On the Impact of Submarine Cable Deployments on Multi-cloud Network Latencies

Joseph Colton
University of Oregon

Ramakrishnan Durairajan
University of Oregon
Link Oregon

Reza Rejaie
University of Oregon

Walter Willinger
NIKSUN, Inc.

***Abstract*—Enterprises are adopting multi-cloud strategies, establishing overlays atop two or more cloud providers (CP) backbones to connect resources and services, even across continents. Simultaneously, there is a significant increase in submarine cable deployments by the CPs to enhance the reliability and performance of their backbone networks. However, enterprises face challenges in understanding how these deployments impact their adoption of multi-cloud strategies. These challenges include the evolving nature of submarine cable deployments, the high cost (and resulting paucity) of data collection efforts on multi-cloud network paths, and a general unawareness of how these deployments impact multi-cloud network path latencies in practice.**

To address this problem, this work presents a third-party measurement study to corroborate the latency trends/changes characteristics of multi-cloud network paths with submarine cable deployments. To this end, we develop a three-step approach: (1) analyze latency characteristics of multi-cloud paths by comparing two sets of measurements obtained across three major CPs, (2) examine the possible root causes of latency trends by leveraging publicly available data sources on submarine cable deployments, and (3) validate our findings with submarine cable operators and providers. Our study reveals several insights into the impact of submarine cable deployments on multi-cloud network paths' latency characteristics, helping enterprises make informed decisions regarding their cloud overlays. To promote reproducibility and extension of our work, we will release the code and datasets to the community.

***Index Terms*—Multi-cloud network paths, submarine cable deployments, latency impacts**

I. Introduction

Two developments in cloud computing have gained significant momentum in recent years. First is the adoption of multi-cloud strategies by modern enterprises for a range of application domains (e.g., genomics [1], [2], HPC [3]–[5], finance [6], [7], science and technology [8]–[10], public policy [11], [12], etc.). These strategies involve connecting resources (e.g., virtual machines or VMs) and services (e.g., analytics) by establishing *overlays* atop federated backbones of individual public cloud providers (CPs) to reap benefits such as competitive pricing, global expansion opportunities, and high reliability. We refer to paths that traverse these federated backbones of individual CPs as multi-cloud network paths.

Second is the significant growth rate in recent years in submarine cable deployments by the CPs (as part of their *underlays* or physical infrastructures). Concretely, major CPs such as Google, Amazon, and Microsoft have been actively investing in and deploying submarine cables. For example, Google has been part of numerous submarine cable projects, including the Curie cable [13] connecting the U.S. and Chile, the Dunant cable [14] connecting the U.S. and France, and the Grace Hopper cable [15] connecting the U.S. to the U.K. and Spain. Similarly, Microsoft has invested in the MAREA cable [16], connecting the U.S. and Spain, and Dunant which is a joint venture with Google. Amazon has also participated in submarine cable projects like the Hawaiki cable [17] connecting the U.S., Australia, and New Zealand. A key reason behind this growth is the need to enhance network capacity, reliability, and latency for delivering cloud services to enterprises worldwide.

However, these two developments notwithstanding, enterprises interested in reaping the benefits of multi-clouds face a formidable problem: *what are the impacts of the ongoing submarine cable deployments on the latency characteristics of multi-cloud network paths?*[1]. Answering this question is fraught with three key challenges. Primary among them is the evolving nature of submarine cable deployments. For example, as many as 60-70 submarine cables became "live" (i.e., ready for service) in the last three years [18]. Second, the lack of community-wide efforts to collect suitable cloud measurements further exacerbates the problem. Third, the operational expenses associated with continuously running tools to gather latency measurements for the myriad of available multi-cloud paths act as a significant deterrent for enterprises. As a result, many aspects of multi-cloud paths are unknown to enterprises, including a basic understanding of how the latency characteristics vary across CPs over time, detailed explanations of the root causes of observed latency features (i.e., evolution of submarine cable deployments), and practical advice on how the performance of multi-cloud deployments affects their selection of multi-cloud strategies.

Our goal in this paper is to address these issues by bringing *awareness of submarine underlays to multi-cloud overlays*. To this end, we build on [19] and present a third-party, cloud-centric measurement study that focuses on corroborating the latency characteristics of multi-cloud network paths with the evolving submarine cable deployment landscape. In particular, we complement a 2019 snapshot of latency measurements of multi-cloud paths considered in [19] with a corresponding

[1] In this work, our focus is on latency but note that the impacts of submarine cables on metrics such as throughput, loss, etc. remain largely unknown and require further studies.

2022 snapshot to perform a comparative analysis of obtained latencies. Both snapshots were gathered through the deployment of VMs across three global CPs, namely AWS, Azure, and GCP, spanning 26 availability regions and encompassing about 448 unique multi-cloud paths (see Table I).

At the core of our study is a three-step approach to assess the latency characteristics of multi-cloud paths. The first step entails analyzing the fundamental characteristics and trends that one can observe in the two snapshots across various geographic granularities, including continents, countries, and cloud availability regions. The second step involves examining the root causes of those trends, leveraging publicly available data sources (i.e., TeleGeography's Submarine Cable Map [18], Global Submarine Cable Systems [20]). The third step is to validate the root causes with submarine cable builders and operators. To this end, we leveraged our partnership with submarine cable operators [21] and performed extensive consistency checking of the root causes.

Cloud Provider (CP)	**Region**
AWS	us-east-1, us-west-2, ap-northeast-1, ap-northeast-2, ap-south-1, ap-southeast-2, eu-central-1, eu-west-2, sa-east-1, ca-central-1
GCP	us-east4, us-west1, asia-northeast1, asia-south1, australia-southeast1, europe-west3, europe-west2, southamerica-east1
Azure	eastus, westus2, japanwest, koreasouth, southindia, australiaeast, uksouth, brazil-south

TABLE I: Cloud providers (CPs) and regions considered in this study.

Using this three-step approach, our study reveals several new insights into the latency characteristics of multi-cloud paths including their overall improvements across snapshots, the impacts of recent submarine cable deployments on latency changes, and the unforeseen consequences that accompany some of the deployments. Concretely, we find that:

- Overall, intercontinental multi-cloud network path latencies have improved with an average decrease of ∼11ms across all three CPs between the two snapshots. The largest drop in latency was ∼221ms between `aws.eu-central-1` (Frankfurt) and `gcp.asia-south1` (Mumbai). In contrast, the largest increase (some 34ms) was observed between `azure.australiaeast` (New South Wales) and `aws.ap-northeast-1` (Tokyo).
- Unsurprisingly, all three of the CPs had both intercontinental multi-cloud paths with latency decreases and paths with latency increases. However, most of the multi-cloud network paths that showed an increase in latency involved AWS regions as either a source or destination, with ∼40% of them showing an increase. At the same time, ∼80% of intercontinental intra-cloud paths of AWS (i.e, AWS-to-AWS paths) showed a latency decrease. In general, the CPs' intra-cloud routes showed the most improvements in latency reductions.
- CPs such as Google (GCP) and Amazon (AWS) are listed as owners or partial owners for many of the submarine cables that were deployed between 2019 and 2022. These cables include: (1) JUPITER [22] (AWS) and Pacific Light Cable Network [23] (GCP) connect East Asia to U.S. West coast where there was a ∼40ms improvement in latency. (2) INDIGO-West [24] (GCP) and INDIGO-Central [24] (GCP) help connect East Australia to India where we saw a latency improvement of ∼24ms. (3) Havfrue/AEC-2 [25] (GCP), Grace Hopper [15] (GCP), and Dunant [14] (GCP) all cross from the U.S. East coast to Europe and resulted in a ∼7ms latency reduction.
- Some cables that were deployed during 2019–2022 also listed companies as partial owners or common partners with the CPs such as Meta [26], [27], Verizon [28], and PCCW Global [29]–[32]. Such cables include: (1) Southern Cross Next [33] (Verizon) runs from the U.S. West coast to Australia and showed a ∼23ms improvement in latency. (2) PEACE Cable [34] (PCCW) which connects Europe to southern Asia resulted in improvements as high as ∼221ms. (3) 2Africa [35] (Meta) which connects Brazil to Africa (where submarine cables are already in place to carry traffic to India) and where measured latency improved by ∼60ms.

To promote reproducibility and extension of this work, the code and the datasets are openly available to the community at [36].

II. Background and Related Work

The adoption of multi-cloud strategies by enterprises reflects a recent paradigm shift in cloud computing and enbles them to establish overlays atop federated underlays of cloud providers (CPs) [19], [37]. This shift has resulted in so-called multi-cloud network paths, which are paths that traverse federations of private network backbones from two or more distinct CPs. Due to this paradigm shift, today's CPs (e.g., Google, Microsoft) have experienced enormous growth in both their ingress (i.e., Internet-facing) and mid-gress (i.e., inter-data center) traffic. To meet these demands, CPs are aggressively expanding their presence at new colocation facilities and enhancing their private global-scale backbones, which includes deploying submarine cables [38], [39]. As a result, we are already seeing a substantial portion of enterprise traffic circumventing the public Internet, presumably to reap the performance benefits that these private paths provide [40] and making multi-cloud network paths appealing to enterprises.

However, the relationship between multi-cloud network paths and the evolving landscape of submarine cable deployments by CPs has remained largely unexplored to date. The study that is closest to this problem is Fanou et al. [41] but its focus is on Africa and on routing within the public Internet. Note that several studies have investigated these two related aspects separately. For example, recent efforts study submarine cable deployments [42] and how IP traffic depends on those deployments [43]. Similarly, many efforts focus on measuring the peering locations, serving infrastructures, and routing strategies of the individual cloud and content providers [44]–[52]. Comparing the path performance of CPs with each other as well as with transit providers has been the focus of prior efforts such as [37], [53]–[57]. While considerable efforts have been devoted to measuring various aspects of individual CP paths, multi-cloud network paths and the performance changes resulting from the evolution of underlays have remained elusive for enterprises. In particular, the latency characteristics, the evolution of these characteristics over time, and the underlying reasons for the evolution have received little attention to date. Indeed, these unknowns highlight the critical importance of a measurement study like ours.

III. Datasets & Methodology

A. *Measurement setting and data collection*

In this study, we target the top three global-scale CPs namely, Amazon Web Services (AWS), Microsoft Azure, and Google Cloud Platform (GCP). We create small VM instances within multiple regions of these CPs resulting in a total of 26 regions (10, 8, and 8 for AWS, Azure, and GCP respectively). Our selection of regions targeted cities with multiple CPs. Some regions are dedicated to government agencies and are not available to the public. Such regions are not considered in this study. Additionally, we identify the datacenter's geolocation for each CP. Although CPs are at times secretive with respect to the location of their datacenters, various sources do point to their exact or approximate location [58]–[63]. In the absence of any online information, we resort to the nearest metro area that the CP advertises.

Table II shows the number of unique inter-continental paths between different availability regions of the three CPs. Note that the cells in bold highlight the intra-cloud portion of multi-cloud network paths. We conducted pairwise latency measurements between all these regions in 10 minute rounds for the duration of 5 days in June of 2019 resulting in about 365k latency measurement samples for the 508 unique inter-continental multi-cloud paths. We refer to the dataset collected in 2019 as snapshot 1 (or *s1* for short). Similarly, we performed pairwise latency measurements between all VMs in 10 minute rounds again in November of 2022, but for 14 days. This resulted in over 1 million latency measurements. We refer to the dataset collected in 2022 as snapshot 2 (or *s2* for short).

		Destination Regions		
		AWS	GCP	Azure
Source Regions	AWS	**70**	61	59
	GCP	61	**50**	50
	Azure	59	50	**48**

TABLE II: Inter-continental source/destination region pairs between the availability regions of CPs. The bold-highlighted cells represent the intra-cloud segments of multi-cloud network paths.

When calculating latencies for each of the source/destination region pairs, we took the median round-trip-time (RTT) value for each given hour, then we averaged those medians to get a picture of what the RTT was over a given period. This allowed for a fair comparison because we could eliminate outliers and we could get numbers that represent an average throughout all hours of the day.

B. *Analysis methodology*

We developed a three-step approach to elucidate the latency characteristics of multi-cloud network paths by utilizing two snapshots. In the first step (see Section IV-A), we analyzed the essential characteristics and trends observed in the two snapshots. This includes determining the percentage of paths that experience latency increases, decreases, or remain unchanged. Additionally, we compared the latency characteristics of inter-continental multi-cloud paths between snapshots.

In the second step (see Section IV-B), our objective was to examine the reasons behind the latency changes observed between the two snapshots. We achieve this by utilizing two publicly available data sources that provide information on submarine cable deployments: TeleGeography's Submarine Cable Map [18] and Global Submarine Cable Systems [20]. We leveraged these sources to provide supporting evidence and explanations for the observed latency changes between the snapshots.

For the third step (see Section IV-C), we leveraged our partnership with submarine cable operators and providers [21] to validate the plausible root causes identified in step two.

IV. Results

A. *Two snapshots of multi-cloud network latencies*

Characteristics of inter-continental multi-cloud paths with latency reductions. We start by characterizing all the inter-continental paths. Figure 1 presents the changes in latency of individual inter-continental paths in a scatter plot where each dot indicates a path, and where the x- and y- coordinates show its corresponding latency in snapshots s1 and s2, respectively. Intra-cloud paths are shown with red crosses while inter-cloud paths are shown with blue dots. This figure reveals that the latency for many paths (of both types) has decreased, but there are also paths that experienced an increase in latency over time. This figure has labels for a number of data points associated with specific paths that, with the exception of one set of points, exhibit a relatively larger reduction in latency between the two snapshots. These paths are the focus of our case studies in the next section.

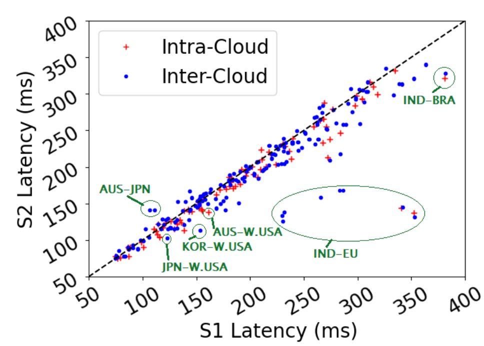

Fig. 1: Scatterplot showing a comparison of latency values for snapshots s1 to s2 for inter-continental paths between regions.

To get a better sense of changes in latency of inter-continental paths between the two snapshots, Figure 2 presents the CDF of the normalized change in latency between s1 and s2, i.e., decreasing or increasing changes (negative or positive values) that are normalized by the latency in s1. Figure 2 further divides these inter-continental paths into the following four categories: one set of intra-cloud provider paths, and three sets of inter-cloud paths between different pairs of the three CPs. This figure illustrates the following noteworthy points: (i) About two-thirds of the paths in each category exhibit a decrease, and the remaining one-third shows an increase in latency between two snapshots. However, the amount of normalized decrease is much larger than that of normalized increase. (ii) There are no significant differences in the distribution of normalized changes across the four categories; i.e., intra- and the different inter-CP paths generally exhibit a similar distribution of normalized changes in latency. (iii) Roughly 1 in 5 paths in each group exhibits more than a 10% drop in latency between the two snapshots.

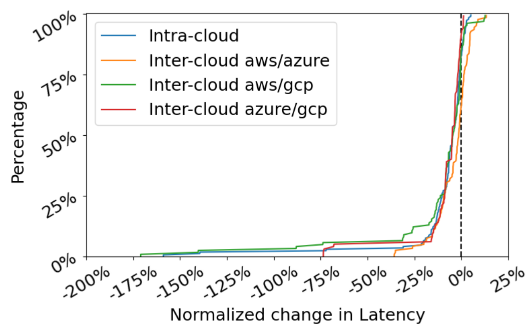

Fig. 2: CDFs showing normalized changes in latency of inter-continental paths between snapshots s1 and s2, for all intra-cloud paths and the three different sets of inter-cloud paths.

Percentage of inter-continental multi-cloud paths with latency reductions. As expected, a significant portion of the intercontinental intra-cloud paths of the three CPs exhibit latency improvements: 80% for AWS, 68% for GCP, and 83% for Azure. Among the inter-cloud portion of the inter-continental multi-cloud paths, over 74% of the paths from Azure to GCP (and vice versa) have seen significant latency improvements. Overall, we observe latency reductions for at least 55% of the inter-continental multi-cloud paths between the two snapshots.

Additional observations. There was an average decrease in the median latency values for all inter-continental paths that directly connect regions in different continents, including both intra- and inter-CP paths. Note that the latency improvements are more pronounced for intra-cloud paths of GCP compared to the other two CPs. Additionally, GCP is the only provider that provided at least an average of 10ms latency savings for enterprises in an intercontinental multi-cloud setting over the two-snapshot period.

Main takeaways. We observed significant latency improvements in many inter-continental paths, leading to a more consistent and normalized latency between global locations. While the choice of CPs played a substantial role in reducing inter-continental latency in 2019, by late 2022, the impact of this choice had diminished considerably. However, regions such as Brazil and Australia still experience latency fluctuations and would benefit from additional submarine cable deployments to further improve their connectivity and, in turn, achieve better and more stable latencies.

B. Case Studies: Observed Latency Changes and their Potential Root Causes

Having described the main trends and noteworthy differences in the measured latencies between the two different snapshots, we next present several case studies selected from our comparison. Our goal is to explain the reasons behind the latency differences we observed between the two snapshots. To this end, we are motivated by prior efforts reported in Fanou et al. [41] and use publicly-available data sources including TeleGeography's Submarine Cable Map [18] and Global Submarine Cable Systems [20]. In particular, we (i) identify a number of submarine cables that became live (i.e., ready for service) between 2019 (when s1 was collected) and 2022 (when s2 was collected); (ii) confirm the providers, partnerships, and locations for each one of those cables and consider their measured pairwise path differences between snapshots; and (iii) validate the results with operators and providers at the SubOptic foundation. Figure 3 shows several such instances identified from [18].

1) India to Europe: We observed the largest reduction in multi-cloud path latencies between India and Europe. Figure 4a shows that the average difference over all CP paths that connect India and the UK between snapshots s1 and s2 is ∼76ms. Additionally, the reduced latency variation in s1, as compared to s2, reflects more consistent performance. This improvement can be attributed to latency decreases by GCP and Azure as they are trying to catch up with AWS. We also examined individual regions. Figures 4b and 4c show that connections from both Frankfurt and London to Mumbai improved on average by ∼200ms. The highest latency connections in s1 were larger than 230-350ms depending on the CPs involved. By the time snapshot s2 was collected, all of the CP paths had similar latency numbers (∼150ms). This indicates that the CPs are all likely now using the same or similar quality routes from India to Europe.

There are a number of possible explanations for this improvement in latency. While the typical strategy GCP employs for routing packets is cold-potato routing, AWS uses hot-potato routing and Azure employs mixed routing (i.e., hot-potato routing for GCP and cold-potato routing for AWS [37]). This means that if GCP does not have a low-latency connection, other CPs will also be affected. The Peace Cable [34], labeled ❸ on Figure 3, was deployed between France and Karachi, Pakistan, allowing more communication between Europe, down through the Middle East, Northeast Africa, and South Asia. Karachi, Pakistan, and Mumbai, India, are already connected through the AAE-1 [64] submarine cable that went live in 2017. Both cables are listed on PCCW Global's Infrastructure Map [29] showing the connection, and

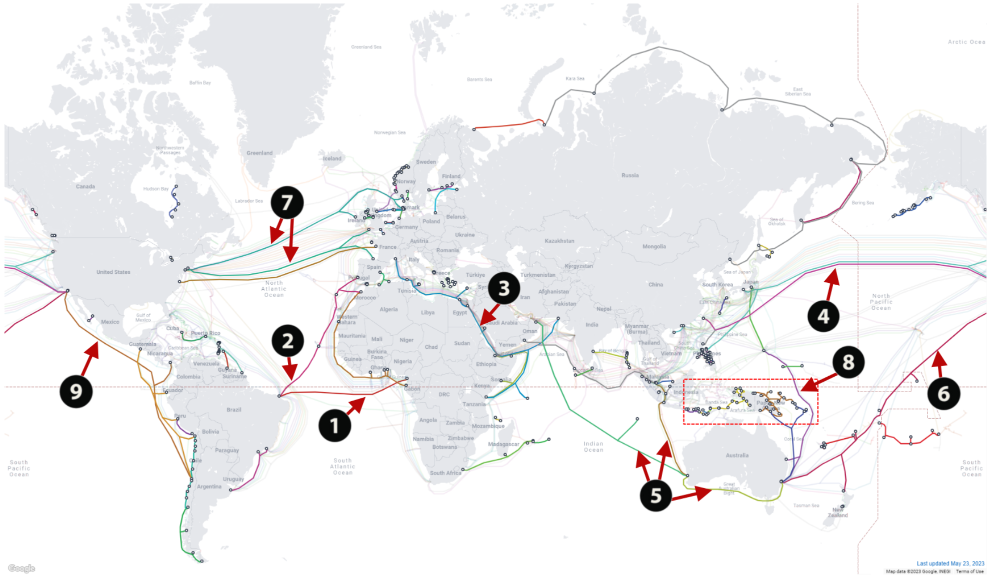

Fig. 3: Annotated submarine fiber-optic map with cables ready-for-service between 2019 and 2022. Source of basemap: submarinecablemap.com [18].

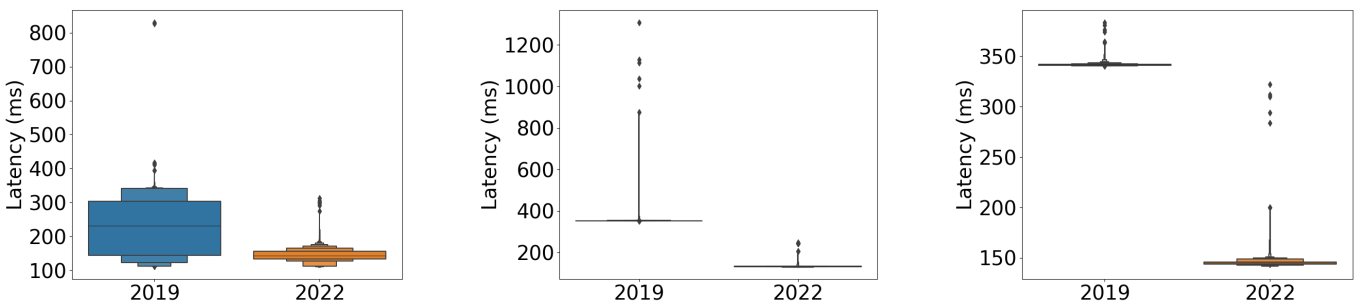

(a) India to UK (b) Frankfurt (AWS) to Mumbai (GCP) (c) London (AWS) to Mumbai (GCP)

Fig. 4: India to Europe

PCCW [30]–[32] is a common partner with GCP and Azure. Additionally, Lightstorm [65], which lists AWS, Google, and Azure as partners, deployed thousands of miles of fiber-optic cable across India connecting the major cities [66]. Because new submarine cable deployments were established that were operated by GCP partners, GCP was able to achieve a lower latency, and other CPs were also able to take advantage thanks to the GCP cold-potato routing.

2) East Asia to the Western United States: The latencies between East Asia and the West Coast of the U.S. improved from s1 to s2 as seen in Figures 5a and 5b. The latency values between the snapshots dropped by ∼20ms for Japan and ∼40ms for South Korea. Now both have similar latency values.

From Figure 3, we can see that South Korea typically connects to the U.S. through submarine cables first to Japan, then across the Pacific. Two submarine cables, JUPITER, which is partially owned by AWS and Meta [22], and connects the U.S. to Japan, and Pacific Light Cable Network, which is partially owned by Google and Meta [23], which connects the U.S. to Taiwan, which is also connected to South Korea and

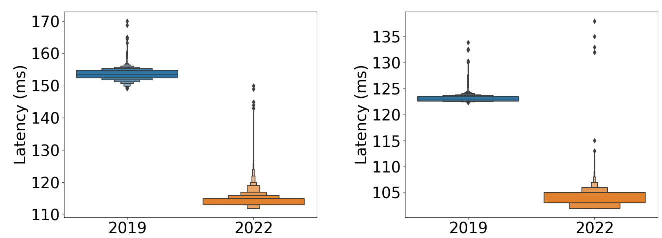

(a) South Korea (Azure) to Oregon (AWS) (b) Japan (Azure) to Oregon (AWS)

Fig. 5: West United States to East Asia (Japan/South Korea)

Japan, were deployed across the Pacific during the period the two snapshots were collected. They are both identified by the label ❹ on Figure 3. These cables from the West Coast of the U.S. to Asia are likely causes of the decrease in latency measurements between s1 and s2.

3) Brazil to India: The latency measurements from Brazil to India saw a decrease of ∼23ms on the inter-cloud paths from GCP to Azure and ∼60ms on a GCP intra-cloud paths. In both cases, the total latency numbers are still high. The largest drop was from ∼380ms to ∼320ms, and all of the current routes are still above 300ms, which indicates there is still room for improvement. These improvements in latency, however, can be seen in Figure 6a and Figure 6b, respectively.

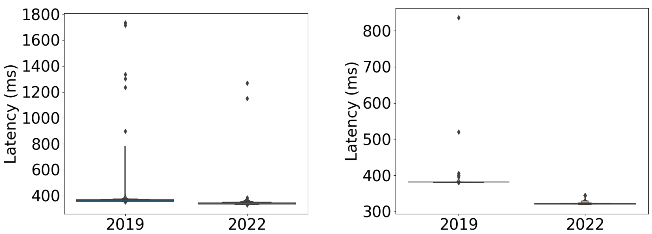

(a) Brazil (GCP) to India (Azure) (b) Brazil (GCP) to India (GCP)

Fig. 6: Brazil to India

Reasoning about the improvement in latency between Brazil and India is not straightforward because there is no direct route between the two countries. However, there are multiple submarine cables that could contribute to the improvement in latency numbers. Two submarine cables were established leaving Brazil. Figure 3 shows the South Atlantic Inter Link (SAIL) cable which connects to Cameroon, where multiple older cables connect to South Africa. From there, the 2Africa cable [35], which is partially owned by Meta [67] ❶ connects around the south east side of the continent. Also, the EllaLink cable [68] ❷ crosses to Africa where there were current submarine cables already in service. Additionally, the Curie cable ❾ connected South America to North America where there are cables running to India.

4) Eastern United States to Europe: Even though the U.S. and Europe have a history of being connected and have good latency, we see an improvement between s1 and s2 of ∼7ms from the East Coast of the U.S. to both London, United Kingdom, and Frankfurt, Germany. The connections from the Eastern U.S. to Frankfurt are depicted in Figure 7a and the ones from the Eastern U.S. to London are shown in Figure 7b.

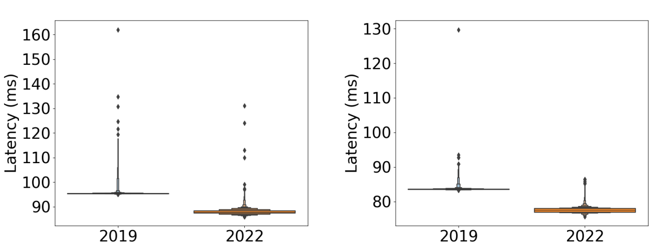

(a) East US (Azure) to Frankfurt (GCP) (b) East US (Azure) to London (AWS)

Fig. 7: East United States to Europe

This improvement is likely because of submarine cables installed across the Atlantic. Specifically, there are three cables: the Havfrue/AEC-2 [25], Grace Hopper [15], and Dunant [14], all of which likely contributed to the latency reduction and are all partially or completely owned by Google. These are all shown in Figure 3 as ❼.

5) Western United States to Australia: Examining the latency between the West Coast of the U.S. and Australia, Figures 8a and 8b shows an improvement of about 20ms.

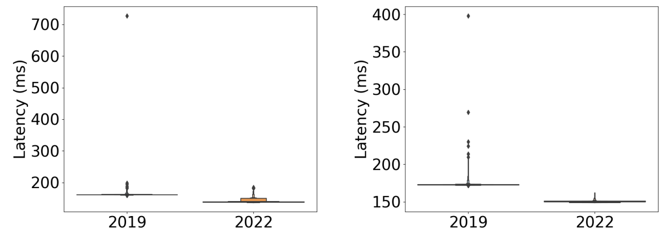

(a) Australia (GCP) to Oregon (GCP) (b) Oregon (GCP) to Australia (Azure)

Fig. 8: West United States to Australia

The 20ms improvement between Oregon (Western U.S.) and Australia could be explained by the deployment of the Southern Cross NEXT submarine cable [33], listed as ❻ in Figure 3. This cable connects the West Coast of the U.S. to New South Wales, Australia, as well as some other islands along the way. Verizon [28] was a partner for this cable and is also a common partner with the major cloud providers.

6) Australia to India: As shown in Figures 9a and 9b, we see a decrease in latency of ∼10ms from Mumbai (GCP) to Australia (Azure) and ∼24ms from Australia (Azure) to India (AWS).

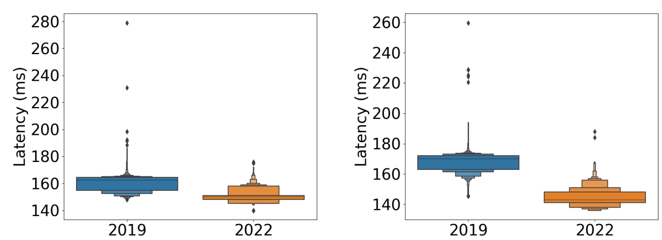

(a) India (GCP) to Australia (Azure) (b) Australia (Azure) to India (AWS)

Fig. 9: Australia to India

To explain the latency changes on the connection between Mumbai, India (GCP), and East Australia (Azure), it is necessary to look at the submarine cables installed along the west side of Australia. The Oman Australia Cable (OAC) [69] and INDIGO Central [70] submarine cables connect Australia to Oman. The INDIGO-West [24] completes a path to Singapore. These cables are all identified in Figure 3 as label ❺. Google is listed as a partial owner for both the INDIGO cables, and all three are listed as cables that SUBCO, the owner of OAC operates. From Oman, the Asia Africa Europe-1 (AAE-1) cable [64], which is partly owned by common partner PCCW [30]–[32], connects directly to Mumbai, India. From Singapore, there are multiple submarine cables connecting to Mumbai. These could contribute to the ∼10ms drop in latency numbers and could be used to move traffic from Eastern Australia along the south side and across to India.

7) Brazil to Japan: The latency measurements exhibit inconsistencies across countries when examined across various regions. In the case of communication between Brazil and

Japan, the latency either increased or decreased depending on the CPs involved. Figure 10a shows a ∼15ms increase in latency from AWS to Azure, while Figure 10b shows a ∼5ms decrease in latency from GCP to AWS.

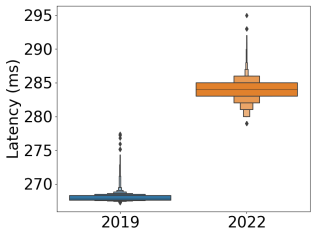

(a) (AWS) to Japan (Azure)

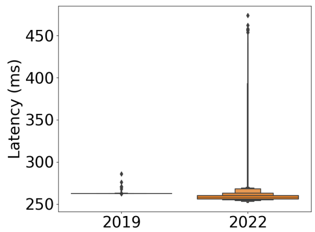

(b) Brazil (GCP) to Japan (AWS)

Fig. 10: Brazil to Japan

To gain a better understanding of the latency measurements and the inconsistencies between Brazil and Japan, we mention the following relevant details. The Curie cable [13], listed as ❾ on Figure 3, runs from California to Valparaíso, Chile, and is owned and operated by Google. From Chile, PCCW Global [29], a Google partner, operates overland fiber cables that cross Argentina, then drop into submarine cables to reach Sao Paulo, Brazil. From California, there are multiple cables that cross to Japan, including the newly deployed Pacific Light Cable Network (PLCN) [23], which is jointly owned by Google and Meta.

Recall that AWS typically uses hot potato routing, Azure employs hot-potato routing for GCP and cold-potato routing for AWS, while GCP typically employs cold potato routing [37]. Figure 10a shows a ∼15ms increase in latency from AWS to Azure, which could be attributed to hot potato routing. This is a clear case of unintended consequences of submarine cable deployment and CPs' hot potato routing. In contrast, Figure 10b shows a ∼5ms decrease in latency from GCP to AWS, which is possibly the result of Google's cold potato routing. Concretely, it appears that Google routes traffic through their own cables, including the Curie cable, rather than letting the traffic find its own way to Japan.

8) Australia to East Asia: The latency measurements in Figure 11b between aws.ap-northeast-1 (Tokyo) and gcp.australia-southeast1 (Sydney) show an increase of ∼30ms in latency. This increase was usually present in s2, but at times, latency measurements were actually at or below those from s1; that is, latency is typically higher, but sometimes actually lower in 2022 than in 2019.

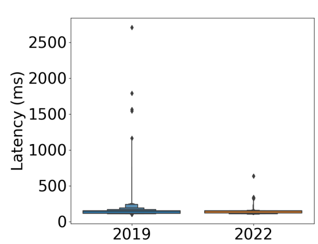

(a) Asia to Australia

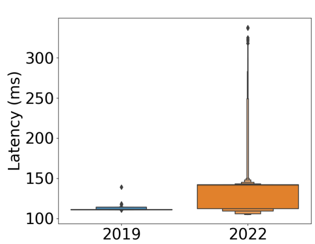

(b) Tokyo (AWS) to Sydney (GCP)

Fig. 11: Asia to Australia

The increase may be another case of an unintended consequence of submarine cables and hot potato routing followed by one of the CPs (i.e., AWS in this case). One possible explanation is the deployment of submarine cables in the Papua New Guinea, Indonesia and Melanesian area (❽ in Figure 3). The following cables were deployed between 2019 and 2022: Coral Sea Cable System (CS2) [71], Kumul Domestic Submarine Cable System [72], Palapa Ring East [73], and three by Moratelindo Broadband Company [74] which include Kupang-Alor Cable Systems [75], Sape-Labuan Bajo-Ende-Kupang Cable System [76], and Denpasar-Waingapu Cable System [77]. These cables connected the islands of Indonesia and Papua New Guinea to the Australian mainland and created additional traffic that could account for the higher latency numbers leaving Australia for Tokyo.

Main takeaways. Our observations over the 3.5-year period between s1 and s2 revealed significant latency reductions in many cases. For one, some CPs held strategic advantages over their competitors; however, the deployment of submarine cables has led to latency decreases, resulting in a more level playing field among the major CPs. This is particularly evident in communication between Europe and India. Second, latency has improved as CPs employing cold potato routing have expanded their submarine cable infrastructure. Notable improvements can be seen in regions where GCP has deployed cables, such as East Asia to the Western United States, Europe to India, and Australia to India. Finally, regions with less stable latency, such as Australia to Japan, Brazil to Africa, and Australia to Brazil, would benefit from additional submarine cable deployments.

C. Validation

To validate the likely root causes found in step two, we relied on our partnership with submarine cable operators and providers [21]. In particular, for the identified root cause for each of the considered case studies, we asked the operators the following two questions: (**q1**) Is our identified root cause accurate? (**q2**) If confidentiality is a concern, are the corroborations reasonable (i.e., provide a plausible explanation)? While we consider answers to (q1) to be a strong indicator of how sound our methodology is, we believe that answers to (q2)—which is mindful of operators privacy concerns—are still useful for the purpose of our study.

	1	2	3	4	5	6	7	8
q1					y	y		
q2	y	y	y	y				

TABLE III: Validation: Expert-provided answers for the considered case studies.

Table III presents the expert-provided answers to the two questions. The experts verified two cases (i.e., 5 and 6) and were cautiously optimistic about our reasoning for the first four cases. They could not comment on the last two cases (i.e., 7 and 8 below) due to missing relevant routing information for the involved CPs. Overall, validating our results with SubOptic experts significantly enhances the confidence of our findings.

V. Summary & Future Work

The cloud computing landscape has witnessed two significant developments: the adoption of multi-cloud strategies by enterprises and the rapid growth of submarine cable deployments by major CPs. While multi-cloud strategies offer several benefits, we posit that understanding the impact of submarine cable deployments on the performance of federated network paths would amplify the benefits of these strategies further. In light of this, this work proposes a three-step approach to assess the latency characteristics and trends of multi-cloud paths observed in the two snapshots; examines the root causes of those trends using publicly available data sources on submarine cable deployments; and validates the root causes with submarine cable operators and providers. Our study reveals several new insights into the latency characteristics of multi-cloud paths, including their improvements across snapshots and the impacts of recent submarine cable deployments on latency changes.

While this study seeks to shed light on the latency changes observed on multi-cloud network paths due to the evolving submarine cable landscape, efforts to investigate these federated private network paths from the perspective of different traffic and application workloads, resiliency issues, egress transit costs, and metrics remain as "unknowns" to enterprises today. Key unknowns include a comprehensive assessment of multi-cloud performance characteristics, insights into how these characteristics differ across different CPs, the identification of performance invariants in federated cloud backbones, shared infrastructure risks of multi-cloud backbones, dynamic variations in the performance of multi-cloud paths underlays, and the implications of these changes for enterprises and applications operating in a multi-cloud environment. We intend to focus on such unknowns as part of future work.

Acknowledgements

We thank the anonymous reviewers as well as the submarine cable operators and providers for their constructive feedback. This work is supported by the Internet Society (ISOC) Foundation and the National Science Foundation (NSF) through CNS-2145813. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of ISOC, or NSF.

References

[1] H. Elshazly, Y. Souilmi, P. J. Tonellato, D. P. Wall, and M. Abouelhoda, "MC-GenomeKey: A Multicloud System for the Detection and Annotation of Genomic Variants," *BMC bioinformatics*, vol. 18, no. 1, p. 49, 2017.

[2] K. Ausmees, A. John, S. Z. Toor, A. Hellander, and C. Nettelblad, "BAMSI: A Multi-cloud Service for Scalable Distributed Filtering of Massive Genome Data," *BMC bioinformatics*, vol. 19, no. 1, p. 240, 2018.

[3] I. Sfiligoi, D. Schultz, B. Riedel, F. Wuerthwein, S. Barnet, and V. Brik, "Demonstrating a pre-exascale, cost-effective multi-cloud environment for scientific computing," *arXiv preprint arXiv:2004.09492*, 2020.

[4] "Using 50k GPUs across multiple Clouds for IceCube Science," https://www.linkedin.com/pulse/using-50k-gpus-across-multiple-clouds-icecube-science-igor-sfiligoi.

[5] "SDSC and IceCube Center Conduct GPU Cloudburst Experiment," https://opensciencegrid.org/news/2019/11/22/gpu-cloudburst.html.

[6] "Multi-Cloud Strategy Drives Financial Services Digital Transformation," https://www.intel.com/content/www/us/en/financial-services-it/solutions/multi-cloud-for-financial-industry-transformation.html.

[7] "Gartner Says Finance Is Moving to the Cloud Much Faster Than Expected," https://www.gartner.com/en/newsroom/press-releases/2017-09-13-gartner-says-finance-is-moving-to-the-cloud-much-faster-than-expected.

[8] "Cloudbank," https://www.cloudbank.org/.

[9] N. S. Foundation and I. to Explore Cloud Computing to Accelerate Science Frontiers, https://www.internet2.edu/news/detail/16908/.

[10] "Nsf fabric testbed," https://whatisfabric.net/.

[11] "Reps. Eshoo, Gonzalez, Sherrill Introduce Bipartisan, Bicameral Legislation to Develop National AI Research Resource," https://eshoo.house.gov/media/press-releases/reps-eshoo-gonzalez-sherrill-introduce-bipartisan-bicameral-legislation-develop.

[12] "Universities and tech giants back national cloud computing project," https://www.nytimes.com/2020/06/30/technology/national-cloud-computing-project.html.

[13] "Curie Cable," https://www.submarinecablemap.com/submarine-cable/curie, 2020, [Online; accessed May 25, 2023].

[14] "Dunant Cable," https://www.submarinecablemap.com/submarine-cable/dunant, 2021, [Online; accessed May 25, 2023].

[15] "Grace Hopper Cable," https://www.submarinecablemap.com/submarine-cable/grace-hopper, 2022, [Online; accessed May 25, 2023].

[16] "MAREA Cable," https://www.submarinecablemap.com/submarine-cable/marea, 2018, [Online; accessed May 25, 2023].

[17] "Hawaiki Cable," https://www.bw-digital.com/, 2018, [Online; accessed May 25, 2023].

[18] "Submarine Cable Map from TeleGeography," https://www.submarinecablemap.com/, 2023.

[19] B. Yeganeh, R. Durairajan, R. Rejaie, and W. Willinger, "A Case for Performance- and Cost-aware Multi-cloud Overlays," in *In Proceedings of IEEE International Conference on Cloud Computing*, 2023.

[20] "Global Submarine Cable Systems," https://www.submarinenetworks.com/en/systems, 2023.

[21] "SubOptic Foundation," https://suboptic.org/page/our-members.

[22] "JUPITER Cable," https://www.submarinecablemap.com/submarine-cable/jupiter, 2022, [Online; accessed May 25, 2023].

[23] "Pacific Light Cable Network Cable," https://www.submarinecablemap.com/submarine-cable/pacific-light-cable-network-plcn, 2022, [Online; accessed May 25, 2023].

[24] "INDIGO-West Cable," https://www.submarinecablemap.com/submarine-cable/indigo-west, 2019, [Online; accessed May 25, 2023].

[25] "Havfrue/AEC-2 Cable," https://www.submarinecablemap.com/submarine-cable/havfrueaec-2, 2020, [Online; accessed May 25, 2023].

[26] "Meta Selects AWS as Key, Long-Term Strategic Cloud Provider," https://press.aboutamazon.com/2021/12/meta-selects-aws-as-key-long-term-strategic-cloud-provider, 2021, [Online; accessed May 25, 2023].

[27] "Meta selects Azure as strategic cloud provider to advance AI innovation and deepen PyTorch collaboration," https://azure.microsoft.com/en-us/blog/meta-selects-azure-as-strategic-cloud-provider-to-advance-ai-innovation-and-deepen-pytorch-collaboration/, 2022, [Online; accessed May 25, 2023].

[28] J. Burt, "Google and Verizon: The Latest Cloud Provider and Mobile Operator Pairing," https://www.enterprisenetworkingplanet.com/news/google-verizon-collaboration/, 2022, [Online; accessed May 25, 2023].

[29] "PCCW Global Infrastructure Map," https://www.pccwglobalinc.com/wp-content/uploads/2022/05/cc-network-map-2022.pdf, 2023, [Online; accessed May 25, 2023].

[30] "Console Connect by PCCW Global: Connecting multinationals, business partners, clouds, and apps all over the world," https://cloud.google.com/customers/pccw-global, 2023, [Online; accessed May 25, 2023].

[31] "Microsoft Hong Kong and Lenovo PCCW Solutions partner to speed up cloud innovation and adoption in Hong Kong," https://news.microsoft.com/en-hk/2022/12/01/microsoft-hong-kong-and-lenovo-pccw-solutions-partner-to-speed-up-cloud-innovation-and-adoption-in-hong-kong/, 2022, [Online; accessed May 25, 2023].

[32] "PCCW Global, NTT, Amazon Join Forces on Trans-Pacific Network," https://www.channelfutures.com/emea/pccw-global-ntt-amazon-join-forces-on-trans-pacific-network, 2023, [Online; accessed May 25, 2023].

[33] "Southern Cross NEXT Cable," https://www.southerncrosscables.com/, 2022, [Online; accessed May 25, 2023].

[34] "PEACE Cable," http://www.peacecable.net/, 2022, [Online; accessed May 25, 2023].

[35] "2 Africa Cable," https://www.2africacable.net/, 2023, [Online; accessed May 25, 2023].

[36] "Link to the source code and datasets." https://gitlab.com/onrg/multicloud-subsea.

[37] B. Yeganeh, R. Durairajan, R. Rejaie, and W. Willinger, "A First Comparative Characterization of Multi-cloud Connectivity in Today's Internet," in *International Conference on Passive and Active Network Measurement*. Springer, 2020, pp. 193–210.

[38] S. Mandal, "Lessons learned from b4, google's sdn wan," *presentation slides, USENIX ATC*, vol. 15, 2015.

[39] Y. Khalidi, "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.

[40] T. Arnold, J. He, W. Jiang, M. Calder, I. Cunha, V. Giotsas, and E. Katz-Bassett, "Cloud Provider Connectivity in the Flat Internet," in *Proceedings of the ACM Internet Measurement Conference*, 2020, pp. 230–246.

[41] R. Fanou, B. Huffaker, R. Mok, and K. C. Claffy, "Unintended consequences: Effects of submarine cable deployment on internet routing," in *Passive and Active Measurement: 21st International Conference, PAM 2020, Eugene, Oregon, USA, March 30–31, 2020, Proceedings 21*. Springer, 2020, pp. 211–227.

[42] S. Liu, Z. S. Bischof, I. Madan, P. K. Chan, and F. E. Bustamante, "Out of sight, not out of mind: A user-view on the criticality of the submarine cable network," in *Proceedings of the ACM Internet Measurement Conference*, 2020, pp. 194–200.

[43] A. Ramanathan and S. A. Jyothi, "Nautilus: A framework for cross-layer cartography of submarine cables and ip links," *arXiv preprint arXiv:2302.14201*, 2023.

[44] R. Anwar, H. Niaz, D. Choffnes, Í. Cunha, P. Gill, and E. Katz-Bassett, "Investigating Interdomain Routing Policies in the Wild," in *proceedings of ACM Internet Measurement Conference*, 2015.

[45] M. Calder, A. Flavel, E. Katz-Bassett, R. Mahajan, and J. Padhye, "Analyzing the Performance of an Anycast CDN," in *proceedings of ACM Internet Measurement Conference*, 2015.

[46] Y.-C. Chiu, B. Schlinker, A. B. Radhakrishnan, E. Katz-Bassett, and R. Govindan, "Are we one hop away from a better internet?" in *proceedings of ACM Internet Measurement Conference*, 2015.

[47] Í. Cunha, P. Marchetta, M. Calder, Y.-C. Chiu, B. V. Machado, A. Pescapè, V. Giotsas, H. V. Madhyastha, and E. Katz-Bassett, "Sibyl: A practical internet route oracle," in *proceedings of USENIX Symposium on Networked Systems Design and Implementation*, 2016.

[48] 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 *ACM SIGCOMM*, 2017.

[49] 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 *proceedings of USENIX Symposium on Networked Systems Design and Implementation*, 2018.

[50] F. Wohlfart, N. Chatzis, C. Dabanoglu, G. Carle, and W. Willinger, "Leveraging Interconnections for Performance: The Serving Infrastructure of a Large CDN," in *proceedings of ACM SIGCOMM*, 2018.

[51] T. Arnold, E. Gürmeriçliler, G. Essig, A. Gupta, M. Calder, V. Giotsas, and E. Katz-Bassett, "(How Much) Does a Private WAN improve Cloud Performance?" in *IEEE INFOCOM 2020-IEEE Conference on Computer Communications*. IEEE, 2020, pp. 79–88.

[52] B. Yeganeh, R. Durairajan, R. Rejaie, and W. Willinger, "How cloud traffic goes hiding: A study of amazon's peering fabric," in *Proceedings of the Internet Measurement Conference*, 2019, pp. 202–216.

[53] A. Li, X. Yang, S. Kandula, and M. Zhang, "CloudCmp: Comparing Public Cloud Providers," in *proceedings of ACM Internet Measurement Conference*, 2010.

[54] O. Haq, M. Raja, and F. R. Dogar, "Measuring and Improving the Reliability of Wide-area Cloud Paths," in *proceedings of WWW*, 2017.

[55] H. Zhang, K. Chen, W. Bai, D. Han, C. Tian, H. Wang, H. Guan, and M. Zhang, "Guaranteeing deadlines for inter-data center transfers," *Transactions on Networking*, 2017.

[56] M. Calder, X. Fan, Z. Hu, E. Katz-Bassett, J. Heidemann, and R. Govindan, "Mapping the Expansion of Google's Serving Infrastructure," in *proceedings of ACM Internet Measurement Conference*, 2013.

[57] "Cloudharmony, transparency for the cloud," https://cloudharmony.com/.

[58] Build Azure, "Microsoft Azure Region Map," https://map.buildazure.com/, 2019.

[59] Google, "Data center locations," https://www.google.com/about/datacenters/inside/locations/index.html, 2019.

[60] R. Miller, "Regional Data Center Clusters Power Amazon's Cloud," https://datacenterfrontier.com/regional-data-center-clusters-power-amazons-cloud/, 2015.

[61] M. Williams, "Amazon's central Ohio data centers now open," http://www.dispatch.com/content/stories/business/2016/10/18/amazon-data-centers-in-central-ohio-now-open.html, 2016.

[62] G. Plaven, "Amazon keeps building data centers in Umatilla, Morrow counties," http://www.eastoregonian.com/eo/local-news/20170317/amazon-keeps-building-data-centers-in-umatilla-morrow-counties, 2017.

[63] I. Burrington, "Why Amazon's Data Centers Are Hidden in Spy Country," https://www.theatlantic.com/technology/archive/2016/01/amazon-web-services-data-center/423147/, 2016.

[64] "Asia Africa Europe-1 (AAE-1) Cable," https://www.aaeone.com/aaeportal/, 2017, [Online; accessed May 25, 2023].

[65] "Lightstorm," https://www.lightstorm.net/, 2023, [Online; accessed May 25, 2023].

[66] D. Goovaerts, "Lightstorm runs fresh fiber for AWS, Google data centers in India," https://www.fiercetelecom.com/telecom/lightstorm-hooks-aws-google-fresh-fiber-india, 2022, [Online; accessed May 24, 2023].

[67] "2Africa Subsea Cable Makes First Landing in Genoa, Italy," https://about.fb.com/news/2022/04/2africa-subsea-cable-makes-first-landing-in-genoa-italy/, 2023, [Online; accessed May 25, 2023].

[68] "EllaLink Cable," https://www.ella.link/, 2021, [Online; accessed May 25, 2023].

[69] "Oman Australia Cable (OAC)," https://sub.co/, 2022, [Online; accessed May 25, 2023].

[70] "INDIGO-Central Cable," https://www.submarinecablemap.com/submarine-cable/indigo-west, 2019, [Online; accessed May 25, 2023].

[71] "Coral Sea Cable System (CS2)," https://coralseacablecompany.com/, 2020, [Online; accessed May 25, 2023].

[72] "Kumul Domestic Submarine Cable System," https://www.pngdataco.com/, 2019, [Online; accessed May 25, 2023].

[73] "Palapa Ring East," https://palaparing.id/, 2019, [Online; accessed May 25, 2023].

[74] "Moratelindo Broadband Company," http://moratelindo.co.id, [Online; accessed May 25, 2023].

[75] "Kupang-Alor Cable System Cable," https://www.submarinecablemap.com/submarine-cable/kupang-alor-cable-systems, 2019, [Online; accessed May 25, 2023].

[76] "Sape-Labuan Bajo-Ende-Kupang Cable System," https://www.submarinecablemap.com/submarine-cable/sape-labuan-bajo-ende-kupang-cable-systems, 2021, [Online; accessed May 25, 2023].

[77] "Denpasar-Waingapu Cable System," https://www.submarinecablemap.com/submarine-cable/denpasar-waingapu-cable-systems, 2019, [Online; accessed May 25, 2023].