

Detecting Network Neutrality Violations with Causal Inference

Mukarram Bin Tariq, Murtaza Motiwala, Nick Feamster, Mostafa Ammar
School of Computer Science, College of Computing.
Georgia Institute of Technology, Atlanta, GA
{mtariq, murtaza, feamster, ammar}@cc.gatech.edu

ABSTRACT

We present NANO, a system that detects when ISPs apply policies that discriminate against specific classes of applications, users, or destinations. Existing systems for detecting discrimination are typically specific to an application or to a particular discrimination mechanism and rely on active measurement tests. Unfortunately, ISPs can change discrimination policies and mechanisms, and they can evade these tests by giving probe traffic higher priority. NANO detects ISP discrimination by passively collecting performance data from clients. To distinguish discrimination from other causes of degradation (*e.g.*, overload, misconfiguration, failure), NANO establishes a causal relationship between an ISP and observed performance by adjusting for confounding factors. NANO agents deployed at participating clients across the Internet collect performance data for selected services and report this information to centralized servers, which analyze the measurements to establish causal relationship between an ISP and performance degradations. We have implemented NANO and deployed clients in a controlled environment on Emulab. We run a combination of controlled experiments on Emulab and wide-area experiments on PlanetLab that show that NANO can determine the extent and criteria for discrimination for a variety of discrimination policies and applications.

Categories and Subject Descriptors: C.2.3 [Computer Communication Networks]: Network Operations, Network Management

General Terms: Management, Measurement

Keywords: Network Neutrality, Causal Inference

1. INTRODUCTION

Network neutrality states that ISPs remain neutral to how they forward user traffic, irrespective of content, application, or sender [9]; ISPs may *discriminate* against certain subsets of users or services. Rather than taking a stance in this debate, we aim to make the policies of Internet service providers more *transparent* to end users, so that they can detect when ISPs degrade performance or connectivity for some subset of users or applications.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CoNEXT'09, December 1–4, 2009, Rome, Italy.

Copyright 2009 ACM 978-1-60558-636-6/09/12 ...\$10.00.

Because discrimination can take many forms, detecting it is difficult. ISPs have been interfering with TCP connections for BitTorrent and other peer-to-peer applications [8]; recently, British Telecom throttled video content [4] after previously demanding compensation from content providers such as BBC for increased traffic due to their content [3]; Cox Communications also said that it planned to begin throttling peer-to-peer traffic [1]. Other types of discrimination may include blocking specific ports, shaping traffic for specific services, or enforcing traffic quotas.

Existing mechanisms for detecting ISP discrimination actively probe ISPs to test for specific cases: Glasnost detects spurious TCP reset packets of BitTorrent connections [8], Beverly *et al.* study port-blocking [24], and NVLens detects prioritization by observing the type-of-service field in ICMP time-exceeded messages [29]. These tools detect specific classes of discrimination, but they have several drawbacks. First, they are *specific* to either the application (*e.g.*, BitTorrent) or the mechanism that the ISP is using to discriminate (*e.g.*, resetting TCP connections, or setting TOS bits). Second, they rely primarily on *active probes*, which are typically detectable, making it possible for an ISP to either block or prioritize them. Because discrimination may vary depending on the application or the mechanism, and ISPs can evade detection mechanisms that rely on active probes, users need detection tools that rely primarily on observations of *in situ* network traffic.

We present the design, implementation, and controlled evaluation of Network Access Neutrality Observatory (NANO), a system that infers the extent to which an ISP's policy causes performance degradations for a particular service. Instead of trying to determine whether an ISP is discriminating using a particular mechanism, NANO tries to infer whether there are differences in performance achieved through a particular ISP when compared to other ISP(s) for a given service. NANO tries to establish a causal relationship between an ISP's policy and the *observed* degradation of performance for a service using only passively collected data. Because NANO directly uses the observed performance of the service, it is difficult for ISPs to evade NANO inference, while at the same time discriminate against a service to degrade its performance. NANO's techniques apply to general performance metrics and can thus apply to many services and applications. For example, throughput can be used to characterize the performance for both Web traffic (including pages, embedded content, video, etc.) and non-Web traffic (*e.g.*, FTP, BitTorrent). Similarly, jitter and loss rate can characterize the performance of many real-time services, such as interactive voice, video, or gaming traffic.

NANO's design draws inspiration from statistical epidemiology: Just as epidemiologists seek to determine whether a particular drug might be responsible for the improved health of a patient, we seek to determine whether a particular ISP affects performance degradation. The challenge in establishing causality is that many *confound-*

