

NANO: Network Access Neutrality Observatory

Mukarram Bin Tariq, Murtaza Motiwala, Nick Feamster
{mtariq,murtaza,feamster}@cc.gatech.edu

ABSTRACT

We present *NANO*, a system that establishes whether performance degradations that services or clients experience are caused by an ISP’s discriminatory policies. To distinguish discrimination from other causes of degradation (e.g., overload, misconfiguration, failure), *NANO* uses a statistical method to estimate causal effect. *NANO* aggregates passive measurements from end-hosts, stratifies the measurements to account for possible confounding factors, and distinguishes when an ISP is discriminating against a particular service or group of clients. Using simulation we demonstrate the promise of *NANO* for both detecting discrimination and absolving an ISP when it is not discriminating.

1. Introduction

In late 2005, Ed Whitacre sparked an intense debate on *network neutrality* when he decried content providers “using his pipes [for] free”. Network neutrality says that end users must be in “control of content and applications that they use on the Internet” [1], and that the ISPs respect that right by remaining *neutral* in treating the traffic, irrespective of its content or application. This paper does not take a stance in this debate, but instead studies a technical question: Can users in access networks detect and quantify discriminatory practices of an ISP against a particular group of users or services? We define any practice by the ISP that degrades performance or connectivity for a service as discrimination or violation of network neutrality. We refer to such violations as discrimination. We argue that, regardless of whether or not discrimination is ultimately deemed to be acceptable, the network should be *transparent*; that is, users should be able to ascertain the behavior of their access ISPs.

Unfortunately, because ISP discrimination can take many forms, detecting it is difficult. Several ISPs have been interfering with TCP connections for BitTorrent and other peer-to-peer applications [3]. Other types of discrimination may include blocking specific ports, throttling bandwidth, or shaping traffic for specific services, or enforcing traffic quotas. Existing detection mechanisms *actively* probe ISPs to test for specific cases of discrimination: Glasnost [3], automates detection of spurious TCP reset packets, Beverly *et al.* [8] present a study of port-blocking, and NVLens [10] detects the use of packet-forwarding prioritization by ISPs by examining the TOS bits in the ICMP time exceeded messages. The main drawback of these mechanisms is that each is specific to one type of discrimination; thus, each form of discrimination requires a new test. Worse yet, an ISP may either block or prioritize active probes associated with these tests, making it difficult to run them at all.

If end users could instead somehow detect discrimination by observing the effect on service performance using pas-

sive, in-band methods, the detection mechanism would be much more robust. Unlike active measurements, ISPs cannot prioritize, block, or otherwise modify in-band measurements. To achieve this robustness, we use a black-box approach: we make no assumptions about the mechanisms for implementing discrimination and instead use statistical analysis primarily based on *in situ* service performance data to quantify the *causal relationship* between an ISP’s policy and the observed service degradation.

In this paper, we present the design for Network Access Neutrality Observatory (*NANO*), a system that infers the extent to which an ISP’s policy causes performance degradations for a particular service. *NANO* relies on participating end-system clients that collect and report service performance measurements for a service. Establishing such a causal relationship is challenging because many *confounding factors* (or variables) that are unrelated to ISP discrimination can also 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-the-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).

A necessary condition for inferring a valid causal relationship is to show that *when all the other factors are equal*, a service performs poorly when accessed from an ISP compared to another ISP. The main challenge in designing *NANO* is to create an environment where all other factors are in fact equal. Creating such an environment requires (1) enumerating the confounding factors; (2) establishing a “baseline” level of performance where all factors besides the confounding variables are equal. Unfortunately, the nature of many confounding factors makes it difficult to create an environment on the real Internet where all other factors, except for an ISP’s discriminative policy and service, would be equal. Instead, to correctly infer the causal relationship, we must adjust for the confounding factor by creating strata of clients that have “similar” values for all factors except for their access network. Our approach is based on the theory of causal inference, which is applied extensively in other fields, including epidemiology, economics, and sociology.

This paper is organized as follows. In Section 2 we overview necessary background for establishing a causal relationship between ISP policy and service performance degradation and formalize the problem. In Section 3, we describe the steps in the causal inference, the confounding variables for the problem, and the *NANO* architecture for collecting and processing the necessary data. Section 4 presents simulation-based results, and Section 5 concludes with a discussion of open issues and a research agenda.

2. Background and Problem Formulation

In this section, we formalize the definitions and basic concepts used for establishing ISP discrimination as the cause of service degradation. We describe the concept of service, service performance, ISP, and discrimination; the inference of causal effect, how it relates to association and correlation; and finally, approaches for quantifying causal effect. We also formalize the application of causality to detecting ISP discrimination.

2.1 Definitions

Service and Performance. A *service* is the “atomic unit” of discrimination. An ISP may discriminate against traffic for a particular service, e.g., Web search, traffic for a particular domain, or particular type of media, such as video. Such traffic may be identifiable using the URL or the protocol. Similarly, ISPs may target specific applications, e.g., VoIP, or peer-to-peer file transfers. *Performance*, the outcome variable, is specific to the service. For example, we use server response time for HTTP requests, loss, and jitter for VoIP traffic, and average throughput for peer-to-peer traffic.

ISP and Discrimination. *Discrimination* against a service is a function of ISP policy. The performance for a service depends on both the properties of the ISP’s network, e.g., its location, as well as the policy of treating the traffic differently. Thus, an objective evaluation of ISP discrimination must adjust for the ISP’s network as a confounding factor. To differentiate an ISP’s network from its discrimination policy, we use the ISP *brand* or *name* as the causal variable referring to the ISP’s discrimination policy. In the rest of the paper, when we use ISP as the cause, we are referring to the ISP policy or the brand with which the policy is associated.

We aim to detect whether a certain practice of an ISP results in poorer performance for a service compared to other similar services or performance for the same service through other ISPs. If an ISP’s policy of treating traffic differently does not result in degradation of performance, we do not consider it as discrimination.

2.2 Background for Causal Inference

Statistical methods offer tools for causal inference that have been used in observational and experimental studies [6, 7]. *NANO* draws heavily on these techniques. In this section, we review basic concepts and approaches for causal inference, and how they relate to inferring ISP discrimination.

Causal Effect. The statement “ X causes Y ” means that if there is a change in the value of variable X , then we expect a change in value of variable Y . We refer to X as the *treatment variable* and Y as the *outcome variable*.

In the context of this paper, accessing a particular service through an ISP is our treatment variable (X), and the observed performance of a service (Y) is our outcome variable. Thus, treatment is a binary variable; $X \in \{0, 1\}$, $X = 1$ when we access the service through the ISP, and $X = 0$ when we do not (e.g., access the service through an alternative ISP). The value of outcome variable Y depends on the performance metric and the service for which we are mea-

suring the performance.

The goal of causal inference is to estimate the effect of the treatment variable (the ISP) on the outcome variable (the service performance). Let’s define *ground-truth* value for the outcome random variable as G_X , so that G_1 is the outcome value for a client when $X = 1$, and G_0 is the outcome value when $X = 0$. We will refer to the outcome when not using the ISP ($X = 0$) as the *baseline*—we can define baseline in a number of ways, as we describe in more detail in Section 3.1.2.

We can quantify the *average causal effect* of using an ISP as the expected difference in the ground truth of service performance between using the ISP and the baseline.

$$\theta = \mathbb{E}(G_1) - \mathbb{E}(G_0) \quad (1)$$

Note that to compute the causal effect, θ , we must observe values of the outcome both under the treatment and without the treatment.

Association vs. Causal Effect. In a typical *in situ* dataset, each sample presents only the value of the outcome variable either under the treatment, or under the lack of the treatment, but not both; e.g., a dataset about users accessing a particular service through one of the two possible ISPs, ISP_a and ISP_b , will comprise data of the form where, for each client, we have performance data for either ISP_a or ISP_b , but not both. Such a dataset may thus be incomplete and therefore not sufficient to compute the causal effect, as shown in Equation 1.

Instead, we can use such a dataset to compute correlation or association. Let’s define *association* as simply the measure of observed effect on the outcome variable:

$$\alpha = \mathbb{E}(Y|X = 1) - \mathbb{E}(Y|X = 0) \quad (2)$$

It is well known that association is not a sufficient metric for causal effect, and in general $\alpha \neq \theta$.

Example. Tables 1(a) and (b) illustrate the difference between association and causal effect using an example of eight clients ($a-h$). The treatment variable X is binary; 1 if a user uses a particular ISP, and 0 otherwise. For simplicity, the outcome (Y) is also binary, 1 indicating that a client observes good performance and 0 otherwise; both α and θ are in the range $[-1, 1]$.

Table 1(a) shows an *in situ* dataset. In this dataset, clients $a-d$ do not use the ISP in question and clients $e-h$ use the ISP. Note that for each sample, only one or the other outcome is observable. The association value in this dataset is $\alpha = -3/4$. If we use the association value as an indicator of causal effect, we would infer that using the ISP causes a significant negative impact on the performance.

Table 1(b), on the other hand, presents the ground-truth performance values, G_0 and G_1 , as the performance when not using the ISP and performance when using the ISP for the same client, respectively. These values could be obtained by either subjecting the client to the two cases, or through an oracle. For this set of clients, the true average causal effect $\theta = 1/8$, which is quite small, implying that in reality, the choice of ISP has no or little effect on the performance for these clients. Although the *in situ* dataset is consistent with the ground-truth, i.e., $Y = G_X$, there is a clear discrepancy between the observed association and the true causal effect.

	(a)		(b)		(c)	
	Original Dataset		Ground Truth (Oracle)		Random Treatment	
	X	Y	G_0	G_1	X	Y
a	0	1	1	1	1	1
b	0	1	1	1	0	1
c	0	1	1	1	1	1
d	0	1	1	1	0	1
e	1	0	0	0	1	0
f	1	0	0	0	0	0
g	1	0	0	0	0	0
h	1	1	0	1	1	1
	$\alpha = -3/4$		$\theta = 1/8$		$\alpha = 0$	

Table 1: (a) Observed Association (α) in a passive dataset (b) True causal effect (θ) in an example dataset: $\alpha \neq \theta$. (c) Association converges to causal effect under random treatment assignment: $\alpha \approx \theta$.

2.3 Approaches for Estimating Causal Effect

This section presents two techniques for estimating the causal effect, θ . The first, random treatment, involves an active experiment, where we randomly assign the treatment to the clients and observe the association. The second, adjusting for confounding variables, is a passive technique, where we work with only an *in situ* dataset and estimate the overall causal effect by aggregating the causal effect across several small strata.

1. Random Treatment. Because the ground-truth values (G_0, G_1) are not simultaneously observable, we cannot estimate the true causal effect (Eq. 1) from an *in situ* dataset alone. Fortunately, if we assign the clients to the treatment in a way that is independent of the outcome, then under *certain conditions*, association is an unbiased estimator of causal effect. This property holds because when X is independent of G_X , then $\mathbb{E}(G_X) = \mathbb{E}(G_X|X) = \mathbb{E}(Y|X)$; see [9, pp. 254–255] for a proof. In Table 1(c) we randomly assign a treatment, 0 or 1, to the clients and see that association, α , converges to the true causal effect, θ .

For association to converge to causal effect with random treatment, all other variables in the system that have a causal association with the outcome variable must remain the as we change the treatment. In the case of the example above, association will converge to true causal effect under random treatment, if and only if the original ISP and the alternative ISP are both similar except for their discrimination policy.

Random treatment is difficult to emulate in the Internet for two reasons. First, it is difficult to make users switch to an arbitrary ISP, because not all ISPs may offer services in all geographical areas, the users may be contractually bound to a particular ISP, and asking users to switch ISPs is inconvenient for users. Second, if changing the ISP brand also means that the users must access the content through a radically different network which could affect the service performance, then we cannot use the mere difference of performance seen from the two ISPs as indication of interference: the association may not converge to causal effect under these conditions because the independence condition is not satisfied. This situation is called *operational confounding*: changing the treatment inadvertently or unavoidably changes a confounding variable.

2. Adjusting for Confounding Variables. Because it is difficult to emulate random treatment on the real Internet and control operational confounding, we need to find a way to

adjust for the effects of confounding variables. *NANO* uses the well-known stratification technique for this purpose [6].

Confounding variables are the extraneous variables in the inference process that are correlated with both the treatment and the outcome variables. As a result, if we simply observe the association between the treatment and the outcome variables, we cannot infer causation or lack of it, because we cannot be certain whether the change is due to change in the treatment variable or a change in one or more of the confounding variables.

With stratification, all samples in a stratum are *similar* in terms of values for the confounding variables. As a result, X and G_X are independent of the confounding variables within the stratum, essentially creating conditions that resemble random treatment. Thus, the association value within the stratum converges to causal effect, and we can use association as a metric of causal effect within a strata.

2a. Challenges. This approach presents several challenges. First, we must enumerate the confounding variables and collect sufficient data to help disambiguate the true causal effect from the confounding effects. Second, we must define the stratum boundaries in a way that satisfies the above conditions. Unfortunately, there is no automated way to enumerate all the confounding variables for a given problem; instead, we must rely on domain knowledge. Section 3 addresses these challenges.

2b. Formulation. In the context of *NANO*, we have multiple ISPs and services; we wish to calculate the causal effect $\theta_{i,j}$ that estimates how much the performance of a service j , denoted by Y_j , changes when it is accessed through ISP i , versus when it is not accessed through ISP i . Let Z denote the set of confounding variables, and s a stratum as described above. The causal effect $\theta_{i,j}$ is formulated as:

$$\theta_{i,j}(s; x) = \mathbb{E}(Y_j | X_i = x, Z \in \mathbb{B}(s)) \quad (3)$$

$$\theta_{i,j}(s) = \theta_{i,j}(s; 1) - \theta_{i,j}(s; 0) \quad (4)$$

$$\theta_{i,j} = \sum_s \theta_{i,j}(s) \quad (5)$$

$\mathbb{B}(s)$ represents the range of values of confounding variables in the stratum s . $\theta_{i,j}(s)$ represents the causal effect within the stratum s . A key aspect is the term $\theta_{i,j}(s; 0)$ in Equation 4: it represents the *baseline* service performance, or the service performance when the ISP is *not* used; we define this concept in more detail in Section 3.1.2. Note that the units for causal effect are same as for service performance, so we can apply simple thresholds to detect discrimination.

2c. Sufficiency of Confounding Variables. Although there is no simple or automatic way to enumerate all the confounding variables for a problem, we can test whether a given list is sufficient in the realm of a given dataset. To do so, we predict the value of the outcome variable using a non-parametric regression function, $f()$, of the treatment variable, X , and the confounding variables, Z , as $\hat{y} = f(X; Z)$. We then compare the predicted value with the value of outcome variable observed in the given dataset, y , using relative error, $|y - \hat{y}|/y$. If X and Z are sufficient to define the outcome Y , then the prediction error should be small.

3. NANO System Design

This section explains how *NANO* performs causal inference, enumerates the confounding variables required for this inference, and describes the system architecture for collecting and processing the relevant data.

3.1 Establishing the Causal Effect

Estimating the causal effect for a service degradation involves three steps. First, we stratify the data. Next, we estimate the extent of causal impact of possible ISP interference within each stratum and across the board. Finally, we try to infer the criteria that the ISP is using for discrimination.

3.1.1 Stratifying the data

To stratify the data, *NANO* creates bins (i.e., ranges of values) along the dimensions of each of the confounding variables, such that the value of the confounding variable within the bin is (almost) constant. The bin size depends on the nature of the confounding variable. As a general rule, we create strata such that there is a bin for every unique value of categorical variables; for the continuous variables, the bins are sufficiently small, that the variable can be assumed to have essentially a constant value within the stratum. For example, for a confounding variable representing the client browser, all the clients using a particular version and make of the browser are in one stratum. Similarly, we create one hour strata along the time-of-the-day variable.

We use simple correlation to test whether the treatment variable and the outcome variable are independent of the confounding variable within a stratum. We combine adjacent strata if the distribution of the outcome variable conditioned on the treatment variable is identical in each of the stratum; this reduces the total number of strata and the number of samples needed.

3.1.2 Establishing the baseline performance

A thorny aspect of Equation 4 is the term $\theta_{i,j}(s; 0)$, which represents the *baseline* service performance, or the service performance when the ISP is *not* used. This aspect raises the question: What does it mean to not use ISP i to access service j ? It could mean using another ISP, k , but if ISP k is also discriminating against service j , then $\theta_{k,j}(s; 1)$ will not have the (neutral) ground-truth baseline value. To address this problem, *NANO* takes $\theta_{i,j}(s; 0)$ as the average service performance when not using ISP i , calculated as: $\sum_{k \neq i} \theta_{k,j}(s; 1) / (n_s - 1)$, where $n_s > 2$ is the number of ISPs for which we have clients in stratum s .

An important implication of defining the baseline in this way is that *NANO* is essentially comparing the performance of a service through a particular ISP against the average performance achieved through other ISPs, while adjusting for the confounding effects. If all or most of the ISPs across which *NANO* obtains measurements are discriminating against a service, it is not possible to detect such discrimination using the above definition of baseline; in this case, discrimination becomes the *norm*. In such cases, we might consider using other definitions of discrimination, such as the comparing against the best performance instead of the average, or using a performance model of the service ob-

tained from laboratory experiments or mathematical analysis as the baseline.

3.1.3 Inferring the discrimination criteria

NANO can infer the discrimination criteria that an ISP uses by using simple decision-tree based classification methods. For each stratum and service where *NANO* detects discrimination, *NANO* assigns a negative label, and for each stratum and service where it does not detect discrimination, it assigns a positive label. *NANO* then uses the values of the confounding variables and the service identifier as the feature set and uses the discrimination label as the target variable, and uses a decision-tree algorithm to train the classifier.

The rules that the decision tree generates indicate the discrimination criteria that the ISP uses, because the rules indicate the boundaries of maximum information distance between discrimination and the lack of it.

3.2 Confounding Variables

Confounding variables are the extraneous factors in inferring whether an ISP’s policy is discriminating against a service; these variables correlate, either positively or negatively, with both the ISP brands and service performance. Because there is no automated way to enumerate these variables for particular problem, we must rely on domain knowledge. In this section, we describe three categories of confounding variables and explain how they correlate with both the ISP brands and the service performance. In Section 3.3, we describe the specific variables that we collect to adjust for these confounding variables.

Client-based. Client-side applications, as well as system and network setup, can confound the inference. The particular application that a client uses for accessing a service might affect the performance. For example, in the case of HTTP services, certain Web sites may be optimized for a particular Web browser and perform poorly for others. Similarly, certain Web browsers may be inherently different; for example, at the time of this writing, Opera, Firefox, and Internet Explorer use different number of simultaneous TCP connections, and only Opera uses HTTP pipelining by default. For peer-to-peer traffic, various client software may experience different performance. Similarly, the operating system and the configuration of the client’s computer and local network, as well as a client’s service contract, can affect the performance that the client perceives for a service.

We believe that the above variables also correlate with ISP brand, primarily because the ISP may serve particular communities or localities. As an example, we expect that Microsoft’s Windows operating system may be more popular among home users, while Unix variants may be more common in academic environments. Similarly, certain browsers may be more popular among certain demographics and localities than other.

Network-based. Various properties of the Internet path, such as location of the client or the ISP relative to the location of the servers on the Internet, can cause performance degradation for a service; such degradation is not discrimination. Similarly, a path segment to a particular service




Figure 1: NANO System Architecture.

provider might not be sufficiently provisioned, which could degrade service. If we wish to not treat these effects as discrimination, we should adjust for the path properties.

Time-based. Service performance varies widely with time-of-day due to changes in utilization. Further, the utilization may affect both the ISPs and the service providers, thus confounding the inference.

3.3 Data Collection and Analysis Architecture

An important aspect of *NANO* is collecting the necessary data for facilitating inference. This section describes the criteria for data collection, the features that we collect, and, finally, the data collection mechanism.

1. Criteria. The criteria for data collection has two parts. First, the feature should quantify the treatment variable, the outcome variable, or the values of the confounding factors. Second, the data should be unbiased.

The first criterion helps us determine a set of features for which to collect data; this list is explained below. We can collect many of these features through active or passive monitoring. The second criterion, however, suggests that we must take care that the measurements are not biased. As we discussed in Section 1, ISPs may have the incentive to interfere with identifiable active measurements to deter inference of discrimination or improve their rankings. Similarly, we believe that while we could use the data directly from service providers as the “baseline” service performance, such information could be biased in the favor of the service provider. Therefore, to the extent possible, *NANO* relies on passive measurements to determine the values of the features.

2. Mechanism. Given the nature of the confounding factors and the desire to collect data passively, we believe the best place to collect this data is using monitoring agents at clients. Figure 1 shows the system architecture. The primary source of data for *NANO* are client-side agents installed on computers of voluntarily participating clients (*NANO*-Agents). Each agent continuously monitors and reports the data to the *NANO* servers. We are developing two versions of this agent. The first is a Web-browser plug-in that can monitor Web-related activities, and the second is a packet-level sniffer that can access more fine-grained information from the client machines.

3. Dataset Features. The *NANO*-Agent collects three sets of features for the confounding factors, corresponding to the three classes of confounding variables (Section 3.2)

First, the *NANO*-Agent collects features that help identify the client setup, including the operating system, basic system configuration and resource utilization on the client machine. Second, *NANO*-Agents perform active measurements to a corpus of diverse benchmark sites (PlanetLab nodes) to establish the topological location of the clients and their ISPs. These measurements include periodic short and long transfers with the benchmark sites. These measurements are similar in spirit to ones used by many Internet coordinate systems [5]. *NANO* uses this information to establish the topological properties of the ISP and stratify ISPs with similar topological location to adjust path properties factor. Finally, all the data is time-stamped to allow adjustment for the time-of-day factor.

To identify the treatment variable, i.e., ISP brand, we use the IP address of the client and look it up against *whois* registry servers. To determine the performance for each service, the *NANO*-Agent monitors and logs the information about the ongoing traffic from the client machine for the services that *NANO* monitors. The sniffer version of the agent logs the network round-trip time (RTT) measurements to various destinations for small and large packets. It also collects unsampled flow statistics for the ongoing flows to determine the throughput, and also maintains the applications associated with each flow. The latter is used to disambiguate the performance differences that might be associated with particular applications. The *NANO*-Agent tags this information with a service identifier that it infers by inspecting the packet payloads (e.g., by looking for regular-expression `google.com/search?q=` in the HTTP request message to identify search service), or, if possible, by looking at the protocol and port numbers.

4. Simulation

To illustrate how *NANO* can detect ISP discrimination against a particular service, we evaluate the technique in simulation for a specific example. A rigorous validation of the approach will ultimately require a real deployment where there is less control over confounding variables; we discuss this issue and various others in more detail in Section 5.

Our simulation setup comprises three ISPs, ISP_A , ISP_B , and ISP_C , that provide connectivity for two services, S_1 and S_2 for their respective clients. The clients use one of the two types of applications, App_1 and App_2 to access the services S_1 and S_2 . Performance for the service S_1 is sensitive to the choice of application: clients perceive better performance when using application App_1 , compared to using App_2 . The performance of service S_2 is agnostic to the choice of application. The distribution of clients using the two types of applications is different across the ISPs (e.g., due to demographics). In particular, we set the distribution of App_1 to be 60%, 10%, and 90%, across the three ISPs, respectively. This distribution makes the client application a confounding variable because its distribution correlates with both the ISP and service performance.

We set up the experiment such that ISP_B throttles the bandwidth for all of its clients that access service S_1 . To achieve this, we configure the request rates from the clients and the available throttled bandwidth between the ISPs and the services to achieve certain expected utilization levels. In

(a) Association				
	Service S ₁		Service S ₂	
Baseline	7.68		2.67	
ISP _B	8.60		2.71	
Association	0.92		0.04	
(b) Causal Effect				
	Service S ₁		Service S ₂	
Confounding Var.	App ₁	App ₂	App ₁	App ₂
Baseline	9.90	2.77	2.61	2.59
ISP _B	11.95	7.95	2.67	2.67
Causal Effect	2.0	5.18	0.06	0.12

Table 2: Simulation Results: Estimating the causal effect between ISP_B and S₁. All the numbers are request completion times in seconds.

particular, we configure the links so that the expected utilization on all links is about 40%, except for the traffic from ISP_B to service S₁, which faces about 90% utilization.

We aim to detect discrimination by ISP_B against the service S₁, and establish a causal relationship. We compute the association and causal effect using Equation 2 and Equation 4, respectively. We use the average response time seen through ISP_A and ISP_C, combined, as the baseline.

Table 2 presents the association and causal effect estimated for this simulation. Table 2(a) shows that ISP_B has little to no association for either service S₁ or S₂. However, as we stratify on the application variable, we see in Table 2(b) that ISP_B has significant causal effect for both applications for service S₁, but there is still no causal effect for service S₂. This example shows that, for this case, *NANO* can establish a causal relationship where one exists (ISP_B and service S₁), and rule out where one does not exist, e.g., between ISP_B and S₂.

5. Summary and Research Agenda

We presented Network Access Neutrality Observatory (*NANO*), a system for inferring whether an ISP is discriminating against a particular service. We have examined only basic criteria for discrimination in a simulation environment; in future work, we will evaluate *NANO*'s effectiveness in the wide area, for a wider range of possible discrimination activities, and in adversarial settings where ISPs may attempt to evade detection. In this section, we pose several questions that are guiding our ongoing work.

How can *NANO*-agents be deployed? *NANO* relies on participating clients to collect the data needed to perform causal inference. PlanetLab and CPR [4] nodes are our initial deployment candidates, but in the long run, we wish to incentivize home users to deploy *NANO*-Agents. Because *NANO* inference engine must exclude all extraneous factors, including transient faults to establish ISP discrimination, *NANO*-Agents can additionally act as a network troubleshooting and diagnostics application that users might find useful.

How can *NANO*-agents protect user privacy while still exposing enough client-side information to expose discrimination? Because some of the measurements that *NANO*-Agent collects can lead to invasion of user privacy, *NANO* stores the data in a stratified form and does not maintain any client-identifiable data (e.g., client IP addresses or

search queries). Further, we are instrumenting *NANO* to give users full control over the services that the a *NANO*-agent can monitor. Finally, we are investigating ways of distributed inference, were we may be able to mitigate the need for aggregating the data at a central repository for inference, instead, the *NANO* server can act as a coordinator and the clients infer the discrimination in a peer-to-peer fashion.

Is *NANO* general enough to detect diverse, evolving, and adversarial discrimination policies? ISPs may continue to evolve their policies for distinguishing between various services. One such policy is imposing quotas for Internet traffic from a client [2]; in the future, ISPs may extend this policy by exempting traffic to certain *partner* sites from such quotas, thereby creating a discriminatory environments. We believe that *NANO* can quickly detect such new policies and infer the criteria if we measure sufficient metrics about the state of the network and service.

How much data is needed to perform inference? *NANO* requires a collection of sample data inputs from each stratum to be able to adjust for each confounding variable. The greater the variance of the confounding variables, the more data samples are needed to adjust for each of them and establishing confidence bounds. While statistics theory does help determine the number of samples needed for each stratum, each variable will be distributed differently across the set of clients and ISPs, so it is difficult to determine how many clients would need to run *NANO*-Agents to allow sufficient confidence levels for inference. We expect to understand this better as we deploy these agents.

Acknowledgements

We would like to thank Mostafa Ammar at Georgia Tech and Joseph Hellerstein at UC Berkeley for their valuable feedback that helped improve several aspects of our work. This work is supported in part by NSF Awards CNS-0643974, CNS-0721581, and CNS-0721559.

REFERENCES

- [1] A Guide to Network Neutrality for Google Users. <http://www.google.com/help/netneutrality.html>.
- [2] Comcast Excessive Use Policy. <http://help.comcast.net/content/faq/Frequently-Asked-Questions-about-Excessive-Use>.
- [3] Glasnost: Bringing Transparency to the Internet. <http://broadband.mpi-sws.mpg.de/transparency/>.
- [4] CPR: Campus-Wide Network Performance Monitoring and Recovery. <http://www.rnoc.gatech.edu/cpr/>, 2006.
- [5] F. Dabek, R. Cox, F. Kaashoek, and R. Morris. Vivaldi: A decentralized network coordinate system. In *Proc. ACM SIGCOMM*, Portland, OR, Aug. 2004.
- [6] N. Jewell. *Statistics for Epidemiology*. Chapman & Hall/CRC, 2004.
- [7] J. Pearl. *Causality: Models, Reasoning, and Inference*. Cambridge University Press, 2000.
- [8] S. B. Robert Beverly and A. Berger. The internet's not a big truck: Toward quantifying network neutrality. In *Passive & Active Measurement (PAM)*, Louvain-la-neuve, Belgium, Apr. 2007.
- [9] L. Wasserman. *All of Statistics: A Concise Course in Statistical Inference*. Springer, 2003.
- [10] Y. Zhang, Z. M. Mao, and M. Zhang. Ascertaining the Reality of Network Neutrality Violation in Backbone ISPs. In *Proc. 7th ACM Workshop on Hot Topics in Networks (Hotnets-VII)*, Calgary, Alberta, Canada., Oct. 2008.