

Balachander Krishnamurthy

AT&T Labs – Research

11 Aprile Ore 11:30 Aula 1.6

Cat and Mouse: Content Delivery Tradeoffs in Web Access

Generating a privacy footprint on the Internet

11 Aprile Ore 15:00 Aula 1.6

Trends in unwanted traffic and possible solutions

Cat and Mouse: Content Delivery Tradeoffs in Web Access

This talk examines extraneous material (e.g., advertisements) on Web pages that may be viewed as undesirable by some users and the increasing requirement to register before accessing a Web site. Users have predictably reacted to this by avoiding downloading and/or rendering of extraneous material and bypassing registration. Such tension between content owners and users has resulted in a cat and mouse game. I will report on a recent measurement-based study characterizing the nature of extraneous content, manner of distribution, its impact on performance as perceived by users, the effectiveness of blocking mechanisms, and attempts to circumvent registration. Preliminary results show that extraneous content exists on a majority of popular pages and that a 25-30% reduction in downloaded objects and bytes with corresponding latency reduction can be attained by blocking such content. A majority of popular sites require some form of registration and they can be bypassed. I will discuss possible future stages in this cat and mouse game.

[this is joint work with Craig Wills, WPI. i presented this at WWW in May 2006]

Generating a privacy footprint on the Internet

The issue of privacy on the Internet is an important topic. In prior work we characterized traffic deemed as unwanted by Web clients examining the impact of extraneous content, primarily advertisements, and the resulting increase in servers contacted, objects, bytes and latency. Here, we examine how information related to individual users is aggregated as a result of browsing seemingly unrelated Web sites. Our goal is to examine the privacy diffusion on the Internet, the rising number of hidden transactions, and the potential for a few sites to be able to construct a profile of individual users.

We carry out an extensive study of thousands of Web sites across numerous categories, countries, and languages. We present a freely downloadable extension for the Firefox browser, that we developed to carry out the measurements for our study. For a given set of Web sites, we define and generate a privacy footprint, which allows us to assess and compare the diffusion of privacy information across a wide variety of sites. We also study the effectiveness of existing and new techniques to reduce this diffusion. Our results show that the size of the privacy footprint is a legitimate cause for concern across all sets of sites that we study.

[this is joint work with Craig Wills and I presented this at Internet Measurement Conference Oct 2006]

Trends in unwanted traffic and possible solutions

Unwanted packets have been at the heart of many of the problems on the Internet in the last few years. Small scale attacks on individual network nodes have since spread to every layer of the protocol stack. Large numbers of compromised machines are used to launch a wide range of distributed attacks. Spam has clogged networks and tied up the productivity of many users while lowering the value of email communication. Over the last several years a wide range of solutions have been proposed to the various attacks in the literature. This tutorial presents a taxonomy of the attacks as well as the variety of existing and proposed mechanisms to deter the attacks. As targets of attacks we examine routers, links, the protocol infrastructure, and popular applications. We explore the different forms of attacks: probes, denial of service, worms, spam etc. For each of the attacks, we examine a range of solutions. While there have been legal and social solutions offered, we concentrate on the technical portion of the solution space ranging from prevention, establishing identity, intrusion and anomaly detection, deflection, filtering, and traceback.

Balachander Krishnamurthy has been with AT&T Labs--Research since his PhD. His main focus of research of late is in the areas of unwanted traffic, Internet measurements, and Internet protocols. He has authored and edited a dozen books, published over 70 papers, holds eighteen patents, and has given invited talks in over thirty countries. He co-founded the successful Internet Measurement Conference in 2001 and created the new workshop SRUTI (Steps to Reducing Unwanted Traffic on the Internet, <http://www.usenix.org/events/byname/sruti.html>). He has been on the thesis committee of several PhD students, on the board of several journals, collaborated with over sixty researchers worldwide, and given tutorials at several industrial sites and conferences. His most recent book, co-authored with Mark Crovella, "Internet Measurements: Infrastructure, Traffic and Applications" (525pp, John Wiley) was published in July 2006 and is the first detailed book on Internet Measurement. His last book 'Web

Protocols and Practice' (672 pp, Addison-Wesley, co-authored with Jennifer Rexford) is the first in-depth book on the technology underlying the World Wide Web, and has been translated into Portuguese, Japanese, Russian, and Chinese.

Bala is homepageless but many of his papers can be found at <http://www.research.att.com/~bala/papers>