

The University of Chicago Computer Science Department  
**PRESENTS:**

**More Users, More Data... More Problems:  
Three Grand Challenges for the Modern Internet**



**Nick Feamster**  
Princeton University

**Abstract:**

The Internet was designed to make it easy for users and devices to connect and communicate. This ease of connectivity has resulted in more users, more devices, more traffic, more applications, and more user data. In this talk, I will describe three new challenges that the Internet faces as a result of these developments, and how we have begun to tackle them with new approaches and systems for network data analytics, ranging from lower-level measurement to higher-level inference and prediction.

First, Internet applications should be **usable, for everyone**. Internet access speeds continue to increase, with gigabit broadband access now a reality in many regions. Yet, faster access speeds does not imply that everyone is having a good experience. The digital divide has widened, performance bottlenecks have moved from the access network to other parts of the network, and the reliability and performance of the applications we use consistently fail to meet our expectations. Better measurement and inference capabilities can help, from improving diagnostics to informing public policy. I will describe data analytics systems that we have built that shed light on everything from the speed of a user's Internet service provider to the causes of poor streaming video quality.

Second, the Internet should be **secure and trustworthy**. The Internet's design has enabled malicious communications from denial-of-service to disinformation. Many citizens now rely on the Internet as their primary source of information and means to communicate. Yet, the protocols and platforms that the Internet uses to deliver this information to us remain extremely insecure and vulnerable to manipulation. I will describe data-driven systems that we have built to detect—and predict—Internet attacks, from Internet scams to disinformation.

Finally, the Internet should be **private**. Internet-connected “smart” infrastructure, from homes to cities, is poised to revolutionize personal health to urban planning. Devices from phones to fridges are recording everything from our locations to our eating habits. In light of these trends, we must ensure that the increasing amount of data that is collected about us remains private and is used only for the purposes that we expect and intend. I will describe new measurement tools and techniques that shed light on the data collection practices in smart homes, the parties involved in collecting our data, what they can learn from our data, and methods we are developing to help protect our privacy in these environments.

Sometimes, these three goals can be at odds, such as when the data that is needed to improve network performance or security contains private information. I will briefly mention several areas where we are redesigning Internet protocols and systems to facilitate a better balance between performance, security and privacy. Finally, the solutions to these challenges are not exclusively the purview of Internet engineers: they ultimately lie at the intersection of systems, networking, machine learning, and Internet policy. I will touch on how each of these areas can play a role as we embark on this new set of exciting grand challenges for the modern Internet.

**Thursday, January 31, 2019**  
**2:30 pm, JCL 390**  
**Host: Mike Franklin**