

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-*




Figure 1. NANO Architecture.

ing factors may be the underlying cause for the observed outcome. Many factors other than the ISP may affect the performance of a particular service or application. For example, a service may be slow (e.g., due to overload at a particular time of day). A service might be poorly located relative to the customers of the ISP. Similarly, a service may be fundamentally unsuitable for a particular network (e.g., Internet connectivity is not suitable for VoIP applications in many parts of the world). Similarly, the performance might depend on software or hardware, or other network peculiarities.

NANO identifies when service performance differs across ISPs but confounding factors are equal. A big challenge in designing NANO is to identify the confounding factors and create an environment where all confounding factors are equal or independent of the ISP or service performance. Although these goals are difficult to achieve, NANO can infer causal relationships by adjusting for confounding variables on passively observed data. Applying this approach has two main requirements: (1) enumerating the confounding factors and collecting data for the possible values of these variables, and (2) establishing a “baseline” level of service performance for a given set of values for the confounding variables that serves as a point of comparison. NANO’s client-side software, NANO-Agent, collects and reports performance data to NANO-Servers regarding their traffic, as well as various meta-data (e.g., the CPU load on the machine at the time, the operating system, the type of connection, etc.) as shown in Figure 1. NANO then analyzes this performance data to quantify the causal relationship between an ISP’s policy and the observed service degradation.

We have implemented the NANO-Agent and Server and made the NANO-Agent available for download [20]. We have evaluated NANO in a controlled environment; we emulate access network ISPs on Emulab, where clients perform HTTP and BitTorrent downloads from hundreds of PlanetLab nodes across the Internet; some ISPs in our setup discriminate while others remain neutral. We demonstrate that, even when the distribution of performance from the discriminating ISPs may look similar to the distribution of performance from the neutral ISPs, NANO detects discrimination, estimates the total causal effect on the performance of the services, and determines the discrimination criteria. Our goal in this paper is to describe the NANO techniques and describe the implementation and controlled evaluation as a proof-of-concept. We do not yet have a sufficient deployment to infer ISP discrimination in real networks, but the NANO project Web site [20] provides the participating clients with other useful performance statistics. With a more extensive deployment, we hope to ultimately report on general discrimination practices across ISPs.

This paper is organized as follows. Section 2 defines and motivates the problem, provides definitions, and articulates the chal-

lenges. Section 3 provides background on causal inference and formulates ISP discrimination detection as a statistical causal inference problem. Section 4 describes the design of NANO and Section 5 describes the implementation. Section 5.2 evaluates the accuracy, sensitivity, and scalability of NANO. Section 6 lists various open issues with NANO. Section 7 discusses related work, and Section 8 concludes.

2. PROBLEM AND MOTIVATION

Problem statement. We aim to detect whether an ISP causes performance degradation for a service when compared to performance for the same service through other ISPs.

Definitions. A *service* is an “atomic unit” of discrimination (e.g., a group of users or a network based application). *Discrimination* is an ISP policy to treat traffic for some subset of services differently such that it causes degradation in performance for the service. Metrics for performance may be service-specific. We say that an ISP *causes* degradation in performance for some service (i.e., that it discriminates against some service) if we can establish a causal relation between the ISP and the observed degradation. For example, an ISP may discriminate traffic for a particular application (e.g., Web search), traffic for a particular domain, or traffic carrying particular type of media, such as video or audio, such that performance for these services degrades.

Challenges. Detecting discrimination is challenging for the following reasons.

1. *The mechanism for discrimination may not be known.* Although existing tools for detecting network neutrality all assume that either the mechanism for discriminating against traffic or the application being discriminated against is known, this is generally not the case. Users often do not even know whether an ISP might be discriminating certain subsets of traffic. These users need methods for detecting discrimination that do not rely on testing for specific discrimination types.
2. *The baseline performance for a service in an ISP is not known.* Users do not know what the “baseline” performance is for a given service through their ISP, so detecting when the performance is degraded, potentially as a result of discrimination is difficult. We propose one approach to establish baseline performance in Section 4.2.
3. *Many factors can cause performance degradation.* Any tool that detects discrimination must identify the ISP—as opposed to any other possible factor—as the underlying cause of discrimination. An industry source recently expressed skepticism about the effectiveness of existing tools: “However, one ISP industry source, who asked not to be identified, questioned whether the tools would accurately point to the cause of broadband problems. ‘Spyware or malware on computers can affect browser performance, and problems with the wider Internet can cause slowdowns, the source said.’” [11]. It is precisely this problem—adjusting for such external causes—that we tackle.

We believe that NANO is the first technique that can isolate such discrimination from other confounding factors, without *a priori* knowledge of an ISP’s discrimination policy. NANO relies on knowledge of confounding variables, but these are not difficult to enumerate using domain knowledge. NANO uses monitoring agents to collect values for the confounding variables.

