

(How Much) Does a Private WAN Improve Cloud Performance?

Todd Arnold[†], Ege Gürmeriçliler[†], Georgia Essig[†], Arpit Gupta[†],
Matt Calder[‡], Vasileios Giotsas[°], Ethan Katz-Bassett[†]

[†]Columbia University, [‡]Microsoft, [°]Lancaster University

Abstract—The construction of private WANs by cloud providers enables them to extend their networks to more locations and establish direct connectivity with end user ISPs. Tenants of the cloud providers benefit from this proximity to users, which is supposed to provide improved performance by bypassing the public Internet. However, the performance impact of cloud providers’ private WANs is not widely understood.

To isolate the impact of a private WAN, we measure from globally distributed vantage points to two large cloud providers, comparing performance when using their worldwide WAN and when instead using the public Internet. The benefits are not universal. While 48% of our vantage points saw improved performance when using the WAN, 43% had statistically indistinguishable median performance, and 9% had better performance over the public Internet. We find that the benefits of the private WAN tend to improve with client-to-server distance, but the benefits (or drawbacks) for a particular vantage point depend on specifics of its geographic and network connectivity.

I. INTRODUCTION

Internet routing can lead to suboptimal performance (*e.g.*, high latency or low throughput). Internet routes traditionally traverse multiple Autonomous Systems (ASes), and each AS has only limited visibility into or control over the routing of other ASes along a route. The Border Gateway Protocol (BGP), the Internet’s inter-AS routing protocol, provides limited visibility into alternate routes available at ASes and makes decisions agnostic to performance [1]. Combined, the Internet’s structure, the autonomous policies of individual ASes, and the lack of visibility and coordination can lead to inflated latencies [2] and congestion [1].

Partly to overcome the aforementioned inefficiencies, cloud and content providers often build their own private Wide Area Networks (WANs). They deploy Points of Presence (PoPs) around the world, peer promiscuously with other networks [3], and build measurement [4, 5] and control systems [1, 6] to select high performing paths. These private WANs allow cloud and content providers to bypass the public Internet for much of the route between their servers and many users [7, 8], which reduces transit costs, improves control over routing, utilization, and performance [1].

How much performance benefit does the extensive investment and unified control of a private WAN provide compared to the public Internet? Knowledge of the benefits can inform tenants’ decisions on how to deploy cloud-based services and can provide insights into the effectiveness of the performance-aware control systems developed by cloud providers [1, 6].

Measuring the impact of a private WAN on performance has inherent challenges. The first is isolating the impact of using the private WAN. As a straw-man, we could compare the performance of two separate cloud providers, one that sends tenant traffic across its private WAN and one that does not. However, this does not isolate the impact of a private WAN: cloud providers differ across various aspects of their deployment, such as each provider’s peers and datacenter locations. The other factors will impact performance and measurements, possibly more than the use of a private WAN.

The second challenge is achieving representative measurement coverage, as performance can vary widely around the world. Prior studies assessed other aspects of cloud performance [9] from PlanetLab [10], which had hundreds of Vantage Points (VPs), mainly hosted at academic institutions. More recent studies used only a few dozen VPs hosted in datacenters [11, 12]. Cloud providers also provide measurements of private WAN performance versus the public Internet for their services [13], but their results are too aggregated to provide insights into any variation or to explain any observations. While these approaches provide initial insights into the performance of cloud providers, they only provide partial results in narrow and possibly biased contexts.

In this paper, we overcome these challenges and provide, to our knowledge, the first comprehensive study that isolates latency differences when using a private WAN versus using the public Internet. We measure from Speedchecker [14], a commercial platform offering measurements from hundreds of thousands of devices, located in thousands of networks around the world (§V). Our measurements target servers we host in Google Cloud and Amazon Web Services (AWS), which both recently began offering the choice of routing via their private WAN (Google’s Premium Tier networking and AWS Global Accelerator) or routing via the public Internet (Google’s Standard Tier and AWS’s normal cloud offering) (§II).

This paper makes the following contributions. First, we show that our Speedchecker VPs provide more comprehensive coverage than established platforms (§V). They provide access to measurements from ASes estimated to host 91% of the world’s Internet users. Second, we assess the expanse of Google’s WAN (§VI). We find that 80% of routes used by VPs’ traffic enters the WAN within 400 km of the VPs location, and half of our VPs’ traffic spends 10% or less of its end-to-end route on the public Internet. Third, we quantify

the performance difference between using a private WAN versus the public Internet to serve users from the cloud (§VII) for Google and AWS. The benefits are not universal. While 48% of our VPs saw improved performance when using the WAN, 9% had better performance when using the public Internet. The remaining 43% had statistically indistinguishable median performance. For brevity, we focus on latency results but throughput measurements (not included) show similar trends. Fourth, given that some VPs see limited benefit to using the private WAN, we investigate factors that impact the performance differences (§VIII). We find that the benefits of a private WAN tend to improve with client-to-server distance, but the benefits (or drawbacks) to a particular VP depend on factors such as geographic location and network connectivity.

II. BACKGROUND

A. Cloud Provider WANs

In the last 10 years, cloud and content provider networks, including Google, Microsoft, Amazon, and Facebook, expanded their global Internet footprint. They built massive private WANs and deployed PoPs globally where they peer with other ASes to improve performance and reduce cost [3]. The resulting path diversity affords providers additional opportunity to optimize traffic in and out of their network.

To utilize the increased connectivity, providers rely on *traffic engineering* to control traffic flow. Providers can control egress traffic routing behavior by overriding BGP’s default decisions [1], replacing it altogether [6], and using hot/cold potato routing [15]. Ingress traffic engineering works by selectively advertising IP prefixes at PoPs, prepending, and steering requests using DNS or other means [5].

Recently, Google introduced tiered networking services for its cloud infrastructure tenants [16, 17]. It uses traffic engineering to offer two networking tiers: (1) *Premium Tier* attempts to improve performance by maximizing the portion of the end-to-end path spent on Google’s private WAN, and (2) *Standard Tier* relies on the public Internet. For inbound traffic, Google advertises its Premium Tier IP addresses from PoPs worldwide to all peers, whereas it advertises a datacenter’s Standard Tier addresses only from PoPs relatively near the datacenter. For outbound traffic, Google generally carries Premium Tier traffic across its WAN to exit near the client, whereas Standard Tier traffic exits the WAN near the datacenter. §VI-A presents our detailed measurements of this behavior.

AWS provides the ability to choose between their private WAN and the public Internet to reach hosted services as well. Their recent offering, Global Accelerator [18], operates similarly to Google’s Premium Tier. The primary difference is that a given flow always egresses AWS’s network at the same PoP where it ingresses, while with Google traffic may egress at another location (see §VI-A).

B. Measurement Infrastructures

To understand performance differences, we use both data plane measurements and control plane information.

Speedchecker. Speedchecker is a global measurement platform deployed in home routers, mobile phones, and PCs [14]. Speedchecker has over 1M VPs, of which we observed over 56K are available at any given time. It allows users to run measurements (*e.g.*, `ping`, `traceroute`, `HTTP GET`) using an API. Speedchecker is our primary data source in this work.

RIPE Atlas. RIPE Atlas is a measurement platform operated by the RIPE Regional Internet Registry [19]. Its deployment comprises small hardware devices connected to volunteers’ networks. Similar to Speedchecker, it lets users run measurements using an API. Unlike Speedchecker, it does not support `HTTP GET` measurements to specific targets.

Citrix Intelligent Traffic Management. Citrix ITM is a JavaScript-based, in-browser, measurement platform similar to those employed by some cloud providers (*e.g.*, Microsoft’s Odin [5]). It is embedded in popular web pages and executes whenever a visitor loads the page causing, for example, the browser to fetch objects and report download times.

Looking Glass Servers and Periscope. A looking glass is a diagnostic tool which allows remote users to run basic measurements from an Internet Service Provider (ISP)’s router. Common operations include `ping`, `traceroute`, and viewing of BGP routes. We use Periscope [20], which consolidates access for multiple looking glass portals.

III. GOALS AND REQUIREMENTS

In this measurement study, we leverage Google’s tiered networking service offering, AWS’s Global Accelerator, and Speedchecker’s extensive Internet measurement infrastructure to answer the following three questions:

Goal 1: How do cloud providers use their private WAN? (§VI) To better understand the advantages of using private WANs over the public Internet for traffic delivery, it is critical to understand the degree to which cloud providers use their WANs to bypass the public Internet.

Goal 2: How much does performance differ? (§VII) Google regularly publishes the difference in performance between its Standard Tier and Premium Tier using Citrix ITM [13]. However, the results are aggregated across all VPs and are only to a single datacenter, showing only an overall benefit and masking details about how performance varies around the world. A deeper understanding requires identifying how the performance difference varies across users, datacenters, and their relative locations.

Goal 3: What factors contribute to this difference? (§VIII) Understanding the specific performance differences can affect the decision of a tenant service about whether to use the private WAN or the public Internet, or small cloud providers deciding whether they need to invest more in expanding their own private WAN. Thus, it is important to understand what factors contribute to any difference in performance.

Requirements. For this measurement study, we need the ability to isolate the use of a private WAN versus the use of the public Internet as the only variable in our setup. Once we are

able to control which is used, we need to collect data from a set of VPs that are representative of the global user population. The data should measure the performance (*e.g.*, throughput and latency) and routing behavior (*e.g.*, traceroutes) for both the private WAN and the public Internet in ways that let us explain performance differences.

IV. METHODOLOGY

To assess how performance differs between a private WAN and the public Internet, we took advantage of new offerings from Google and Amazon (and not yet available from other providers) that provide tenants the choice of using their WAN or the public Internet. We collected and analyzed a large volume of diverse measurements from a global set of VPs. We focus our analysis on measurements to Google datacenters. Google claims to have the world’s largest network [21], is estimated to carry 25% of the Internet’s traffic [21], and deploys sophisticated performance-based traffic engineering [6]. Hence, our measurements should represent a rough best case of the performance benefit of bypassing the public Internet via a private WAN. We supplement with measurements to Amazon datacenters to show that our conclusions hold beyond Google.

A common deployment model for a cloud-based service with a widespread client base is to deploy across multiple datacenters, use global load balancing to terminate TCP connections close to clients, and direct users to the closest datacenter with capacity. Instead, our measurements direct VPs from around the world to a particular target datacenter, without a global load balancer. Using this approach isolates the use of a private WAN versus the public Internet and provides a basis for comparing performance.

A. Collecting Measurement Data

Cloud servers. We created two VMs in each of four Google and seven AWS datacenters, with one VM using the private WAN and one using the public Internet. Each VM runs a minimal web server hosting files of multiple sizes. When a VM receives a valid HTTP GET, we configure it to issue a `traceroute` towards the source IP address of the request. When we began our measurements, only three Google datacenters were available. We added measurements to the Asia-Northeast datacenter when Standard Tier support was announced [22]. When AWS added support for Global Accelerator, we were able to initiate measurements to all datacenters on the same day.

Speedchecker vantage points. Speedchecker offers VPs worldwide (§V) but we have a limited measurement budget per day. To balance global coverage with reliable results, we choose to issue multiple measurements throughout a day from a limited set of VPs. Each day, we use Speedchecker’s API to select VPs from 800 $\langle \text{City}, \text{AS} \rangle$ locations, changing the set daily to maximize $\langle \text{City}, \text{AS} \rangle$ coverage over time and restarting the process after exhausting all available $\langle \text{City}, \text{AS} \rangle$ locations.

We issue measurements from daily locations in ten rounds spread across the day. In each round, we request measurements

from each $\langle \text{City}, \text{AS} \rangle$ to eight destinations: two VMs at four datacenters. For AWS we reduced the number of daily VPs to have the same number of measurements per day to all seven datacenters. For each destination, a VP issues five ICMP echo requests (`ping`), one `traceroute`, and one HTTP GET download of a 10 MB file. We opted for a large file size to better test for speeds required to support sustained file transfers (*e.g.*, HD video transfers require chunks of 4 MB [23]) and to exercise ISP traffic shaping [24]. To ensure that we are focusing solely on routing performance, the VPs access IPv4 addresses instead of URLs. This removes DNS resolution and load balancing, which could affect our observed performance.

For our measurements to Google, which were from Sep 2018 - Jul 2019 and Dec 2019 - Jan 2020, we averaged 197 pings per $\langle \text{City}, \text{AS} \rangle$ location per datacenter. For pings to AWS, which were from Jul 2019 to Dec 2019, we averaged 28 measurements per $\langle \text{City}, \text{AS} \rangle$ location per datacenter.

Statistics and confidence. Our primary choice of metric for performance is median round-trip latency because, unlike mean, median is a good estimator of expected value in non-normal distributions as it is resilient to skew from outliers, particularly common in latency distributions [25, 26]. Minimum values hide the impact of path changes, transient congestion, and time of day effects. Cloud providers often use median to evaluate their own expected performance [5].

The number of measurements and the variance in the underlying performance determine the statistical significance of observed differences in median latency. To make a statistically confident assessment, we use a method that calculates confidence intervals around the difference in medians of two sets of measurements, without assuming that the underlying distributions are Gaussian or have the same shape [27]. This method allows us to conclude whether using the public Internet or using a private WAN is significantly better when the 95% confidence interval for the difference in their median performance across a VP, or set of VPs, does not cross zero.

Processing traceroutes. To convert IP-level traceroutes to AS-level paths, we use techniques from previous work [3, 28]: identifying AS loops and removing Internet eXchange Points (IXPs), private IP addresses, and unresponsive hops. We use the Cymru IP-to-ASN mapping tool [29]. We discard 0.01% of traceroutes where the AS path appears to enter Google, exit, and later re-enter, which may result from path changes, load balancing, or incorrect IP-to-AS mapping of border routers.

Standard Tier traceroutes. Traceroutes from Standard Tier VMs to Internet hosts do not function properly; the destination is reached regardless of the initial TTL set by traceroute, without revealing any intermediate hops. To determine the root cause of this behavior, we conducted tests (using UDP, TCP, and ICMP packets) from multiple Standard Tier VMs to a remote host outside of Google that we control. At the remote host, we used `tcpdump` to confirm (improper but) successful packet arrival. By manipulating the initial TTL value of test packets, we were able to validate that for Standard Tier, a device in Google increments the TTL by 12.




Fig. 1: Geographical distribution of Speedchecker and RIPE Atlas VPs used in our measurements.

B. IP Address Geolocation

Geolocation databases are known to be inaccurate for routers [30]. Instead, we geolocated traceroute IP addresses with an approach similar to the active geolocation technique from RIPE IPmap [31]:

- 1) To derive candidate locations for IP address x , we find its origin ASN ASN_x and find the set of $\langle facility, city \rangle$ locations listed for ASN_x in PeeringDB [32]. If we can decode a location hint from a reverse DNS lookup of x , we only use candidates that match it.
- 2) For each candidate $\langle facility, city \rangle$, we identify a RIPE Atlas VP that is within 40 km of the city and in an AS that either has a presence in the facility or is in the customer cone [33] of an AS in the facility. If multiple VPs fit the criteria, we select one at random. We apply techniques from existing work to avoid VPs with suspicious locations, such as using RIPE Atlas ground truth data [30].
- 3) We ping x from each VP. If a VP measures a round-trip latency of at most 1 ms (maximum distance of 100 km based on the speed of light in fiber), we assume x to be in the VP’s city.

We are particularly interested in identifying where paths enter/exit Google’s WAN. To augment the method described above, we observed that many inter-AS peerings are between routers in the same facility. To validate this intuition, we analyzed all the traceroute paths where we could geolocate both the Google border and the adjacent IP address outside Google. For 90% of these 135K traceroutes, our geolocation placed the two routers at most 100 km apart. Based on this observation, it is usually correct to locate the edge of Google’s WAN as near the location of the first router outside Google on a traceroute. If our geolocation works for the IP address just outside Google, but not for the Google border IP address, we assign the Google address the same location.

Using this insight and our previously described technique, we geolocated the Google border for 81% of our traceroutes. This coverage suffices for our purposes, as the geolocated borders include 99% of cities where Google has a presence (§V) and enable our case studies (§VIII).

V. COVERAGE

The performance of an Internet route depends on the topologies and policies of various networks, as well as the

	Speedchecker	RIPE Atlas	Intersection
Total VPs	1M	27K	–
ASes w/ VPs, total (used in study)	15.9K (7.8K)	3.5K (2.8K)	1.9K (1.3K)
Count of $\langle City, AS \rangle$ used	42.6K	4.7K	0.6K
Internet Users in ASes used [34]	91%	47%	18%
Google edge cities observed	99%	91%	89%

TABLE I: Coverage of Speedchecker and RIPE Atlas. “used” indicates successful measurements in our study.

underlying geography and physical infrastructure. To obtain a rich understanding of the relative performance of routes over the public Internet versus across a private WAN, it is important that our measurements cover much of the Internet. In this section, we demonstrate that Speedchecker provides us with VPs in ASes hosting 91% of Internet users [34], much more than other available measurement infrastructures and that these VPs observe most Google PoPs.

Speedchecker provides better coverage than established platforms. PlanetLab [10] and RIPE Atlas [19] were used to conduct studies similar to ours [9, 35, 36]. However, they offer only limited coverage. PlanetLab is undergoing a slow, sad death, and only 42 sites are currently operational. RIPE Atlas offers 10K active VPs. While the number of RIPE Atlas VPs is impressive, they still do not suffice to capture performance of many Internet users. To quantify this limitation, we use publicly available user population data hosted by Asia-Pacific Network Information Centre (APNIC) Labs [34]. This data uses ad-based measurements to estimate the fraction of Internet user population per AS. According to the APNIC estimates, RIPE Atlas hosts VPs in ASes representing 56.5% of the total Internet user population.

Speedchecker has $35\times$ as many VPs, hosted in $4\times$ as many ASes, compared to RIPE Atlas (table I). RIPE Atlas has very dense coverage of Europe in particular, but sparse coverage in many other parts of the world (Fig. 1). We have measurements from Speedchecker VPs that reside in ASes that host 91% of Internet users, compared to 47% for RIPE Atlas, according to APNIC estimates. Performance can vary across an AS. Because the performance differences we are interested in stem from topological (*e.g.*, where do ASes have PoPs and how are they interconnected) and policy constraints, we would like to measure to/from as many $\langle City, AS \rangle$ locations as possible. To gather the $\langle City, AS \rangle$ for each RIPE Node, we used the geolocation data provided by each probe, and reverse geolocated the city [37]. Speedchecker nodes provide both geolocation data and a city. Speedchecker covers $10\times$ as many $\langle City, AS \rangle$ locations as RIPE Atlas. APNIC only shares population estimates per ASes, so we cannot estimate user coverage at finer granularities.

Speedchecker routes enter Google in most cities where Google has PoPs, which is perhaps not surprising, given Speedchecker’s coverage of end-user ASes globally. Google lists edge locations in 56 cities [38], which we complement with the list of peering facilities where Google indicates a presence in its PeeringDB entry [32]. This brings the total edge locations to 74 cities. As part of our measurement methodology (§IV-A), we issue traceroutes from Speedchecker, then

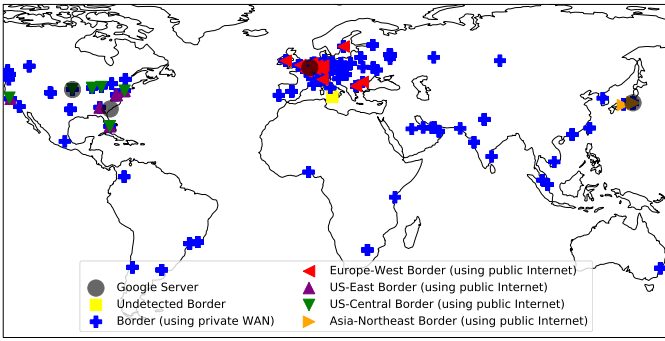


Fig. 2: Google PoP locations observed in our measurements.

identify and geolocate the Google borders they traverse. Collectively, our traceroutes ingress into Google at 73 of the 74 edge cities documented by Google, and 79 total edge cities, meaning that our measurements should exercise much of the performance achievable by Google’s private WAN.

The six cities seen in our measurements but not publicly listed by Google could represent incomplete information published by Google or errors in our approach to geolocating where our measurements enter Google’s WAN. However, these extra cities show up on less than 0.1% of our traceroutes, so they have negligible impact on our overall results. Our approach can introduce errors in two ways. First, our geolocation pins router locations using the speed of light relative to a VP’s location (§IV-B), but it is possible that some VPs are not at the locations they claim despite our efforts to exclude any with suspect locations. Second, the router we geolocate (§IV-B) may not be the border router for Google. For all six cities, we never located a Google IP address in the city, but rather inferred a Google border there because it was the location of the last IP address on a traceroute before Google (§IV-B).

VI. PRIVATE WAN VS INTERNET: ROUTING

A private WAN can shorten the distance traffic travels on the public Internet both geographically by building out PoPs in more locations and in terms of AS hops via increased peering. In this section we measure along these two dimensions, showing both the extensive reach of Google’s WAN and how it shortens paths to clients around the world.

A. Google Edge Locations

Figure 2 depicts the cities where traceroutes from Speedchecker ingress into Google’s network for each datacenter when using Google’s private WAN globally and when using the public Internet to reach the datacenter. Routes that use the public Internet ingress in 2 to 12 cities per datacenter, mainly located near the datacenter, whereas the global private WAN ingresses in 79 cities worldwide.

These geo-distributed ingresses allow traffic from many VPs to enter and exit the WAN near the VP. Figure 3 depicts the distribution of the distance from Speedchecker VPs to the ingress PoP of Google’s network. When using the private WAN, 80% of VPs ingress within 400 km of their location




Fig. 3: Distribution across Speedchecker \langle City, AS \rangle locations of the distance traveled from the VP to the destination Google-hosted server, separated into VP-Ingress (outside Google) and Ingress-Server (inside Google), for our servers hosted in the US-Central datacenter.




Fig. 4: Distribution across Speedchecker \langle City, AS \rangle locations of AS path length for traceroutes using the private WAN versus using the public Internet, for servers hosted in Google’s US-Central datacenter.

and spend most of their route on the WAN. In contrast, paths using the public Internet ingress far from the VPs and close to the server: 90% ingress within 400 km of the server.

Google has direct control over where traffic egresses its WAN, but the ingress location is at the mercy of BGP routing. For Premium Tier, we calculate the distance between the ingress and egress PoP, as a measure of how similar BGP’s performance agnostic selections are to Google’s traffic engineering decisions. In 60% of our traceroutes, traffic ingresses and egresses in the same city and 85% of the time both locations are within 350 km of each other. Standard Tier egress is limited to a small set of PoPs close to the datacenter [17], but outbound traceroutes do not work on Standard Tier (§IV-A).

B. AS Path Lengths

To prevent traffic to Standard Tier VMs from using their private WAN, Google limits the set of available routes and PoPs by selectively advertising the most specific prefixes for VMs’ IP addresses only at regional PoPs. In contrast, Google advertises IP addresses for Premium Tier VMs worldwide. We verified this behavior using 147 Looking Glass servers [20]. For example, in the case of the Standard Tier server in Google’s US-Central datacenter, we observe the most specific IP prefixes advertised in Chicago; for Europe-West it is in Frankfurt.

Figure 4 shows the impact of the advertising policies on logical path length. For Premium Tier, 60% of paths enter directly from the VP’s AS to Google’s network. An additional

32% use a single intermediate AS (e.g., when Google peers with the provider of a VP’s ISP).

Conversely, selective advertisement forces Standard Tier traffic to take the public Internet and traverse longer logical paths. Traceroutes from Standard Tier VMs to VPs do not work (§IV-A) and are not depicted. The majority (60%) of Standard Tier paths traverse two or more intermediate ASes between the VP’s AS and Google, while an additional 34% traverse a single intermediate AS. For the remaining six percent, we observe two main patterns: (1) some IXPs provide layer 2 peering fabrics across remote facilities; and (2) a large AS hosting the VPs, where the VP itself is far from the datacenter, but the AS peers with Google at a regional PoP.

VII. PRIVATE WAN VS. INTERNET: LATENCY

This section presents the performance difference observed by our VPs when accessing servers using the private WAN versus the public Internet. We only present the results for latency, but throughput measurements showed similar trends.

A. Latency by $\langle \text{City}, \text{AS} \rangle$

Figure 5 (red line) shows the cumulative distribution of difference in median latency (as measured by ping), per $\langle \text{City}, \text{AS} \rangle$ location, when using the public Internet versus a private WAN, from our Speedchecker VPs to our server in Google’s US-Central datacenter. Negative values indicate that the public Internet outperformed the private WAN, and positive values indicate that the private WAN performed better. The multi-colored shaded regions in the graph shows the *distributions* of the lower and upper bounds of the 95% confidence intervals for the differences in median latency (§IV-A).

In our measurements, 48% of $\langle \text{City}, \text{AS} \rangle$ locations have a difference with a confidence interval fully above zero (the blue region), indicating better performance via the private WAN, with 22% seeing a latency improvement of at least 15 ms. However, 9% of $\langle \text{City}, \text{AS} \rangle$ locations have a confidence interval fully below zero (the red region), indicating that the public Internet provides lower median latency than the private WAN, with 3% of $\langle \text{City}, \text{AS} \rangle$ locations seeing latency improvements of 15 ms or more. For 43% of $\langle \text{City}, \text{AS} \rangle$ locations (the yellow region), the confidence interval included 0 so our measurements do not establish either as better. While the private WAN provides latency indistinguishable from or better than the public Internet for the majority of our VPs, the latency improvement remains modest until the tail, where 6% of $\langle \text{City}, \text{AS} \rangle$ locations improve by 50 ms or more. Other AWS datacenters show similar trends (not shown).

To demonstrate that large confidence intervals are not overly influencing our results, the black line in Figure 5 includes just the 60% of $\langle \text{City}, \text{AS} \rangle$ locations with a 95% confidence interval width of 20 ms or less. This subset of $\langle \text{City}, \text{AS} \rangle$ locations displays a similar trend to the full set.

Our AWS measurements show similar results, but with fewer $\langle \text{City}, \text{AS} \rangle$ locations benefiting from the WAN. For measurements to the AWS US-East datacenter, 34% of $\langle \text{City}, \text{AS} \rangle$ locations showed better performance using the




Fig. 5: Distribution across Speedchecker $\langle \text{City}, \text{AS} \rangle$ locations of difference in median latency, in ms, to the US-Central datacenter when using the private WAN vs. the public Internet. Negative values indicate the public Internet performed better, and positive values indicate private WAN performed better. The multi-colored shaded area is bounded by the *distributions* of the lower and upper bounds of the 95% confidence intervals for difference of medians (rather than showing confidence intervals around individual points on the CDF). The blue region indicates 95% confidence that the median latency on the private WAN is lower than on the public Internet. The red region indicates 95% confidence that the median latency on the private WAN is higher than on the public Internet. The yellow region indicates confidence intervals that cross 0, where our measurements do not establish either as better with 95% confidence. The gray shaded area shows similar distributions, but limited to $\langle \text{City}, \text{AS} \rangle$ locations where the width of confidence interval is $\leq 20\text{ms}$.

private WAN (confidence interval fully above zero), 18% see a latency improvement of 15 ms or better, and 5% improved by at least 50 ms. The public Internet performed better for 12% of $\langle \text{City}, \text{AS} \rangle$ locations, and 5% saw an improvement of at least 15 ms. For the remaining 54%, the confidence intervals around the difference in median performance crossed or was equal to zero, and so our measurements did not support concluding that either was better than the other. Other AWS datacenters showed similar trends.

B. Latency by Country

The performance of a private WAN compared to the public Internet can vary by region, depending on various factors (§VIII). Figure 6 shows, per country, the difference in median latency when using the public Internet versus the private WAN for three Google and three Amazon datacenters, one each in Europe, the US, and Asia.

Across all three Google datacenters, we can see that virtually all locations in Europe, North America, and South America have small but significant improvements from using the private WAN, and the improvements generally increase as the distance from the VP to the datacenter increases. Africa has a couple countries that show substantial improvement from the use of the private WAN over the public Internet, and a few where the public Internet is significantly better, but most have only a slight improvement. The results may stem from the internal structure of the African continent’s Internet [39]. Surprisingly, Asia and Oceania show both the most benefit from using the private WAN (Figs. 6b, 6c) and

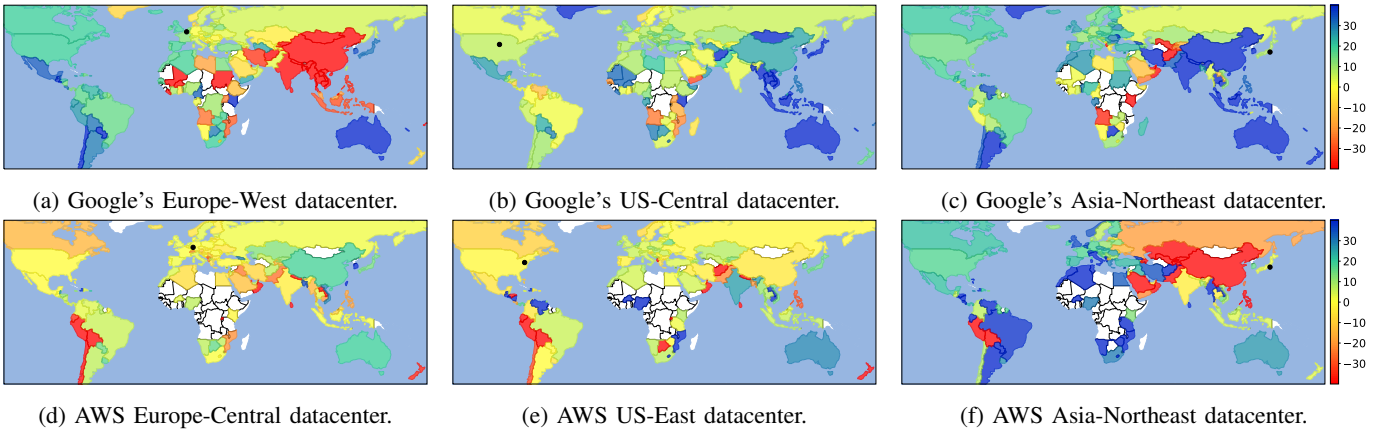


Fig. 6: Difference in median latency, in ms, per country when using the public Internet versus a private WAN to reach servers in different datacenters (black dot) for Google and AWS. Negative values (depicted in red and orange) indicate that the public Internet performed better; positive values (depicted in green and blue) indicate that the private WAN performed better; little to no difference is depicted in yellow. Countries depicted in white did not have enough measurements to make an assessment.

the most penalty (Fig. 6a). Many island nations in these regions have limited routing options, and so performance can be determined by the directness of these routes relative to the datacenter locations. We explain these findings in §VIII. The AWS datacenters, as shown in Figures 6d, 6e, and 6f, do not see such a dramatic benefit or penalty in this region.

C. Comparison with Public Results


Google teamed with Citrix to provide publicly available performance measurements of Standard Tier and Premium Tier for the US-Central datacenter [13]. While the results are publicly available, they are aggregated and do not provide details beyond the percentile measurements. Compared to the results published by Citrix [13], our global Speedchecker measurements observed higher absolute latency values, but smaller differences between the private WAN and the public Internet. The range of latency values reported by Citrix are similar to those we collected from North America and Europe VPs by both RIPE Atlas and Speedchecker, suggesting that Citrix may be biased towards these regions. However, since we cannot break down the Citrix measurements further, we can only say that we see neither the universal nor uniform benefit from use of a private WAN over the public Internet implied by the aggregate result from Citrix.

VIII. CONTRIBUTING FACTORS


A combination of factors impacts whether and by how much a private WAN outperforms the public Internet for a location.

A. WAN Benefit Tends to Increase with Distance

Inter-AS routing can lead to circuitous routes, with the amount of indirectness tending to increase with the AS path length [2]. A global private WAN places most of the path under the control of one AS, which tends to lead to more direct paths [2]. Since latency increases with distance, it makes sense that the absolute inflation incurred by indirect multi-AS paths would tend to increase with distance.




(a) For the AS path length using the public Internet, the regression model has R^2 score of 0.022 and slope of 3.098 ms/hop.




(b) For geographic distance between VP and datacenter, the regression model has R^2 score of 0.018. The slope is 0.002 ms/km, a public Internet “penalty” of 20% relative to the speed of light in fiber, which equates to .01 ms/km. Areas of low data density are due to oceans.

Fig. 7: Difference in latency between the public Internet and private WAN versus path length and geographical distance between $\langle \text{City}, \text{AS} \rangle$ locations and the server, for measurements to Google’s US-Central datacenter.

Figure 7 quantifies the relationship between distance and the performance difference between a private WAN and the public Internet. Both y-axes show the difference in median latency as measured from the VP using the private WAN versus using the public Internet. The x-axes show (a) AS path length of traceroutes from VP to datacenter when using the public Internet (Fig. 7a), and (b) the geographic distance from VP to datacenter (Fig. 7b). Figure 7a shows that the private WAN yields an expected improvement of 3.747 ms per AS



(a) $\langle \text{City}, \text{AS} \rangle$ locations in New Zealand and Australia to Europe datacenters.



(b) $\langle \text{City}, \text{AS} \rangle$ locations in Hong Kong, Singapore, Indonesia, Malaysia, Thailand, Vietnam, India Myanmar, Macao, Cambodia and Philippines to Google's Europe-West and US-Central datacenters.

Fig. 8: Latency comparisons when using the public Internet versus a private WAN for case studies in §VIII-B (Fig. 8a), and §VIII-C (Fig. 8b). The shaded areas are distributions of the lower and upper bounds of the confidence intervals.

hop on the public Internet path. Figure 7b shows that there is an expected benefit of 0.002 ms/km for the private WAN, 20% relative to what can be achieved at the speed of light in fiber (200 km/ms equates to an RTT of 0.01ms/km of client-server distance). At 6000 km (approximately the distance from New York to London), this benefit translates to 12 ms improvement from using the private WAN, a meaningful improvement.

Following the trend in Figure 7, Figure 6 shows, for both Google and AWS, that most countries close to the datacenter have similar performance from either path option. As distance increases, so does the performance difference. There are some notable exceptions, particularly in Figures 6a and 6c, which we will explain in subsequent paragraphs.

B. PoP Proximity Can Make a Difference

Whether a VP is close to a Google PoP, a PoP of a large Internet provider, or both can impact the relative performance between Google's private WAN and the public Internet.

Case study: A nearby PoP helps performance for VPs hosted in Uninet, AS 8151, in Mexico, which observe better performance using the private WAN compared to the public Internet. Uninet uses Telia as a provider, which peers with Google at multiple locations, including in a Kio Networks colocation facility in Mexico. Measurements using the WAN enter Google from Telia in this facility, then Google routes them via Dallas to Chicago. In contrast, the route on the public Internet goes from Mexico to Los Angeles, then Dallas, Kansas City, and Chicago, all via Telia. The traffic finally enters Google's network in Chicago. The Google PoP in Mexico extends its WAN, keeping traffic sourced in Mexico from having to detour to Los Angeles to enter the WAN.

Case study: The lack of a nearby PoP hurts performance

for VPs in Zagreb, Croatia hosted in Croatian Academic and Research Network (CARN), AS 2108. For these VPs, the median latency to Google's US-Central datacenter using the public Internet is less than the median latency using the private WAN. As its provider, CARN uses GÉANT, AS 21320, which routes Premium Tier traffic from Croatia to Milan, Italy via Vienna, Austria. In contrast, Standard Tier measurements pass from GÉANT to its provider, Cogent, in Croatia, and Cogent takes the traffic to US-Central. When available, peer routes (like the one GÉANT learns from Google) are generally preferred to provider routes (like the one learned from Cogent). Because GÉANT and Google only peer in a limited number of locations (and not in Croatia), these peer routes incur a geographic detour and latency penalty. In such cases, global providers like Cogent may have more direct routes.

Case study: The WAN helps Australia but hurts New Zealand.

Whereas most Australian $\langle \text{City}, \text{AS} \rangle$ locations see substantially lower latency using the WAN instead of the public Internet to access European datacenters (purple line Fig. 8a), measurements from New Zealand are slower (Figs. 6a, 6d, red line Fig. 8a). All routes from both countries reach Europe via the US. Google has a PoP in Sydney, Australia, but not in New Zealand (Fig. 2), and all measurements from both countries towards Premium Tier ingress there. However, routing via Australia represents a detour for $\langle \text{City}, \text{AS} \rangle$ locations in New Zealand, as they reach the Standard Tier via either Telstra or Vocus, which both have direct connectivity between New Zealand and the US [40, 41]. The distance from Wellington, New Zealand, to Sydney is 2,224 km, equivalent to 22 ms RTT in fiber, which roughly approximates the performance penalty a number of the $\langle \text{City}, \text{AS} \rangle$ locations experience in our measurements.

For AWS, which also has a PoP in Sydney, measurements from Australia to Europe show a similar behavior, but with less improvement from the private WAN (green line Fig. 8a). $\langle \text{City}, \text{AS} \rangle$ locations in New Zealand have a pronounced improvement from using the public Internet when accessing AWS's datacenter in Europe (blue line Fig. 8a).

C. Impact of Undersea Cables

In recent years, Google invested heavily in the expansion of its undersea cables [21], facilitating their global presence. Our measurements found that Google's cable from Australia yielded performance benefits to users there (Fig. 8a). However, Google's internal routing paths are not necessarily better than those offered by tier-1 transit providers.

Case study: Large transit providers have comparable undersea cable networks to Google's

for VPs in Bergen hosted in GET Norway, AS 41164. We observed similar latency (around 120 ms) to US-Central for both the Standard Tier and Premium Tier traffic. GET's provider is TDC, AS 3292, who peers with Google in Stockholm, Sweden where Premium Tier traffic ingresses into Google's network, whereas the traffic for Standard Tier ingresses Google's network in the US via Telia,

a tier-1 ISP. Telia carries the Standard Tier traffic across the Atlantic Ocean, and Google carries the Premium Tier traffic. In this case, we observe similar performance for both networks.

Case study: Google’s undersea cables can lead to circuitous routes. VPs in Southeast Asia (*e.g.*, Hong Kong, Indonesia, *etc.*) observe better performance to the Europe-West datacenter on the public Internet compared to on Google’s private WAN (Figs. 8b, 6a). Although the traffic for Premium Tier enters Google’s WAN nearby at regional PoPs (*e.g.*, Malaysian VPs enter Google’s WAN within Malaysia), it takes a circuitous route from Southeast Asia to Europe. Many of Google’s undersea cables land in the US; Figures 6b and 8b shows that traffic destined for the US-Central datacenter does experience improved performance from using Google’s private WAN and undersea cables. However, the Premium Tier traffic destined for the Europe-West datacenter goes all the way to the US across the Pacific Ocean before reaching Europe after crossing the Atlantic Ocean (return traffic takes a similarly long return path). In contrast, large ASes in the region (*e.g.*, Tata Communications, GTT, *etc.*) use a shorter path to Europe via West Asia and the Indian Ocean [42, 43]. In general, Google’s US-centric undersea cabling can affect the performance for services hosted in Europe. Google’s network maps [44] show a lack of interconnectivity between Southeast Asia and Europe, as verified in other recent work [11, 12].

Case study: Undersea cables can improve performance depending on the source and destination. Google’s investment in undersea cables in South America and Southeast Asia helps their WAN provide better performance than the public Internet in those areas for VPs accessing datacenters in North America (Fig. 6b) and Asia (Fig. 6c). Figure 6 shows most countries in Europe and North America see little improvement from either provider’s private WAN compared to the public Internet. However, Southeast Asia and South America have small but significant performance improvements from Google’s private WAN over the public Internet. In the same regions, AWS’s private WAN (Figs. 6e, 6f) has mixed results with most countries experiencing little performance gain.

IX. RELATED WORK

Several past studies sought to compare the performance of cloud provider offerings based on common features across multiple locations [9] or how cloud applications react to different what-if scenarios (*e.g.*, dynamic load) based on available hardware configurations [45]. ThousandEyes publishes a report on latency and throughput cloud providers [11, 12], but with orders of magnitude fewer VPs than ours, and their VPs were deployed in cloud datacenters. Other approaches to cloud provider measurement use the provider’s own end-users [5, 46] or CDN to prefix measurements [47], and are based on a random sampling of traffic. While these works have some similar goals, we specifically focus on the impact of a private WAN against using the public Internet.

Previous studies analyzed the performance benefits of multihoming [48, 49]. Our work analyzes the current state of the art strategy of a private WAN with traffic engineering [1, 6].

Several studies examined how routing policy relates to performance. One investigated how routing performance compares across ISPs [36], another analyzed the intradomain intercontinental paths of several cloud networks and found they are more reliable and predictable than Internet paths [35]. Other studies explored circuitous routes on the Internet in general [2] and between mobile networks and cloud providers [50]. We demonstrate that a private WAN decreases the length of AS paths for most users, which tends to improve performance.

One part of a recent study explored private WAN and public Internet performance differences. The study only examined a subset of our dataset: client networks with direct connections to Google on the Premium Tier that reach the Standard Tier via intermediate ASes [51]. We examine the broader performance implications of a private WAN versus the public Internet, and the underlying causes, from multiple cloud providers.

A prior study showed Google’s network moving “closer” to users with respect to AS hops [3]. Our findings in Premium Tier traceroute characteristics confirm this trend, and show Google increased its interconnectivity and now peers with a greater number of user ISPs than seen in previously.

X. CONCLUSIONS

Cloud and content providers make use of private WANs in an effort to improve the network performance of services. We present the first comprehensive measurement study to compare the performance differences between the use of a private WAN and public Internet for large cloud providers. We accomplished this by evaluating Google Cloud’s Premium and Standard network tiers, as well as AWS’s Global Accelerator, from globally distributed VPs. These cloud providers’ wide peering, aggressive WAN expansion, and sophisticated traffic engineering [6] provide an opportune platform to understand the performance a private WAN can offer.

Our results show that while many VPs see performance improvement using a private WAN, the gains are not ubiquitous or uniform – we find that geographic location, network connectivity, and distance from the VP to the datacenter play critical roles in the benefit, or lack thereof, observed by a VP. Having a robust private WAN does not automatically result in providing improved performance over the public Internet; it must also be paired with improved fiber paths and smart routing policy. We find that, although Google and AWS’s private WANs do provide improved performance for many users, there is still room for improvement.

Acknowledgements. We would like to acknowledge the contributions of Jia He and Siao-Ting Wang, whose efforts to support this work we greatly appreciate. We appreciate the valuable feedback from the INFOCOM reviewers. We thank Speedchecker, especially Janusz Jezowicz, for providing us access to their measurement platform. We also thank Chris Blake and Adney Cardoza for their contributions. This work was partly funded by NSF awards CNS-1835253, CNS-1413978, and CNS-1836872. Research supported, in part by, Security Lancaster, H2020 EC CONCORDIA GA #830927.

REFERENCES

- [1] B. Schlinker, H. Kim, T. Cui, E. Katz-Bassett, H. V. Madhyastha, I. Cunha, J. Quinn, S. Hasan, P. Lapukhov, and H. Zeng, "Engineering Egress with Edge Fabric: Steering Oceans of Content to the World," in *Proc. SIGCOMM*, 2017.
- [2] N. T. Spring, R. Mahajan, and T. E. Anderson, "The Causes of Path Inflation," in *Proc. SIGCOMM*, 2003.
- [3] Y.-C. Chiu, B. Schlinker, A. B. Radhakrishnan, E. Katz-Bassett, and R. Govindan, "Are We One Hop Away from a Better Internet?" in *Proc. IMC*, 2015.
- [4] B. Schlinker, I. Cunha, Y.-C. Chiu, S. Sundaresan, and E. Katz-Bassett, "Internet Performance from Facebook's Edge," in *Proc. IMC*, 2019.
- [5] M. Calder, R. Gao, M. Schröder, R. Stewart, J. Padhye, R. Mahajan, G. Ananthanarayanan, and E. Katz-Bassett, "Odin: Microsoft's Scalable Fault-Tolerant CDN Measurement System," in *Proc. USENIX NSDI*, 2018.
- [6] K.-K. Yap, M. Motiwala, J. Rahe, S. Padgett, M. Holliman, G. Baldus, M. Hines, T. Kim, A. Narayanan, A. Jain *et al.*, "Taking the Edge off with Espresso: Scale, Reliability and Programmability for Global Internet Peering," in *Proc. SIGCOMM*, 2017.
- [7] P. Gill, M. Arlitt, Z. Li, and A. Mahanti, "The Flattening Internet Topology: Natural Evolution, Unightly Barnacles or Contrived Collapse?" in *Proc. PAM*, 2008.
- [8] E. Carisimo, C. Selmo, J. I. Alvarez-Hamelin, and A. Dhamdhere, "Studying the Evolution of Content Providers in the Internet Core," in *Proc. TMA*, 2018.
- [9] A. Li, X. Yang, S. Kandula, and M. Zhang, "CloudCmp: Comparing Public Cloud Providers," in *Proc. IMC*, 2010.
- [10] PlanetLab, <https://www.planet-lab.org/>.
- [11] ThousandEyes, "Public Cloud Performance Benchmark Report," 2018, <https://marketo-web.thousandeyes.com/rs/thousandeyes/images/ThousandEyes-2018-Public-Cloud-Performance-Benchmark-Report.pdf>.
- [12] —, "Public Cloud Performance Benchmark Report," 2019, <https://marketo-web.thousandeyes.com/rs/thousandeyes/images/ThousandEyes-2019-Public-Cloud-Performance-Benchmark-Report.pdf>.
- [13] Citrix, "Google Report, Network Tiers," <https://itm.cloud.com/google-reports/>.
- [14] Speedchecker, "ProbeAPI," <http://probeapi.speedchecker.com/>.
- [15] D. McPherson and K. Patel, "Experience with the BGP-4 Protocol," RFC 4277, Jan 2006.
- [16] Prajakta Joshi, "Introducing Network Service Tiers: Your Cloud Network, Your Way," <https://cloudplatform.googleblog.com/2017/08/introducing-Network-Service-Tiers-your-cloud-network-your-way.html>.
- [17] Google, "Google Network Service Tiers," <https://cloud.google.com/network-tiers/>.
- [18] Amazon Web Services, "Introducing AWS Global Accelerator," <https://aws.amazon.com/about-aws/whats-new/2018/11/introducing-aws-global-accelerator/>.
- [19] RIPE NIC, "RIPE Atlas," <https://atlas.ripe.net/>.
- [20] V. Giotsas, A. Dhamdhere, and k. claffy, "Periscope: Unifying Looking Glass Querying," in *Proc. PAM*, 2016.
- [21] B. T. Sloss, "Expanding Our Global Infrastructure With New Regions and Subsea Cables," 2018, <https://blog.google/topics/google-cloud/expanding-our-global-infrastructure-new-regions-and-subsea-cables/>.
- [22] Google, "Network Service Tiers Documentation," <https://cloud.google.com/network-tiers/docs/using-network-service-tiers>.
- [23] F. Li, J. W. Chung, X. Jiang, and M. Claypool, "TCP CUBIC versus BBR on the Highway," in *Proc. PAM*, 2018.
- [24] T. Flach, P. Papageorge, A. Terzis, L. Pedrosa, Y. Cheng, T. Karim, E. Katz-Bassett, and R. Govindan, "An Internet-Wide Analysis of Traffic Policing," in *Proc. SIGCOMM*, 2016.
- [25] R. Padmanabhan, P. Owen, A. Schulman, and N. Spring, "Timeouts: Beware Surprisingly High Delay," in *Proc. IMC*, 2015.
- [26] R. Fontugne, C. Pelsser, E. Aben, and R. Bush, "Pinpointing Delay and Forwarding Anomalies Using Large-Scale Traceroute Measurements," in *Proc. IMC*, 2017.
- [27] R. M. Price and D. G. Bonett, "Distribution-Free Confidence Intervals for Difference and Ratio of Medians," *Journal of Statistical Computation and Simulation*, vol. 72, no. 2, 2002.
- [28] Z. M. Mao, J. Rexford, J. Wang, and R. H. Katz, "Towards an Accurate AS-level Traceroute Tool," in *Proc. SIGCOMM*, 2003.
- [29] Team Cymru, "IP-to-ASN Mapping," <http://www.team-cymru.com/IP-ASN-mapping.html>.
- [30] M. Gharaibeh, A. Shah, B. Huffaker, H. Zhang, R. Ensafi, and C. Papadopoulos, "A Look at Router Geolocation in Public and Commercial Databases," in *Proc. IMC*, 2017.
- [31] RIPE NCC, "RIPE IPmap," <https://ipmap.ripe.net/>.
- [32] PeeringDB, <https://peeringdb.com>.
- [33] M. Luckie, B. Huffaker, A. Dhamdhere, V. Giotsas, and k. claffy, "AS Relationships, Customer Cones, and Validation," in *Proc. IMC*, 2013.
- [34] APNIC, "Visible ASNs: Customer Populations (Est.)," <https://stats.labs.apnic.net/aspop/>.
- [35] O. Haq, M. Raja, and F. R. Dogar, "Measuring and Improving the Reliability of Wide-Area Cloud Paths," in *Proc. of the 26th International Conference on World Wide Web (WWW)*, 2017.
- [36] Ratul Mahajan and Ming Zhang and Lindsey Poole and Vivek Pai, "Uncovering Performance Differences among Backbone ISPs with Netdiff," in *Proc. USENIX NSDI*, 2008.
- [37] K. Esmukov, "Geopy," <https://github.com/geopy/geopy>.
- [38] Google, "Network Edge Locations," <https://cloud.google.com/vpc/docs/edge-locations>.
- [39] A. Formoso, J. Chavula, A. Phokeer, A. Sathiseelan, and G. Tyson, "Deep Diving into Africa's Inter-Country Latencies," in *Proc. INFOCOM*, 2018.
- [40] Telstra, "Network Map," <https://www.telstraglobal.com/company/our-network/network-map>.
- [41] Vocus, "Network Map," <https://www.vocus.co.nz/our-network>.
- [42] GTT, "Network map," <https://www.gtt.net/us-en/our-network/>.
- [43] Tata Communications, "Network Map," <https://www.tatacommunications.com/map/>.
- [44] Google, "Global Locations: Meet our Network."
- [45] Y. Jiang, L. R. Sivalingam, S. Nath, and R. Govindan, "WebPerf: Evaluating What-If Scenarios for Cloud-hosted Web Applications," in *Proc. SIGCOMM*, 2016.
- [46] Calder, Matt and Flavel, Ashley and Katz-Bassett, Ethan and Mahajan, Ratul and Padhye, Jitendra, "Analyzing the Performance of an Anycast CDN," in *Proc. IMC*, 2015.
- [47] R. Krishnan, H. V. Madhyastha, S. Srinivasan, S. Jain, A. Krishnamurthy, T. Anderson, and J. Gao, "Moving Beyond End-to-end Path Information to Optimize CDN Performance," in *Proc. SIGCOMM*, 2009.
- [48] A. Akella, B. Maggs, S. Seshan, A. Shaikh, and R. Sitaraman, "A Measurement-based Analysis of Multihoming," in *Proc. SIGCOMM*, 2003.
- [49] A. Akella, J. Pang, B. Maggs, S. Seshan, and A. Shaikh, "A Comparison of Overlay Routing and Multihoming Route Control," in *Proc. SIGCOMM*, 2004.
- [50] K. Zarifis, T. Flach, S. Nori, D. R. Choffnes, R. Govindan, E. Katz-Bassett, Z. M. Mao, and M. Welsh, "Diagnosing Path Inflation of Mobile Client Traffic," in *Proc. PAM*, 2014.
- [51] T. Arnold, M. Calder, I. Cunha, A. Gupta, H. V. Madhyastha, M. Schapira, and E. Katz-Bassett, "Beating BGP is Harder than we Thought," in *Proc. HotNets*, 2019.