A First Comparative Characterization of Multi-cloud Connectivity in Today's Internet

Bahador Yeganeh[1], Ramakrishnan Durairajan[1], Reza Rejaie[1], and
Walter Willinger[2]

[1] University of Oregon {byeganeh,ram,reza}@cs.uoregon.edu
[2] NIKSUN Inc. wwillinger@niksun.com

Abstract. Today's enterprises are adopting multi-cloud strategies at an unprecedented pace. Here, a multi-cloud strategy specifies end-to-end connectivity between the multiple cloud providers (CPs) that an enterprise relies on to run its business. This adoption is fueled by the rapid build-out of global-scale private backbones by the large CPs, a rich private peering fabric that interconnects them, and the emergence of new third-party private connectivity providers (*e.g.,* DataPipe, HopOne, etc.). However, little is known about the performance aspects, routing issues, and topological features associated with currently available multi-cloud connectivity options. To shed light on the tradeoffs between these available connectivity options, we take a cloud-to-cloud perspective and present in this paper the results of a cloud-centric measurement study of a coast-to-coast multi-cloud deployment that a typical modern enterprise located in the US may adopt. We deploy VMs in two regions (*i.e.,* VA and CA) of each one of three large cloud providers (*i.e.,* AWS, Azure, and GCP) and connect them using three different options: (i) transit provider-based best-effort public Internet (BEP), (ii) third-party provider-based private (TPP) connectivity, and (iii) CP-based private (CPP) connectivity. By performing active measurements in this real-world multi-cloud deployment, we provide new insights into variability in the performance of TPP, the stability in performance and topology of CPP, and the absence of transit providers for CPP.

1 Introduction

Modern enterprises are adopting multi-cloud strategies at a rapid pace. Defined here as end-to-end connectivity between multiple cloud providers (CPs)[3], multi-cloud strategies are critical for supporting distributed applications such as geo-distributed analytics [57,69,68,33,35] and distributed genome sequencing studies at universities [25,12]. Other benefits that result from pursuing such strategies are competitive pricing, vendor lockout, global reach, and requirements for data sovereignty. According to a recent industry report, more than 85% of enterprises have already adopted multi-cloud strategies [39].

Fueled by the deployment of multi-cloud strategies, we are witnessing two new trends in Internet connectivity. First (see Figure 1 (bottom)), there is the emergence of new Internet players in the form of third-party private connectivity

---

[3] This is different from hybrid cloud computing, where a direct connection exists between a public cloud and private on-premises enterprise server(s).

providers (*e.g.,* DataPipe, HopOne, among others [29,51,5]). These entities offer direct, secure, private, layer 3 connectivity between CPs (henceforth referred to as *third-party private* or TPP), at a cost of a few hundreds of dollars per month[4]. TPP routes bypass the public Internet at Cloud Exchanges [21,19,71] where they operate virtualized routers allowing their customers to form virtualized peering sessions with the participating CPs and offer additional benefits to users (*e.g.,* enterprise networks can connect to CPs without owning an Autonomous System Number, or ASN, or physical infrastructure). Second (see Figure 1 (top)), the large CPs are aggressively expanding the footprint of their serving infrastructures, including the number of direct connect locations where enterprises can reach the cloud via direct, private connectivity (henceforth referred to as *cloud-provider private* or CPP) using either new CP-specific interconnection services (*e.g.,* [4,50,28]) or third-party private connectivity providers at colocation facilities. Of course, as shown in Figure 1 (middle), a multi-cloud user can forgo the TPP and CPP options altogether and rely instead on the traditional, best effort connectivity over the public Internet—henceforth referred to as *(transit provider-based) best-effort public (Internet)* (BEP). In terms of routing, CPP and BEP connectivity is offered through default route configurations while TPP routes are enforced via BGP configurations that customers of the TPP network install on their virtual routers.

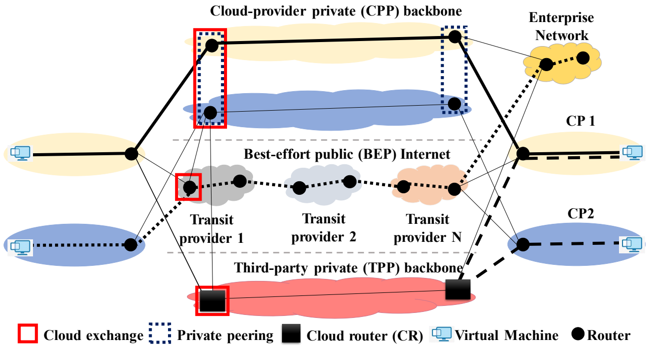

Fig. 1: **Overview of three different multi-cloud strategies. Sample end-to-end measurement paths highlighted using thicker solid, dashed, and dotted lines for CPP, TPP, and BEP options.**

With multi-cloud connectivity being the main focus of this paper, we note that existing measurement techniques are a poor match in this context. For one, they fall short of providing the data needed to infer the type of connectivity (*i.e.,* TPP, CPP, and BEP) between (two or more) participating CPs. Second, they are largely incapable of providing the visibility needed to study the topological properties, performance differences, or routing strategies associated with different connectivity options. Third, while mapping the connectivity from cloud/content providers to users has been considered in prior work (*e.g.,* [9,15,17,20,60,16]

---

[4] See Section 3.4 for more details.

and references therein), multi-cloud connectivity from a cloud-to-cloud (C2C) perspective has remained largely unexplored to date.

This paper aims to empirically examine the different types of multi-cloud connectivity options that are available in today's Internet and investigate their performance characteristics using non-proprietary cloud-centric, active measurements. In the process, we are also interested in attributing the observed characteristics to aspects related to connectivity, routing strategy, or the presence of any performance bottlenecks. To study multi-cloud connectivity from a C2C perspective, we deploy and interconnect VMs hosted within and across two different geographic regions or availability zones (*i.e.,* CA and VA) of three large cloud providers (*i.e.,* Amazon Web Services (AWS), Google Cloud Platform (GCP) and Microsoft Azure) using the TPP, CPP, and BEP option, respectively.

Using this experimental setup as a starting point, we first compare the stability and/or variability in performance across the three connectivity options using metrics such as delay, throughput, and loss rate over time. We find that CPP routes exhibit lower latency and are more stable when compared to BEP and TPP routes. CPP routes also have higher throughput and exhibit less variation compared to the other two options. Given that using the TPP option is expensive, this finding is puzzling. In our attempt to explain this observation, we find that inconsistencies in performance characteristics are caused by several factors including border routers, queuing delays, and higher loss-rates of TPP routes. Moreover, we attribute the CPP routes' overall superior performance to the fact that each of the CPs has a private optical backbone, there exists rich inter-CP connectivity, and that the CPs' traffic *always* bypasses (*i.e.,* is invisible to) BEP transits. In summary, this paper makes the following contributions:

- To the best of our knowledge, this is one of the first efforts to perform a comparative characterization of multi-cloud connectivity in today's Internet. To facilitate independent validation of our results, we will release all relevant datasets [1] (properly anonymized; *e.g.,* with all TPP-related information removed).
- We identify issues, differences, and tradeoffs associated with three popular multi-cloud connectivity options and elucidate/discuss the underlying reasons. Our results highlight the critical need for open measurement platforms and more transparency by the multi-cloud connectivity providers.

2 Background and Related Work

Measuring and understanding the connectivity ecosystem of the Internet has been the subject of a large number of studies over the years [52, and references therein]. Efforts include mapping the (logical) connectivity of the public Internet at the router level (*e.g.,* [13,44,64,10,11]), the POP-level (*e.g.,* [62,65,63]), and the Autonomous System or AS-level (*e.g.,* [73,45]). Other efforts have focused on issues such as the rise of Internet Exchange Points (IXPs) and their effects on inaccuracies of network-layer mapping (*e.g.,* [11,2]), the "flattening" of the Internet's peering structure (*e.g.,* [27,40,22]), and the Internet's physical infrastructure (building repositories of point of presence (POP), colocation, and datacenter locations (*e.g.,* [61,37]), the long-haul and metro connectivity between them (*e.g.,* [24,23,38]), and interconnections with other networks (*e.g.,* [46,43,3]).

More recently, enterprise networks have been able to establish direct connectivity to cloud providers—even without owning an AS number—at Open Cloud Exchanges [21,19] (shown in the red box in Figure 1) via a new type of interconnection service offering called virtual private interconnections [71]. With the advent of such interconnection services, today's large cloud (and content) providers (*e.g.,* Google, Facebook, Microsoft) have experienced enormous growth in both their ingress (*i.e.,* Internet-facing) and mid-gress (*i.e.,* inter-datacenter) traffic. To meet these demands, they are not only aggressively expanding their presence at new colocation facilities but are also simultaneously building out their own private optical backbones [26,36] (see CPP in Figure 1). In addition, connectivity to the CPs at colocation facilities are also available via third-party providers [51,29,5] (TPP in Figure 1) for additional costs (*e.g.,* thousands of dollars for a single, dedicated, private link to CP).

While measuring the peering locations, serving infrastructures and routing strategies of the large content providers has been an active area of research [9,15,17,20,60,16,70] and comparing the performance of CPs and their BEP properties has been the focus of prior efforts [41,30,74,14,18], to the best of our knowledge, ours is one of the first studies to (a) examine and characterize the TPP, CPP, and BEP connectivity options from a C2C perspective, and (b) elucidate their performance tradeoffs and routing issues.

3 Measurement Methodology

In this section, we describe our measurement methodology to examine the various multi-cloud connectivity options, the cloud providers under consideration, and the performance metrics of interest.

3.1 Measurement Setting

As shown in Figure 1, we explore three different types of multi-cloud connectivity options: *TPP* connectivity between CP VMs that bypasses the public Internet, *CPP* connectivity enabled by private peering between the CPs, and *BEP* connectivity via transit providers. To establish TPPs, we deploy cloud routers via a third-party connectivity provider's network. At a high level, this step involves (i) establishing a virtual circuit between the CP and a connectivity partner, (ii) establishing a BGP peering session between the CP's border routers and the partner's cloud router, (iii) connecting the virtual private cloud gateway to the CP's border routers, and (iv) configuring each cloud instance to route any traffic destined to the overlay network towards the configured virtual gateway. To establish CPP connectivity, participating CPs automatically select private peering locations to stitch the multi-cloud VMs together. Finally, we have two measurement settings for BEP. The first setting is between a non-native colocation facility in Phoenix AZ and our VMs through the BEP Internet; the second form of measurement is through the BEP Internet towards Looking Glasses (LGs) residing in the colocation facility hosting our cloud routers.

We conduct our measurements in a series of rounds. Each round consists of path, latency, and throughput measurements between all pairs of VMs (in both directions to account for route asymmetry). Furthermore, the measurements are performed over the public BEPs as well as the two private options (*i.e.,* CPP and TPP). Each connectivity path is enforced by the target address for our measurements (i.e., public IP address for BEP and CPP paths and private IPs VM instances in the TPP case). We avoid cross-measurement interference by tracking the current state of ongoing measurements and limit measurement activities to one active measurement per cloud VM.

3.2 Measurement Scenario & Cloud Providers

For this study, we empirically measure and examine one coast-to-coast, multi-cloud deployment in the US. Our study focuses on connectivity between three major CPs (AWS, Azure, and GCP) as they collectively have a significant market share and are used by many clients concurrently [72]. Using these CPs, we create a realistic multi-cloud scenario by deploying two cloud routers using one of the top third-party connectivity providers' networks; one of the cloud routers is in the Santa Clara, CA region, and one is in the Ashburn, VA region. These cloud routers are interconnected with native cloud VMs from the three CPs. The cloud VMs are all connected to cloud routers with 50Mb/s links. We select the colocation facility hosting the cloud routers based on two criteria: (i) CPs offer native cloud connectivity within that colo, and (ii) geo-proximity to the target CPs datacenters. Cloud routers are interconnected with each other using a 150Mb/s link capacity that supports the maximum number of concurrent measurements that we perform (*i.e.,* 3 concurrent measurements in total to avoid more than 1 ongoing measurement per VM). Each cloud VM has at least 2 vCPU cores, 4GB of memory, and runs Ubuntu server 18.04 LTS. Our VMs were purposefully over-provisioned to reduce any measurement noise within virtualized environments. Throughout our measurement experiments, the VMs CPU utilization always remained below 2%. We also cap the VM interfaces at 50Mb/s to have a consistent measurement setting for both public (BEP) and private (TPP and CPP) routes. We perform measurements between all CP VMs within regions (intra-region) and across regions (inter-region). Additionally, we also perform measurements between our cloud VMs and two LGs that are located within the same facility as our cloud routers in California and Virginia, respectively, and use these measurements as baselines for BEP [5] comparisons.

3.3 Data Collection & Performance Metrics

We conducted our measurements for about a month-long period in the Spring of 2019. The measurements were conducted in 10-minute rounds. In each round, we performed latency, path, and throughput measurements between all pairs of relevant nodes. For each round, we measure and report the latency using 10 *ping* probes paced in 1 second intervals. We refrain from using a more accurate one-way latency measurement tool such as OWAMP as the authors of OWAMP caution its use within virtualized environments [34]. Similarly, paths are measured by performing 10 attempts of *paris-traceroute* using *scamper* [42] towards

---

[5] In Section 5 we highlight that our inter-cloud measurements do not exit the source and destination CP's network.

each destination. We used ICMP probes for path discovery as they maximized the number of responsive hops along the forward path. Lastly, throughput is measured using the *iperf3* tool, which was configured to transmit data over a 10-second interval using TCP. We discard the first 5 seconds of our throughput measurement to account for TCP's *slow-start* phase and consider the median of throughput for the remaining 5 seconds. These efforts resulted in about 48k samples of latency, path, and throughput measurements between each unique src/dst pair and connectivity option.

To infer inter-AS interconnections, the resulting traceroute hops from our measurements were translated to their corresponding AS paths using BGP prefix announcements from Routeviews and RIPE RIS [67,59]. Missing hops were attributed to their surrounding ASN if the prior and next hop ASNs were identical. The existence of IXP hops along the forward path was detected by matching hop addresses against IXP prefixes published by PeeringDB [56] and Packet Clearing House (PCH) [55]. We mapped each ASN to its corresponding ORG number using CAIDA's AS-to-ORG mapping dataset [32]. Lastly, the inter-AS interconnection segments are identified using the latest version of bdrmapIT [3].

3.4 Limitations and Ethical/Legal Considerations

Our study is US-centric and limited by the geographic span of our multi-cloud deployment as well as the number of third-party connectivity providers that we examine. The high cost for connecting multiple clouds using TPP connections prevents us from having a global-scale deployment and performing experiments that involve different TPP providers. For example, for each 1 Gbps link to a CP network, third-party providers charge anywhere from about 300 to 700 USD per month [53,48,58][6]. While limited in scale, the deployment that we consider in this study is nevertheless representative of a typical multi-cloud strategy adopted by modern enterprises with a US-wide footprint [49].

Our study does not raise any ethical issues. Overall, since the goal of this study is to measure and improve multi-cloud connectivity without attributing particular features to any of the utilized third-party providers and CPs, we are not in violation of any of their terms of service. In particular, we obfuscate, and wherever possible, we omit all information that can be used to identify the colocation and third-party connectivity providers. This information includes names, supported measurement APIs, costs, time and date of measurements, topology information, and any other potential identifiers.

4 Characteristics of C2C Routes

In this section, we characterize the performance of C2C routes (*i.e.,* latency and throughput) and attribute the observed characteristics to connectivity and routing.

4.1 Latency Characteristics

CPP routes exhibit lower latency than TPP routes and are stable. Figure 2 depicts the distribution of RTT values (using letter-value plots [31]; see

---

[6] Note that these price points do not take into consideration the additional charges that are incurred by CPs for establishing connectivity to their network.

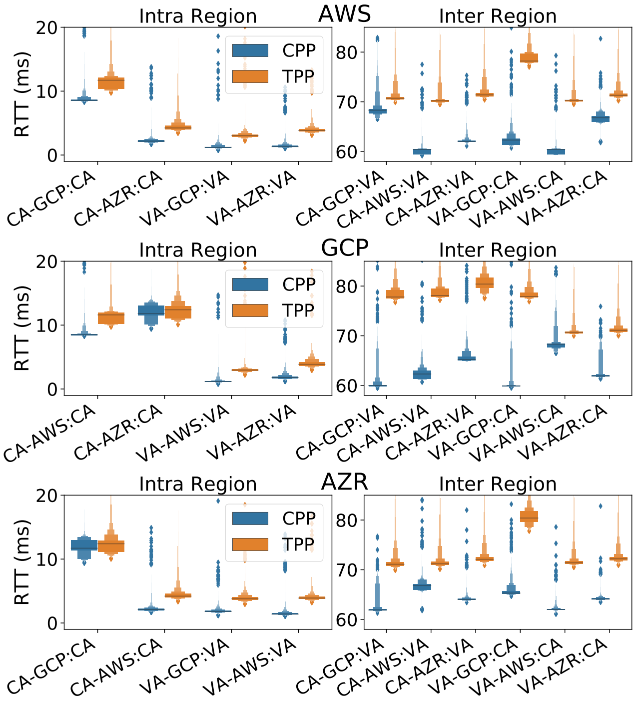

Fig. 2: **Distribution of RTT between AWS, GCP, and Azure for intra (left) and inter (right) region paths.**

Appendix 7.1) between different CPs across different connectivity options. The rows (from top to bottom) correspond to AWS, GCP, and Azure as the source CP, respectively. Intra-region (inter-region) measurements are shown in the left (right) columns, and CPP (TPP) paths are depicted in blue (orange).

The first two characters of the x-axis labels encode the source CP region and the remaining characters encode the destination CP and region. From these figures, we see that CPP routes typically exhibit lower medians of RTT compared to TPP routes, suggesting that CPP routes traverse the CP's optical private backbone. We also observe a median RTT of ∼2ms between AWS and Azure VMs in California which is in accordance with the relative proximity of their datacenters for this region. The GCP VM in California has a median RTT of 13ms to other CPs in California, which can be attributed to the geographical distance between GCP's California datacenter in LA and the Silicon Valley datacenters for AWS and Azure. Similarly, we notice that the VMs in Virginia all exhibit low median RTTs between them. We attribute this behavior to the geographical proximity of the datacenters for these CPs. At the same time, the inter-region latencies within a CP are about 60ms with the exception of Azure which has a higher median of latency of about 67ms. Finally, the measured latencies (and hence the routes) are asymmetric in both directions albeit the median of RTT values shows latency symmetry (<0.1ms). Also, the median of the measured latency between our cloud routers is in line with the published values by third-party connectivity providers, but the high variance of latency indicates that the TPP paths are in general a less reliable connectivity option compared to CPP routes. Lastly, BEP routes for cloud to LG measurements always have an equal or higher median of latency compared to CPP paths with much higher variability (order of magnitude larger standard deviation; results are omitted for brevity).

Why do CPP routes have better latency than TPP routes? In our path measurements, we observe that intra-cloud paths always have a single organization, indicating that regardless of the target region, the CP routes traffic internally towards the destination VM. More interestingly, the majority of inter-cloud paths only observe two organizations corresponding to the source and destination CPs. Only a small fraction (<4%) of paths involves three organizations, and upon closer examination of the corresponding paths, we find that they traverse IXPs and involve traceroutes that originate from Azure and are destined to Amazon's network in another region. We reiterate that single organization inter-CP paths correspond to traceroutes which are originated from GCP's network and do not reveal any internal hops of its network. For the cloud-to-LG paths, we observe a different number of organizations depending on the source CP as well as the physical location of the target LG. The observations range from only encountering the target LG's organization to seeing intermediary IXP hops as points of peering. Lastly, we measure the stability of routes at the AS-level and observe that all paths remain consistently stable over time with the exception of routes sourced at Azure California and destined to Amazon Virginia. The latter usually pass through private peerings between the CPs, and only less than 1% of our path measurements go through an intermediary IXP. In short, we did not encounter any transit providers in our measured CPP routes.

By leveraging the AS/organization paths described in § 3, we next identify the peering points between the CPs. Identifying the peering point between two networks from traceroute measurements is a challenging problem and the subject of many recent studies [46,3,43]. For our study, as mentioned in § 3 above, we utilized *bdrmapIT* [3] to infer the interconnection segment on the collection of traceroutes that we have gathered. Additionally, we manually inspected the inferred peering segments and, where applicable, validated their correctness using (i) IXP address to tenant ASN mapping and (ii) DNS names such as AMAZON.SJC-96CBE-1A.NTWK.MSN.NET which is suggestive of peering between AWS and Azure. We find that *bdrmapIT* is unable to identify peering points between GCP and the other CPs since GCP only exposes external IP addresses for paths destined outside of its network, *i.e., bdrmapIT* is unaware of the source CPs network as it does not observe any addresses from that network on the initial set of hops. For these paths, we choose the first hop of the traceroute as the peering point only if it has an ASN equal to the target IP addresses ASN. Using this information, we measure the RTT between the source CP and the border interface to infer the geo-proximity of the peering point from the source CP. Using this heuristic allows us to analyze each CP's inclination to use hot-potato routing.

Figure 3 shows the distribution of RTT for the peering points between each CP. From left to right, the plots represent AWS, GCP, and Azure as the

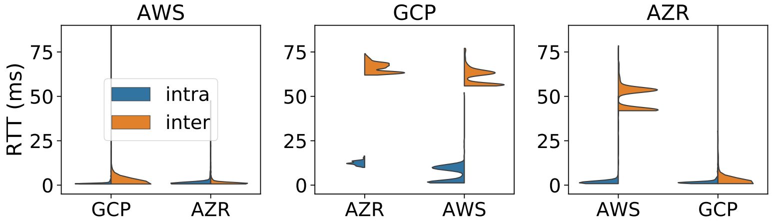

Fig. 3: **Distribution of RTT between source CP and the peering hop.**

source CP. Each distribution is split based on intra (inter) region values into the left/blue (right/orange) halves, respectively. We observe that AWS' peering points with other CPs are very close to their networks and therefore, AWS is employing hot-potato routing. For GCP, we find that hot-potato routing is never employed and traffic is always handed off near the destination region. The bimodal distribution of RTT values for each destination CP is centered at around 2ms, 12ms, 58ms, and 65ms corresponding to the intra-region latency for VA and CA, and inter-region latency to other CPs, respectively. Finally, Azure exhibits mixed routing behavior. Specifically, Azure's routing behavior depends on the target network – Azure employs hot-potato routing for GCP and cold-potato routing for AWS. More specifically, intra-region traffic destined to AWS is delivered through a local peering point while its Virginia-California traffic destined to AWS is handed off in Los Angeles, and for inter-region paths from California to AWS Virginia, the traffic is usually (99%) handed off in Dallas TX and for the remainder is being exchanged through Digital Realty Atlanta's IXP. From these observations, the routing behavior for each path can be modeled with a simple threshold-based method. More concretely, for each path i with an end-to-end latency of l_{ei} and a border latency of l_{bi}, we can infer if source CP employs hot-potato routing if $l_{bi} < \frac{1}{10} l_{ei}$. Otherwise, the source CP employs cold-potato routing (*i.e.,* $l_{bi} > \frac{9}{10} l_{ei}$). The fractions (*i.e.,* $\frac{1}{10}$ and $\frac{9}{10}$) are not prescriptive and are derived based on the latency distributions depicted in Figure 3.

4.2 Throughput Characteristics

CPP routes exhibit higher and more stable throughput than TPP routes. Figure 4 depicts the distribution of throughput values between different CPs using different connectivity options. While intra-region measurements tend to have a similar median and variance of throughput, we observe that for inter-region measurements, TPPs exhibit a lower median throughput with higher variance. Degradation of throughput seems to be directly correlated with higher RTT values as shown in Figure 2. Using our latency measurements, we also approximate loss-rate to be 10^{-3} and 10^{-4} for TPP and CPP routes, respectively. Using the formula of Mathis et al. [47] to approximate TCP throughput[7], we can obtain an upper bound for throughput for our measured loss-rate and latency values.

Using Mathis et al. model, the upper bound of throughput for an MSS of 1460 bytes, a 70ms latency and loss-rate of 10^{-3} (corresponding to the average mea-

---

[7] We do not have access to parameters such as TCP timeout delay and number of acknowledged packets by each ACK to use more elaborate TCP models (*e.g.,* [54]).

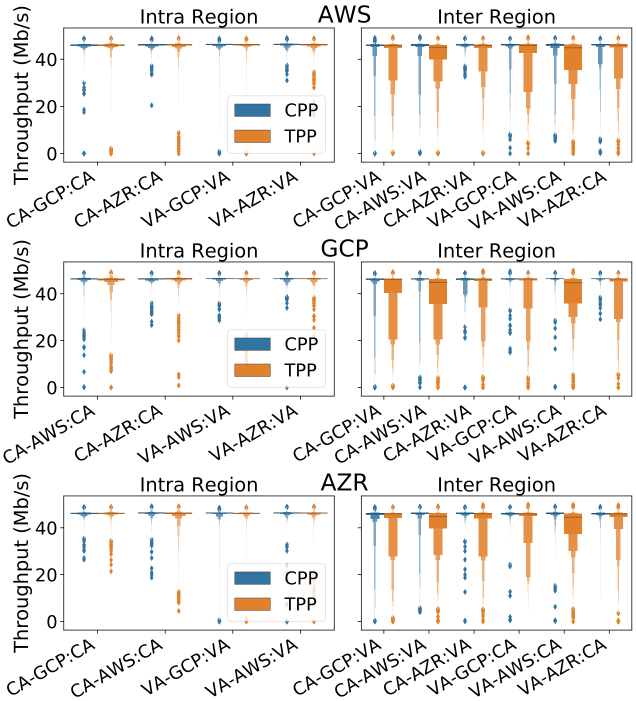

Fig. 4: **Distribution of throughput between AWS, GCL, and Azure for intra (left) and inter (right) region paths.**

sured values for TPP routes between two coasts) is about 53Mb/s. While this value is higher than our interface/link bandwidth cap of 50Mb/s, bursts of packet loss or transient increases in latency could easily lead to sub-optimal TCP throughput for TPP routes.

Why do CPP routes have better throughput than TPP routes? Our initial methodology for measuring loss-rate relied on our low-rate *ping* probes (outlined in § 3.3). While this form of probing can produce a reliable estimate of average loss-rate over a long period of time [66], it doesn't capture the dynamics of packet loss at finer resolutions. We thus modified our probing methodology to incorporate an additional *iperf3* measurement using UDP probes between all CP instances. Each measurement is performed for 5 seconds and packets are sent at a 50Mb/s rate.[8] We measure the number of transmitted and lost packets during each second and also count the number of packets that were delivered out of order at the receiver. We perform these loss-rate measurements for a full week. Based on this new set of measurements, we estimate the overall loss-rate to be $5 * 10^{-3}$ and 10^{-2} for CPP and TPP paths, respectively. Moreover, we experience 0 packet loss in 76% (37%) of our sampling periods for CPP (TPP) routes, indicating that losses for CPP routes tend to be more bursty than for TPP routes. The bursty nature of packet losses for CPP routes could be detrimental to

---

[8] In an ideal setting, we should not experience any packet losses as we are limiting our probing rate at the source.

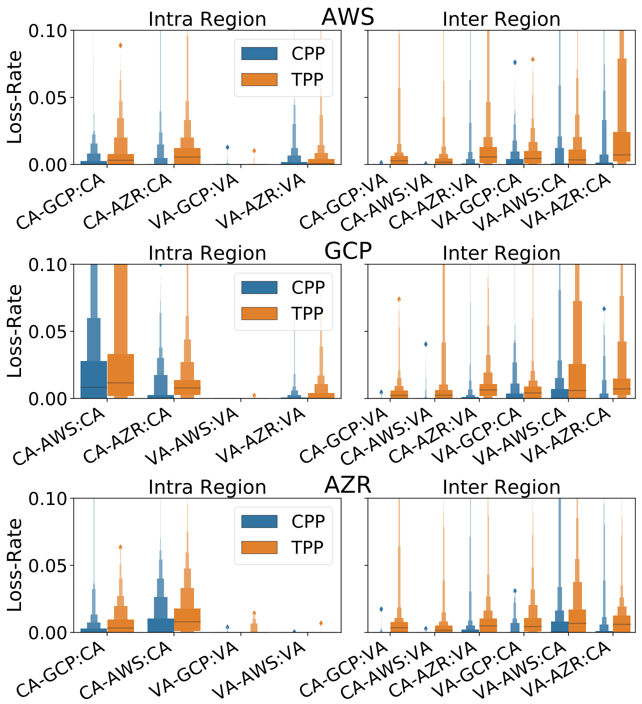

Fig. 5: **Distribution of loss-rate between AWS, GCL, and Azure for intra (left) and inter (right) region paths.**

real-time applications which can *only* tolerate certain levels of loss and should be factored in by the client. The receivers did not observe any out-of-order packets during our measurement period. Figure 5 shows the distribution of loss rate for various paths.

The rows (from top to bottom) correspond to AWS, GCP, and Azure as the source CP, respectively. Intra-region (inter-region) measurements are shown in the left (right) columns, and CPP (TPP) paths are depicted in blue (orange). We observe consistently higher loss-rates for TPP routes compared to their CPP counterparts and lower loss-rates for intra-CP routes in Virginia compared to California. Moreover, paths destined to VMs in the California region show higher loss-rates regardless of where the traffic has been sourced from, with asymmetrically lower loss-rate on the reverse path indicating the presence of congested ingress points for CPs within the California region. We also notice extremely low loss-rates for intra-CP (except Azure) CPP routes between the US east and west coasts and for inter-CP CPP routes between the two coasts for certain CP pairs (*e.g.,* AWS CA to GCP VA or Azure CA to AWS VA).

4.3 Main Findings

Our measurement experiments reveal two interesting findings. First, CPP routes are better than TPP routes in terms of latency as well as throughput. Within a multi-cloud setting, TPPs can serve multiple purposes, including providing connectivity towards CPs from colo facilities that CPs aren't present, lowering inter-cloud traffic costs [8,7], and providing private inter-cloud connectivity over private address spaces. Second, the better performance of CPP routes as compared to their TPP counterparts can be attributed to (a) the CPs' rich (private) connectivity in different regions with other CPs (traffic is by-passing the BEP Internet altogether) and (b) more stable and better provisioned CP (private) backbones.

5 Discussion

CPs are heterogeneous in handling path measurements. Measuring the number of observed AS/organizations (excluding hops utilizing private IP addresses) for inter-cloud, intra-cloud, and cloud-to-LG routes, we observed that of the three CPs, only AWS used multiple ASNs (*i.e.,* ASes 8987, 14618, and 16509) and that there are striking difference between how CPs respond to traceroute probes. In particular, GCP does not expose any of its routers unless the target address is within another GCP region; Azure does not expose its internal routers except for their border routers that are involved in peering with other networks; and AWS relies heavily on private/shared IP addresses for its internal network.

CPs are tightly interconnected with each other in the US. To check the absence of transit ASes along our measured C2C paths more thoroughly, we conducted a more extensive measurement study by launching VM instances within all US regions for our three target CP networks and performing UDP and ICMP *paris-traceroutes* between all VM instances using *scamper*. After annotating the traceroutes as described in § 3.3, in terms of AS/organization-level routes, we only observe organizations corresponding to the three target CPs as well as IXP ASNs for Coresite Any2 and Equinix. All organization-level routes passing through an IXP correspond to paths that are sourced from Azure and are destined to AWS. These measurements further confirm our initial observation regarding the rich connectivity of our three large CPs and their tendency to avoid exchanging traffic through the public Internet.

Taking an Enterprise-to-Cloud (E2C) Perspective. Instead of the C2C perspective shown in Figure 1, we also considered an enterprise-to-cloud (E2C) perspective and report preliminary results for this scenario in Appendix 7.2.

6 Summary

In this paper, we perform a first-of-its-kind measurement study to understand the tradeoffs between three popular multi-cloud connectivity options (CPP vs. TPP vs. BEP). Based on our cloud-centric measurements, we find that CPP routes are better than TPP routes in terms of latency as well as throughput. The better performance of CPPs can be attributed to (a) CPs' rich connectivity in different regions with other CPs (by-passing the BEP Internet altogether) and (b) CPs' stable and well-designed private backbones. In addition, we find that TPP routes exhibit better latency and throughput characteristics when compared with BEP routes. The key reasons include shorter paths and lower loss rates compared to the BEP transits. Although limited in scale, our work highlights the need for more transparency and access to open measurement platforms by all the entities involved in interconnecting enterprises with multiple clouds.

References

1. A First Comparative Characterization of Multi-cloud Connectivity in Today's Internet. https://gitlab.com/onrg/multicloudcmp (2020)
2. Ager, B., Chatzis, N., Feldmann, A., Sarrar, N., Uhlig, S., Willinger, W.: Anatomy of a Large European IXP. In: SIGCOMM. ACM (2012)
3. Alexander, M., Luckie, M., Dhamdhere, A., Huffaker, B., KC, C., Jonathan, S.M.: Pushing the boundaries with bdrmapIT: Mapping router ownership at internet scale. In: Internet Measurement Conference (IMC). ACM (2018)
4. Amazon: AWS Direct Connect. https://aws.amazon.com/directconnect/
5. Amazon: AWS Direct Connect Partners. https://aws.amazon.com/directconnect/partners/
6. Amazon: AWS Transit Gateway. https://aws.amazon.com/transit-gateway/
7. Amazon: AWS Direct Connect Pricing. https://aws.amazon.com/directconnect/pricing/ (2019)
8. Amazon: EC2 Instance Pricing – Amazon Web Services. https://aws.amazon.com/ec2/pricing/on-demand/ (2019)
9. Anwar, R., Niaz, H., Choffnes, D., Cunha, Í., Gill, P., Katz-Bassett, E.: Investigating interdomain routing policies in the wild. In: Internet Measurement Conference (IMC). ACM (2015)
10. Augustin, B., Cuvellier, X., Orgogozo, B., Viger, F., Friedman, T., Latapy, M., Magnien, C., Teixeira, R.: Avoiding Traceroute Anomalies with Paris Traceroute. In: Internet Measurement Conference (IMC). ACM (2006)
11. Augustin, B., Krishnamurthy, B., Willinger, W.: IXPs: Mapped? In: Internet Measurement Conference (IMC). ACM (2009)
12. Ausmees, K., John, A., Toor, S.Z., Hellander, A., Nettelblad, C.: BAMSI: A Multi-cloud Service for Scalable Distributed Filtering of Massive Genome Data. BMC bioinformatics (2018)
13. CAIDA: The Skitter Project. http://www.caida.org/tools/measurement/skitter/ (2007)
14. Calder, M., Fan, X., Hu, Z., Katz-Bassett, E., Heidemann, J., Govindan, R.: Mapping the expansion of google's serving infrastructure. In: Internet Measurement Conference (IMC). ACM (2013)
15. Calder, M., Flavel, A., Katz-Bassett, E., Mahajan, R., Padhye, J.: Analyzing the performance of an anycast CDN. In: Internet Measurement Conference (IMC). ACM (2015)
16. Calder, M., Gao, R., Schröder, M., Stewart, R., Padhye, J., Mahajan, R., Ananthanarayanan, G., Katz-Bassett, E.: Odin: Microsoft's scalable fault-tolerant {CDN} measurement system. In: NSDI. USENIX (2018)
17. Chiu, Y.C., Schlinker, B., Radhakrishnan, A.B., Katz-Bassett, E., Govindan, R.: Are we one hop away from a better internet? In: Internet Measurement Conference (IMC). ACM (2015)
18. CloudHarmony: Cloudharmony, transparency for the cloud. https://cloudharmony.com/
19. CoreSite: The Coresite Open Cloud Exchange. https://www.coresite.com/solutions/cloud-services/open-cloud-exchange
20. Cunha, Í., Marchetta, P., Calder, M., Chiu, Y.C., Machado, B.V., Pescapè, A., Giotsas, V., Madhyastha, H.V., Katz-Bassett, E.: Sibyl: a practical internet route oracle. In: NSDI. USENIX (2016)

21. Demchenko, Y., Van Der Ham, J., Ngo, C., Matselyukh, T., Filiposka, S., de Laat, C., Escalona, E.: Open cloud exchange (OCX): Architecture and functional components. In: International Conference on Cloud Computing Technology and Science. IEEE (2013)
22. Dhamdhere, A., Dovrolis, C.: The Internet is Flat: Modeling the Transition from a Transit Hierarchy to a Peering Mesh. In: CoNEXT. ACM (2010)
23. Durairajan, R., Barford, P., Sommers, J., Willinger, W.: InterTubes: A Study of the US Long-haul Fiber-optic Infrastructure. In: SIGCOMM. ACM (2015)
24. Durairajan, R., Ghosh, S., Tang, X., Barford, P., Eriksson, B.: Internet Atlas: A Geographic Database of the Internet. In: HotPlanet. ACM (2013)
25. Elshazly, H., Souilmi, Y., Tonellato, P.J., Wall, D.P., Abouelhoda, M.: MC-GenomeKey: A Multicloud System for the Detection and Annotation of Genomic Variants. BMC bioinformatics (2017)
26. Facebook: Building express backbone: Facebook's new long-haul network. https://code.fb.com/data-center-engineering/building-express-backbone-facebook-s-new-long-haul-network/ (2017)
27. Gill, P., Arlitt, M., Li, Z., Mahanti, A.: The Flattening Internet Topology: Natural Evolution, Unsightly Barnacles or Contrived Collapse? In: PAM. Springer (2008)
28. Google: GCP Direct Peering. https://cloud.google.com/interconnect/docs/how-to/direct-peering
29. Google: Google supported service providers. https://cloud.google.com/interconnect/docs/concepts/service-providers
30. Haq, O., Raja, M., Dogar, F.R.: Measuring and improving the reliability of wide-area cloud paths. In: WWW. ACM (2017)
31. Hofmann, H., Kafadar, K., Wickham, H.: Letter-value plots: Boxplots for large data. Tech. rep., had.co.nz (2011)
32. Huffaker, B., Keys, K., Fomenkov, M., Claffy, K.: AS-to-Organization Dataset. http://www.caida.org/research/topology/as2org/ (2018)
33. Hung, C.C., Ananthanarayanan, G., Golubchik, L., Yu, M., Zhang, M.: Wide-area analytics with multiple resources. In: EuroSys Conference. ACM (2018)
34. Internet2: One-Way Ping (OWAMP). http://software.internet2.edu/owamp/ (2019)
35. Iyer, A.P., Panda, A., Chowdhury, M., Akella, A., Shenker, S., Stoica, I.: Monarch: gaining command on geo-distributed graph analytics. In: Hot Topics in Cloud Computing (HotCloud). USENIX (2018)
36. Khalidi, Y.: How Microsoft builds its fast and reliable global network. https://azure.microsoft.com/en-us/blog/how-microsoft-builds-its-fast-and-reliable-global-network/ (2017)
37. Klöti, R., Ager, B., Kotronis, V., Nomikos, G., Dimitropoulos, X.: A comparative look into public IXP datasets. SIGCOMM CCR (2016)
38. Knight, S., Nguyen, H.X., Falkner, N., Bowden, R.A., Roughan, M.: The Internet Topology Zoo. In: JSAC. IEEE (2011)
39. Krishna, A., Cowley, S., Singh, S., Kesterson-Townes, L.: Assembling your cloud orchestra: A field guide to multicloud management. https://www.ibm.com/thought-leadership/institute-business-value/report/multicloud
40. Labovitz, C., Iekel-Johnson, S., McPherson, D., Oberheide, J., Jahanian, F.: Internet Inter-domain Traffic. In: SIGCOMM. ACM (2010)
41. Li, A., Yang, X., Kandula, S., Zhang, M.: CloudCmp: Comparing Public Cloud Providers. In: Internet Measurement Conference (IMC). ACM (2010)
42. Luckie, M.: Scamper: a scalable and extensible packet prober for active measurement of the internet. In: Internet Measurement Conference (IMC). ACM (2010)

43. Luckie, M., Dhamdhere, A., Huffaker, B., Clark, D., et al.: bdrmap: Inference of Borders Between IP Networks. In: Internet Measurement Conference (IMC). ACM (2016)
44. Madhyastha, H.V., Isdal, T., Piatek, M., Dixon, C., Anderson, T., Krishnamurthy, A., Venkataramani, A.: iPlane: An Information Plane for Distributed Services. In: OSDI. USENIX (2006)
45. Mao, Z.M., Rexford, J., Wang, J., Katz, R.H.: Towards an Accurate AS-Level Traceroute Tool. In: SIGCOMM. ACM (2003)
46. Marder, A., Smith, J.M.: MAP-IT: Multipass Accurate Passive Inferences from Traceroute. In: Internet Measurement Conference (IMC). ACM (2016)
47. Mathis, M., Semke, J., Mahdavi, J., Ott, T.: The macroscopic behavior of the TCP congestion avoidance algorithm. SIGCOMM CCR (1997)
48. Megaport: Megaport Pricing. https://www.megaport.com/pricing/ (2019)
49. Megaport: Nine Common Scenarios of multi-cloud design. https://knowledgebase.megaport.com/megaport-cloud-router/nine-common-scenarios-for-multicloud-design/ (2019)
50. Microsoft: Azure ExpressRoute. https://azure.microsoft.com/en-us/services/expressroute/
51. Microsoft: Expressroute partners and peering locations. https://docs.microsoft.com/en-us/azure/expressroute/expressroute-locations
52. Motamedi, R., Rejaie, R., Willinger, W.: A survey of techniques for internet topology discovery. Communications Surveys & Tutorials (2014)
53. PacketFabric: Cloud Connectivity. https://www.packetfabric.com/packetcor#pricing (2019)
54. Padhye, J., Firoiu, V., Towsley, D., Kurose, J.: Modeling tcp throughput: A simple model and its empirical validation. SIGCOMM CCR (1998)
55. PCH: Packet Clearing House. https://www.pch.net/ (2019)
56. PeeringDB: PeeringDB. https://www.peeringdb.com/ (2019)
57. Pu, Q., Ananthanarayanan, G., Bodik, P., Kandula, S., Akella, A., Bahl, P., Stoica, I.: Low latency geo-distributed data analytics. SIGCOMM CCR (2015)
58. Pureport: Pricing - Pureport. https://www.pureport.com/pricing/ (2019)
59. RIPE: RIPE RIS (2019)
60. Schlinker, B., Kim, H., Cui, T., Katz-Bassett, E., Madhyastha, H.V., Cunha, I., Quinn, J., Hasan, S., Lapukhov, P., Zeng, H.: Engineering egress with edge fabric: Steering oceans of content to the world. In: SIGCOMM. ACM (2017)
61. Sermpezis, P., Nomikos, G., Dimitropoulos, X.A.: Re-mapping the Internet: Bring the IXPs into Play. CoRR (2017)
62. Shavitt, Y., Shir, E.: DIMES: Let the Internet Measure Itself. In: SIGCOMM CCR. ACM (2005)
63. Sherwood, R., Bender, A., Spring, N.: Discarte: A Disjunctive Internet Cartographer. In: SIGCOMM. ACM (2008)
64. Sherwood, R., Spring, N.: Touring the Internet in a TCP Sidecar. In: SIGCOMM Conference on Measurement. ACM (2006)
65. Spring, N., Mahajan, R., Wetherall, D.: Measuring ISP topologies with Rocketfuel. SIGCOMM (2002)
66. Tariq, M.M.B., Dhamdhere, A., Dovrolis, C., Ammar, M.: Poisson versus periodic path probing (or, does pasta matter?). In: Internet Measurement Conference (IMC). ACM (2005)
67. University of Oregon: Route Views Project. http://www.routeviews.org/

68. Viswanathan, R., Ananthanarayanan, G., Akella, A.: {CLARINET}: Wan-aware optimization for analytics queries. In: Operating Systems Design and Implementation ({OSDI}). USENIX (2016)
69. Vulimiri, A., Curino, C., Godfrey, P.B., Jungblut, T., Karanasos, K., Padhye, J., Varghese, G.: Wanalytics: Geo-distributed analytics for a data intensive world. In: SIGMOD. ACM (2015)
70. Wohlfart, F., Chatzis, N., Dabanoglu, C., Carle, G., Willinger, W.: Leveraging interconnections for performance: the serving infrastructure of a large cdn. In: SIGCOMM. ACM (2018)
71. Yeganeh, B., Durairajan, R., Rejaie, R., Willinger, W.: How cloud traffic goes hiding: A study of amazon's peering fabric. In: Internet Measurement Conference (IMC). ACM (2019)
72. ZDNet: Top cloud providers 2019. https://tinyurl.com/y526vneg
73. Zhang, B., Liu, R., Massey, D., Zhang, L.: Collecting the Internet AS-Level Topology. In: SIGCOMM CCR. ACM (2005)
74. Zhang, H., Chen, K., Bai, W., Han, D., Tian, C., Wang, H., Guan, H., Zhang, M.: Guaranteeing deadlines for inter-data center transfers. Transactions on Networking (TON) (2017)

7 Appendices

7.1 Representation of Results

Distributions in this paper are presented using letter-value plots [31]. Letter-value plots, similar to boxplots, are helpful for summarizing the distribution of data points but offer finer details beyond the quartiles. The median is shown using a dark horizontal line and the $1/2^i$ quantile is encoded using the box width, with the widest boxes surrounding the median representing the quartiles, the 2nd widest boxes corresponding to the octiles, etc. Distributions with low variance centered around a single value appear as a narrow horizontal bar while distributions with diverse values appear as vertical bars.

Throughout this paper we try to present full distributions of latency when it is illustrative. Furthermore, we compare latency characteristics of different paths using the median and variance measures and specifically refrain from relying on minimum latency as it does not capture the stability and dynamics of this measure across each path.

7.2 Preliminary results on E2C perspective

We emulate an enterprise leveraging multi-clouds by connecting a cloud router in the Phoenix, AZ region to a physical server hosted within a colocation facility in Phoenix, AZ.

TPP routes offer better latency than BEP routes. Figure 6a shows the distribution of latency for our measured E2C paths. We observe that TPP routes consistently outperform their BEP counterparts by having a lower baseline of latency and also exhibiting less variation. We observe a median latency of 11ms, 20ms, and 21ms for TPP routes towards GCP, AWS, and Azure VM instances in California, respectively. We also observe symmetric distributions on the reverse path but omit the results for brevity. In the case of our E2C paths, we always observe direct peerings between the upstream provider (*e.g.,*

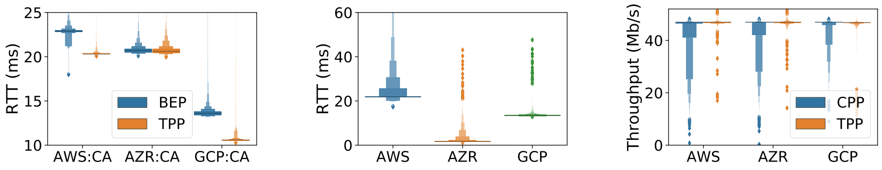

(a) (b) (c)

Fig. 6: **(a) Distribution of latency for E2C paths between our server in AZ and CP instances in California through TPP and BEP routes. Outliers on the Y-axis have been deliberately cut-off to increase the readability of distributions. (b) Distribution of RTT on the inferred peering hop for E2C paths sourced from CP instances in California. (c) Distribution of throughput for E2C paths between our server in AZ and CP instances in California through TPP and BEP routes.**

Cox Communications (AS22773)) and the CP network. Relying on *bdrmapIT* to infer the peering points from the traceroutes associated with our E2C paths, we measure the latency on the peering hop. Figure 6b shows the distribution of the latency for the peering hop for E2C paths originated from the CPs' instances in CA towards our enterprise server in AZ. While the routing policies of GCP and Azure for E2C paths are similar to our observations for C2C paths, Amazon seems to hand-off traffic near the destination which is unlike their hot-potato tendencies for C2C paths. We hypothesize that this change in AWS' policy is to minimize the operational costs via their Transit Gateway service which provide finer control to customers and peering networks over the egress/ingress point of traffic to their network [6]. In addition, observing an equal or lower minimum latency for TPP routes as compared to BEP routes suggests that TPP routes are shorter than BEP paths[9]. We also find (not shown here) that the average loss rate on TPP routes is $6 * 10^{-4}$ which is an order of magnitude lower than the loss rate experienced on BEP routes ($1.6 * 10^{-3}$).

TPP offers consistent throughput for E2C paths. Figure 6c depicts the distribution of throughput for E2C paths between our server in AZ and CP instances in CA via TPP and BEP routes, respectively. While we observe very consistent throughput values near the purchased link capacity for TPP paths, BEP paths exhibit higher variability which is expected given the best effort nature of public Internet paths. Similar to the latency characteristics, we attribute the better throughput of TPP routes to the lower loss rates and shorter fiber paths from the enterprise server to the CPs' instances in CA. Moreover, compared to the CPs' connect locations, the third-party providers are often present in additional, distinct colocation facilities closer to the edge and partially answers the question we posed earlier in § 4.3.

---

[9] In the absence of information regarding the physical fiber paths, we rely on latency as a proxy measure of path length.