Internet Traffic Measurement: from Packets to Insight

Cristian Estan
Computer Science and Engineering Department
University California, San Diego

McBryde Hall - Room 655
4:00 - 5:00pm
Wednesday, March 3, 2004

Abstract:

One of the main reasons for the success of the Internet is its service model that emphasizes flexibility. While this freedom enabled the widespread deployment of the applications popular today such as email and the Web, it has also greatly complicated the task of administering these networks. To understand how a network is being used, or whether it is being abused, an administrator must inspect the flow of packets and "infer" the intent of users and applications. Existing measurement solutions either lack the necessary detail, do not scale up to the speeds of today's networks or are not flexible enough to keep up with the ever changing traffic mix. I will present two approaches to improve the state of the art.

My first approach is to develop fast and accurate algorithmic building blocks that allow routers to collect better measurement data. For example, it is often necessary to identify large flows of traffic, the "heavy hitters". I will present multistage filters which quickly and scalably identify heavy-hitters. A second useful building block scalably estimates the number of active flows or IP addresses using a family of bitmap algorithms. I will show theoretical and experimental evaluations of the effectiveness of these building blocks.

My second approach is to improve the flexibility of offline analysis through a new method of traffic characterization. The conventional approach is a static analysis specialized to capture flows, applications, or network-to-network traffic matrices. By contrast, my analysis dynamically and automatically produces hybrid traffic definitions that match the underlying usage. I will describe a publicly available tool called AutoFocus that I built to implement this analysis, and its use on various production networks to infer such varied phenomena as new worms, denial of service attacks, routing changes, and traffic periodicities.

After receiving his bachelor's and master's degrees in Computer Science from the Technical University of Cluj-Napoca, Cristian Estan worked from 1995 to 1998 at the Romanian Educational Computer Network running the second largest node. After attending the PhD program at Cornell University for a year and working at a Silicon Valley startup developing virtualization based software infrastructure for hosting companies, he joined the PhD program of UCSD in 2000. He received his PhD for work under professor George Varghese for work on Internet traffic measurement and analysis which produced new algorithms and systems with uses in network operations and security.

Candidate for Faculty Position in Computer Science