. Bajpai, A. Snoeren, and k. claffy. In ACM SIGCOMM, Aug 2018.
16. [type:design-research] SDX: A Software Defined Internet Exchange. A. Gupta, L. Vanbever, M. Shahbaz, S. P. Donovan, B. Schlinker, N. Feamster, J. Rexford, S. Shenker, R. Clark, E. Katz-Bassett. ACM SIGCOMM, 2014.

Directing users to content

Oct 21: Anycast

17. [type:design] FastRoute: A Scalable Load-Aware Anycast Routing Architecture for Modern CDNs. Ashley Flavel, Pradeepkumar Mani, David A. Maltz, and Nick Holt, Jie Liu, Yingying Chen, and Oleg Surmachev. In NSDI, 2015.
18. [type:measurement-internal] (short paper) Analyzing the performance of an anycast CDN. Matt Calder, Ashley Flavel, Ethan Katz-Bassett, Ratul Mahajan, and Jitendra Padhye. In Proceedings of the ACM Internet Measurement Conference, Tokyo, Japan, October 2015.

Oct 28: Watch IMC talks and report out

Nov 4: DNS redirection

19. [type:design/type:measurement-internal] End-User Mapping: Next Generation Request Routing for Content Delivery. Fangfei Chen, Ramesh K. Sitaraman, Marcelo Torres. In SIGCOMM 2015.
20. [type:measurement-internal] Moving beyond end-to-end path information to optimize CDN performance. Rupa Krishnan, Harsha V. Madhyastha, Sridhar Srinivasan, Sushant Jain, Arvind Krishnamurthy, Thomas Anderson, and Jie Gao. In IMC, 2009.

Transporting content to users

Nov 11: Split TCP

21. [type:design/type:measurement-internal] (short paper) Measuring and evaluating TCP splitting for cloud services. Abhinav Pathak, Y. Angela Wang, Cheng

Huang, Albert Greenberg, Y. Charlie Hu, Randy Kern, Jin Li, and Keith W. Ross. In PAM, 2010.

22. [type:measurement-external] (short paper) Characterizing roles of front-end servers in end-to-end performance of dynamic content distribution. Yingying Chen, Sourabh Jain, Vijay Kumar Adhikari, and Zhi-Li Zhang. In IMC, 2011.
23. [type:design] Reducing Web Latency: the Virtue of Gentle Aggression. T. Flach, N. Dukkipati, A. Terzis, B. Raghavan, N. Cardwell, Y. Cheng, A. Jain, S. Hao, E. Katz-Bassett, R. Govindan. ACM SIGCOMM, 2013.

Nov 18: Evolving transport

24. [type:design] BBR: Congestion-Based Congestion Control. Neal Cardwell, Yuchung Cheng, C. Stephen Gunn, Soheil Hassas Yeganeh, Van Jacobson Communications of the ACM, vol. 60 (2017), pp. 58-66
25. [type:design] The QUIC Transport Protocol: Design and Internet-Scale Deployment. Adam Langley, Al Riddoch, Alyssa Wilk, Antonio Vicente, Charles 'Buck' Krasic, Cherie Shi, Dan Zhang, Fan Yang, Feodor Kouranov, Ian Swett, Janardhan Iyengar, Jeff Bailey, Jeremy Christopher Dorfman, Jim Roskind, Joanna Kulik, Patrik Göran Westin, Raman Tenneti, Robbie Shade, Ryan Hamilton, Victor Vasiliev, Wan-Teh Chang. SIGCOMM 2017.

Nov 25: Thanksgiving Recess

Stepping out of the box

Dec 2: Blog posts on Dropbox, Riot, and BBR (by Geoff Huston)

Dec 9: Crossing boundaries

26. [type:measurement-internal/type:design-research] Efficiently Delivering Online Services over Integrated Infrastructure. Hongqiang Harry Liu, Raajay Viswanathan, Matt Calder, Aditya Akella, Ratul Mahajan, Jitendra Padhye, Ming Zhang. NSDI 2016.
27. [type:design] Steering Hyper-Giants' Traffic at Scale. Enric Pujol, Ingmar Poese, Johannes Zerwas, Georgios Smaragdakis, and Anja Feldmann. ACM CoNEXT 2019.