A Hop Away from Everywhere: A View of the Intercontinental Long-haul Infrastructure

ESTEBAN CARISIMO, Northwestern University, USA
CALEB J. WANG, Northwestern University, USA
MIA WEAVER, University of Wisconsin-Madison, USA
FABIÁN E. BUSTAMANTE, Northwestern University, USA
PAUL BARFORD, University of Wisconsin-Madison, USA

We present a longitudinal study of intercontinental long-haul links (LHLs) – links with latencies significantly higher than that of all other links in a traceroute path. Our study is motivated by the recognition of these LHLs as a network-layer manifestation of critical transoceanic undersea cables. We present a methodology and associated processing system for identifying long-haul links in traceroute measurements. We apply this system to a large corpus of traceroute data and report on multiple aspects of long haul connectivity including country-level prevalence, routers as international gateways, preferred long-haul destinations, and the evolution of these characteristics over a 7 year period. We identify 85,620 layer-3 links (out of 2.7M links in a large traceroute dataset) that satisfy our definition for intercontinental long haul with many of them terminating in a relatively small number of nodes. An analysis of connected components shows a clearly dominant component with a relative size that remains stable despite a significant growth of the long-haul infrastructure.

CCS Concepts: • **Networks** → **Topology analysis and generation**.

Additional Key Words and Phrases: Long-Haul Links (LHL), intercontinental links, submarine cables

ACM Reference Format:
Esteban Carisimo, Caleb J. Wang, Mia Weaver, Fabián E. Bustamante, and Paul Barford. 2023. A Hop Away from Everywhere: A View of the Intercontinental Long-haul Infrastructure. *Proc. ACM Meas. Anal. Comput. Syst.* 7, 3, Article 47 (December 2023), 26 pages. https://doi.org/10.1145/3626778

1 INTRODUCTION

As part of ongoing work focused on the criticality of the submarine cable network [68], we set ourselves to map traceroute measurements to submarine cables as a first step towards understanding its potential vulnerabilities. We meant to follow an approach, first introduced in a position paper [7], building on the assumption that given a traceroute, one could identify the link with the largest latency and, if the associated routers were found near submarine landings, that this link could be mapped to one or a handful of cables. Alas, a preliminary analysis of traceroute datasets disabused us of this assumption.

Authors' addresses: Esteban Carisimo, esteban.carisimo@northwestern.edu, Northwestern University, Evanston, Illinois, USA; Caleb J. Wang, Northwestern University, Evanston, Illinois, USA; Mia Weaver, University of Wisconsin-Madison, Madison, Wisconsin, USA; Fabián E. Bustamante, Northwestern University, Evanston, Illinois, USA; Paul Barford, University of Wisconsin-Madison, Madison, Wisconsin, USA.

2476-1249/2023/12-ART47
https://doi.org/10.1145/3626778

While we found traceroutes matching these expectations, we also found many others in which the routers associated with a submarine traversing link were far inland from the closest landing points, some as far as 700 km. Further analysis revealed a large number of these long-haul links, covering distances larger than 10,000 km and connecting every country in the world. Figure 1 illustrates an example long-haul link connecting Seattle, US with Singapore. The example illustrates how such links may rely on the concatenation of more than one submarine cable segment, in this case Pacific Crossing-1 (PC-1), the Japan Information Highway (JIH), and the Asia Submarine-cable Express (ASE).

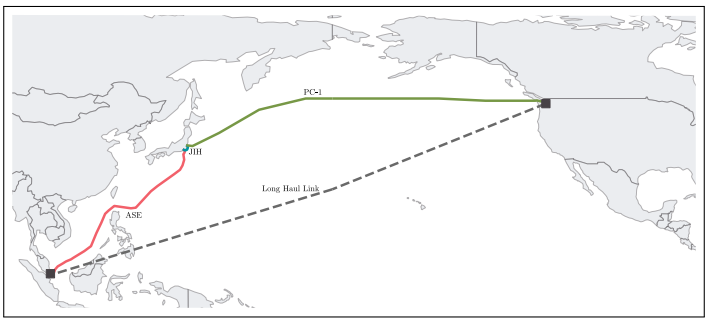

Fig. 1. Example of a long-haul link connecting Seattle, US with Singapore. The example also illustrate the link as a concatenation of several submarine cable segments: PC-1, JIH and ASE in this case.

This paper presents the first in-depth, longitudinal study of intercontinental long-haul links and their preferred destinations, as the network-layer manifestation of critical transoceanic undersea cables. We look to answer several questions, including: How long are these long-haul links (LHLs), particularly by comparison with submarine cables?, Are LHLs more common in certain parts of the world?, Where are their most common end points? How far are end points from submarine cable landings?, Can we infer the network virtualization techniques in use?, and How has the collection of LHLs and termination points changed over time?

To that end, as our *first contribution*, we present a methodology and associated processing system for identifying intercontinental LHLs in large traceroute datasets. Our methodology uses LevelShift [5] to identify potential LHLs, and relies on CAIDA's ITDK [9] for alias resolution and to geolocate associated routers and assign them to autonomous systems, before validating LHLs based on SoL constraints.

As our *second contribution*, we apply this methodology to a large corpus of traceroute data (231.45M traceroutes), collected by the CAIDA Archipelago platform [10], and report on multiple aspects of this LHL network. Focusing on a recent snapshot (composed of three measurement cycles[1]), we identified 85,620 LHLs (out of 2,674,577 total links in a traceroute dataset of 32.83M) directly connecting 31,773 routers in nearly every country of the world (170 or ≈92% of all countries), with 10% of LHLs connecting in single hop routers 193 ms away in countries separated by over 12,500 km. We show that LHLs are substantially longer than any intercontinental submarine cable segment, with a median RTT of 130 ms ms, ≈84% larger than the median RTT of submarine cable segments (70.76 ms).

We discovered preferred destinations for LHLs where 80.4% of them terminate in the US, and 67.69% interconnect nodes in the North Atlantic. We explore properties of these nodes and identify high-degree vertices, or super routers, connecting up to one-fourth of all countries in our dataset

[1]Each Ark cycle is a traceroute campaign covering all /24 subnets from a whole team of vantage points [12].

(28 countries), with most of them operated by Tier-1 transit providers and ≈58% of them located in the US. The termination points of these LHLs are located quite far from the nearest, on route, landing points. Specifically, 64% of them are at a distance of 500 km from the nearest landing point, and in 10% of the cases, that distance exceeds 3,513 km.

Many of the identified LHLs result from the wide adoption of virtualization technologies (e.g., MLPS), which hide physical links in virtual network-layer links connecting pairs of nodes, as far as Sao Paulo and Tokyo, to each other in a single hop. We investigated the adoption of identifiable MPLS tunnels and found a wide range of adoption levels across networks, with an average of 2.54% but with networks with adoptions higher than 90%, such as Vodafone-AS1273 using it in 95.8% of its LHLs. Beyond complicating infrastructure mapping efforts, the adoption of MPLS tunnels challenges routing optimization and debugging in presence of path inflation, in some cases resulting from auto-bandwidth algorithms [59].

As our *third contribution*, we carry out a longitudinal study of the LHL graph, exploring topological changes over a multi-year period starting in 2016. In just seven years, the number of edges grew 2.9x, from 18,026 to 52,066, while the number of vertices grew 2.4x (from 9,560 to 23,267). Despite this growth, some properties of the LHnet has remained stable, including the inter-hop latency distribution has well as the prevalence of the intra-AS LHLs among the total number of LHLs (79% to 89%).

In sum, we make the following key contributions:

- We introduce a methodology and processing system to identify LHLs.
- We apply this methodology to a large corpus of traceroute data and present the first in-depth study of intercontinental LHLs and their preferred destinations.
- We report on the evolution of LHLs properties over a multi-year period starting in 2016.
- We release source code and artifacts to facilitate the reproducibility of our study (available at https://github.com/NU-AquaLab/intercont-LHL-2024).

This study contributes to the community effort to create consistent maps across layers of the Internet, from AS-level, to logical and physical connectivity, critical to a range of important analysis from performance and robustness to security [4, 23, 48].

2 LONG-HAUL LINKS

Key to our analysis is the identification of long-haul links in large traceroute datasets. In the following paragraphs, we present a working definition of these links before discussing our inference methodology.

2.1 A working definition

We define a *long-haul link (LHL)* as a pair of consecutive IP addresses in a traceroute path separated by a latency that, in no-congestion scenarios, differs significantly from other latencies in the path.[2] We identify LHLs in large-scale traceroute campaigns and considered them independently of the underlying physical technologies connecting the consecutive hops (*i.e.,* physical mediums, link layer technologies).

Given our long-term goal of understanding the criticality of the transoceanic submarine cable network, we focus on intercontinental LHLs, which we define as a LHL separated by a set latency threshold (§3.2) and where the pair of routers are in different continents.

[2]That is, in statistical terms the latency of the LHL is an outlier.

2.2 Distances for long-haul links

In related work studying the domestic long-haul infrastructure of the continental US, Durairajan *et al.* [23] defines LHLs as those links connecting major city-pairs, spanning at least 30 miles or connecting population centers of at least 100,000 people. We focus on intercontinental LHLs and our definition relies, instead, on link latencies and the location of both ends of LHLs.

To identify a meaningful latency threshold for intercontinental LHLs, we investigate distances between peering facilities as an estimator of link lengths. We use peering facilities as the importance of these infrastructures makes them strong candidates as end points of intercontinental LHLs. Presence in these facilities (e.g., via remote peering) enable direct peering with thousands of networks simultaneously.

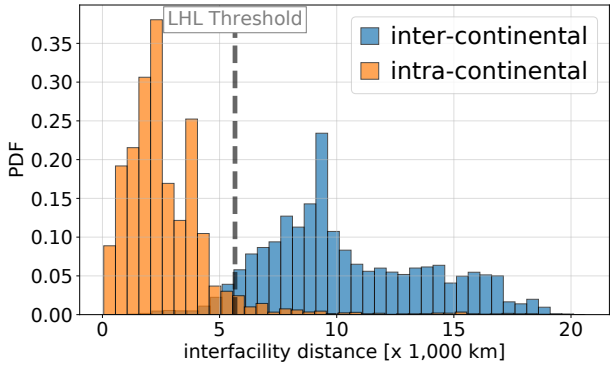

Fig. 2. Histograms of intra-/inter-continental distances, in kilometers, between the peering facilities listed on PeeringDB.

We use a recent PeeringDB snapshot (October 2022) and investigate differences in inter- and intra-continental distance between all peering facilities to identify a suitable threshold for intercontinental LHL. We examine the distribution of distances between all pairs of networks present in different peering facilities (e.g., there would be eight distances associated with two peering facilities with two and four networks present at each). We use great-circle distance[3] to compute the distance between each pair of peering facilities and repeat the process for each network pairs.

Figure 2 shows histograms of the intra- and inter-continental distances (in kilometers) between *all* peering locations. We observe a shift in the intercontinental distribution with little overlap of both distributions beyond 5,700km. Nearly 95% of networks are at a maximum distance of 5,657km from other networks present at peering facilities in the same continent. Within that distance, there are only 5% of networks reachable at inter-continental peering facilities.

Based on this analysis, we set the LHL threshold of 5,700km or the equivalent of 57ms for RTTs propagating at $\frac{2}{3} \cdot c$ in traceroute measurements (assuming optical fibers and symmetric paths). This is a conservative lower bound since cable infrastructure is rarely deployed in straight lines [8]. Thus, we restate our definition of intercontinental LHLs as *a LHL with a latency of at least 57ms RTTs, where the pair of routers at each side of the LHL are in different continents.* The appendix (§C) includes a sensitivity analysis of this threshold.

[3]https://mathworld.wolfram.com/GreatCircle.html

3 IDENTIFYING LONG-HAUL LINKS

To identify intercontinental LHLs in traceroute datasets, we first select candidate LHLs by detecting significant changes in RTTs (§3.1), with latencies over the *LHL threshold* (§3.2). We augment the list of candidate links with topological information (§3.3), and apply a two-stage filter to keep intercontinental LHLs using router geolocation (§3.4), and Speed-of-Light constraints (§3.5). Figure 3 illustrates this process.

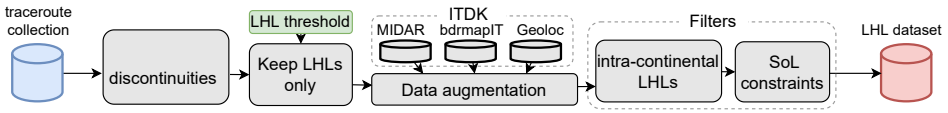

Fig. 3. Identifying LHLs in traceroute measurements by (1) detecting discontinuities in RTT using anomaly detection, (2) adding topological information and selecting based on (3) latency, (4) connecting intercontinental destinations and (5) with consistent propagation delays.

3.1 Detecting discontinuities in traceroute

We expect any given sequence of hops collected by traceroute to include at most one, and rarely more than one, LHL, making it easy to detect LHL candidates as large latency jumps. A potential risk is the presence of false positives where significant latency jumps result from reasons other than propagation delay such as queuing, latency inversions or the presence of middleboxes. We take several measures to address this, as described in the following paragraphs.

To identify LHL candidates as discontinuities in hop latency we frame the problem as anomaly detection and adopt a commonly used tool for this task, ADTK [5]. The ADTK *LevelShift* inference method detects sequence shifts to a persistent state with higher values. The method compares median values in two adjacent windows (hops) and identifies level shifts when the difference exceeds historical interquartile range adjusted by a constant ($c = 1$ in our implementation). This allows us to exclude transitory latency changes while detecting increases of the propagation delay. To gain further confidence in the links identified based on latency discontinuities, we conservatively keep links where both the previous and subsequent routers are responsive.

To illustrate the steps of our methodology we use the swath of the /24 IP space measured by one probe (`jfk-us` in New York City, US) in one cycle (*6537*, April 1, 2018) of the CAIDA Ark dataset. This dataset includes 56,637 traceroutes to addresses in different /24 networks. After applying our anomaly detection tool we are left with 47,400 candidate LHLs out of 609,230 hops probed in 36,214 traceroutes. Note that we use this slice of the dataset only for illustrative proposes and should not be seen as a representative sample of the full dataset (considering it is part of a single cycle captured by a single vantage point).

3.2 Latency threshold

The first step in the process of LHL identification yields a set of candidate LHLs with latencies significantly different than their preceding and following hops.

In this step we apply the *LHL threshold* (§2.2) to candidate LHLs. We leverage multiple RTT samples, grouping all pairs of IP addresses connecting two routers. Utilizing multiple samples allows us to mitigate temporary latency fluctuations generated by network congestion, overloaded slow paths, and other factors that contribute to measurement noise. We aggregate these samples to obtain the minimum RTT difference between router hops (min_{diff}), and then exclude all pairs of routers where the min_{diff} is below the *LHLs threshold*.

Delay-based filtering has some clear limitations. Using the difference of RTT measurements between hops may result in inaccurate estimations of a hop latencies since packets expiring in different hops may have completely different reverse paths [64]. These estimations may also contain false positive inferences (*i.e.,* incorrectly include short-haul links) in the presence of latency inflation resulting from circuitous paths [42], diurnal congestion patterns [20] or misconfigurations. The next stages use additional information to address some of these limitations.

3.3 Augmenting the dataset

Before selecting intercontinental LHLs from the set of candidates using different filtering conditions, we need to first augment the information associated with candidate LHLs using topological and geolocation information.

We rely on CAIDA's Internet Topology Data Kit (ITDK) [9] for alias resolution [43], and to add router geolocation [55] and router-to-AS mappings [54]. The ITDK is generated from a subset of the traceroute data gathered by Ark and it includes two related IPv4 router-level topologies, router-to-AS assignments, geographic location of each router, and reverse DNS lookups of all observed IP addresses [9]. Using ITDK, out of the 6,619 different routers in the illustrative dataset, we can identify the ASN of 6,553 (99.0%) and the location 6,542 (98.83%) of them.

3.4 Removing intra-continental LHLs

We use router geolocation tags to select LHLs with end points in different continents. This step reduces the impact of latency jumps unrelated to propagation delay, including transient latency increments (*e.g.,* congestion and flash crowds) or other pathological events such as two consecutive probes expiring in the same router, followed by a LHL.

Router geolocation inferences are prone to contain errors [60], which we inherit from the ITDK dataset. We mitigate their impact by focusing most of our analyses at a country-level where geolocation services have been shown to provide better accuracy [37, 49].

After removing intra-continental LHLs from our dataset, we are left with 595 LHLs of the total set of candidate LHLs.

3.5 Imposing speed-of-light constraints

The last step in our process of identifying intercontinental LHLs relies on Speed-of-Light (SoL) constraints, using minimum distances between countries, to address potentially incorrect geolocation inferences.

For this, we use countries' boundaries represented by polygons in the dataset to compute the nearest pair of points between *all* countries and obtain the minimum distance between *all* pairs of countries. We use these minimum distances to identify and remove long-haul links where inter-hop latency differences violate SoL constraints.

After applying our complete process we are left with 571 intercontinental LHLs out of the 609,230 hops in our initial traceroute dataset.

4 DATASET

Our methodology for LHL identification can be applied to any traceroute dataset. In this section, we describe the specific datasets we use for our analysis.

4.1 Traceroute measurements

The core of our analyses leverages traceroute measurement campaigns collected by the CAIDA's Archipelago (Ark) platform [10]. Ark is the the only platform offering constant, network-wide

Table 1. Snapshots of traceroute measurement campaigns collected by CAIDA's Ark platform.

year	cycles	# probes	# traceroutes
2016	4576, 4577, 4578	97	32.72M
2017	5422, 5423, 5424	117	33.08M
2018	6446, 6447, 6448	149	32.71M
2019	7615, 7616, 7617	112	32.72M
2020	8820, 8821, 8822	121	33.99M
2021	9643, 9644, 9645	69	32.84M
2022	10019, 10020, 10021	95	32.83M
7 years	21	244	231.45M

traceroute scans, enabling longitudinal studies of Internet-wide features. Ark's measurement campaigns are topological explorations that use a /24 granularity to cover all IPv4 prefixes announced in BGP routing tables [6]. In each campaign, or *cycle*, all /24 subnets are probed from the set of vantage points, but any single destination /24 will be probed by only one monitor in each probing cycle [12].

Table 1 shows the measurement cycles included in our analysis with details on the number of vantage points and traceroute measurements. We combine three consecutive measurement cycles of the same day (*e.g.,* 6446-6448 for 2018) to capture links that may have been missed due to packet loss and to help identify and remove transitory latency inflation with additional RTT samples.

We use a recent set of cycles (1) from 2022(10019-10020-10021), to study the long-haul infrastructure today (§5). For the longitudinal part of our analysis (§6), we use data collected over a period of 7 years starting in 2016.

4.2 Complementary datasets

As noted in Sec. 3, we augment our long-haul inferences with additional topological information from CAIDA's ITDK kit [9]. To examine the data at different granularities, we include three topological features to our inferences: *(i)* router-to-country mappings using MaxMind-based geolocation inferences [55], *(ii)* IP-to-router geolocation inferences using bdrmapIT [54], and *(iii)* router-to-AS mappings using MIDAR alias resolution [43]. We also include PTR records for the IP addresses (router hostnames) in the LHL data collection since this information can provide insights on geolocation [52], customers [53] or other relevant features of network deployments and structure.

We use daily PeeringDB snapshots, collected during the days of the measurement cycles, to investigate characteristics of peering ecosystems as a driver of intercontinental long-haul connectivity. Despite inaccuracies in PeeringDB records, and its limited and potentially biased coverage, previous studies have shown that it includes a representative picture of the network [44, 50]. Last, we also rely on other datasets for our analysis including CAIDA's AS relationship files [11, 32], CAIDA's IXP dataset [13], and geographic information [56].

5 LONG-HAUL CONNECTIVITY TODAY

We begin our study looking at two basic characteristics of intercontinental LHLs: *length* and *termination points*. Using 2022 traceroute snapshots (CAIDA's Ark cycles 10019-10020-10021), we explore continent pairwise link lengths (§5.1) and preferred termination points at country (§5.4) and router levels (§5.5). We use these empirical observations to contrast our *ex ante hypothesis* — network ingress and egress hops to submarine segments are in the vicinity of physical-layer termination points (landing points).

5.1 LHLs: Lengths and destinations

What is the length (or latency equivalent) of LHLs? The LHL identification process (i.e., described in §3) finds 85,620 LHLs connecting 31,773 routers in 170 countries. To compare the latencies of these LHLs with those of the underlying infrastructure, we compute the propagation delay using the length of intercontinental submarine cable segments reported by Telegeography.

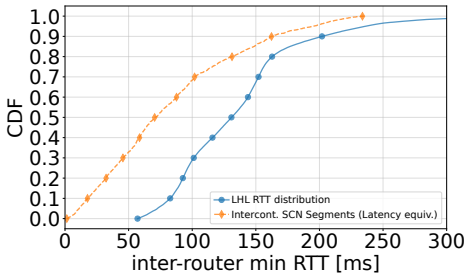

Fig. 4. Cumulative distribution of long-haul inter-router latency.

Figure 4 shows the cumulative distributions of latency (in milliseconds) for both LHLs and the computed latency of submarine cable segments. The inter-router latency distribution shows a somewhat steep distribution with 75% of LHLs having latencies between ~60 ms and ~155 ms. The distribution flattens at 160 ms into long tail for the last 20% of the LHLs. For context, the widest point of the largest ocean in an east-west direction is 19,300 km, extending almost halfway around the world, or an estimated round-trip time of ≈215 ms.

The latency equivalent for submarine cable segments, in contrast, presents a smoother distribution with a median RTT of 70.76 ms, compared with the median RTT of 130 ms for LHLs. At the 70th percentile, the RTT for LHLs is ≈ 1.5x that of RTT of intercontinental submarine cable segments. The discrepancy between network-layer latencies and the latency-equivalent distances is due, in part, to the fact that a single LHL result from the composition of multiple submarine paths (as the example of Fig. 1 illustrates).

What locations are these LHLs connecting? Figure 5 presents the latency distribution of identified LHLs on a per-continent basis. This figure shows that, to a different extents (prevalence is shown in brackets), all continents have LHLs connecting each other. We observe that North America connects to Asia and Oceania where speed-of-light constraints gaps are visible in the distribution, likewise in South America to Asia and Oceania and Europe to Oceania.

The reach of some of the identified LHLs is impressive, with some single hops connecting distant locations such as Los Angeles (US) and Budapest (HU) or Sao Paulo (BR) and Tokyo (JP), some of them up to ~20,000 km apart. There is no submarine infrastructure directly connecting some of these points, such as South America and Asia. This is a another sign of decoupling between the physical infrastructure and the network layer, with network-layer links including multiple physical segments concatenated.

5.2 A compass view of LHLs

What are the most common orientation of LHLs? We assumed that, if only for historical reasons, the majority of the underlying infrastructure behind LHLs would be oriented East-West. On the other hand, the rapid growth of Internet connectivity in the southern hemisphere (*e.g.*, Brazil, Oceania) may mean a growing number of LHLs supporting connections to the large infrastructure (e.g., datacenters, IXPs) hosted predominately in the northern hemisphere.

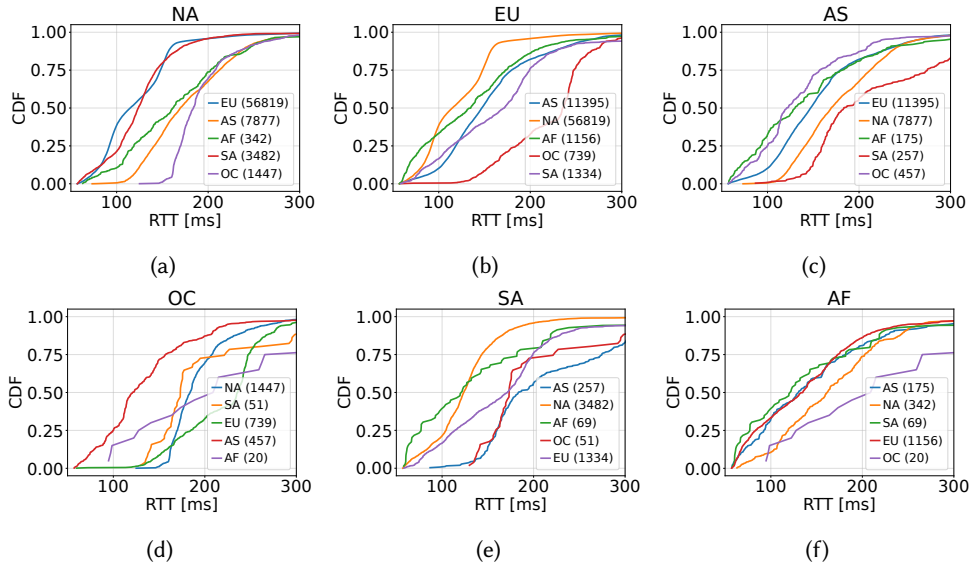

Fig. 5. Inter-router latency distribution (the number of LHLs per distribution is shown between parenthesis).

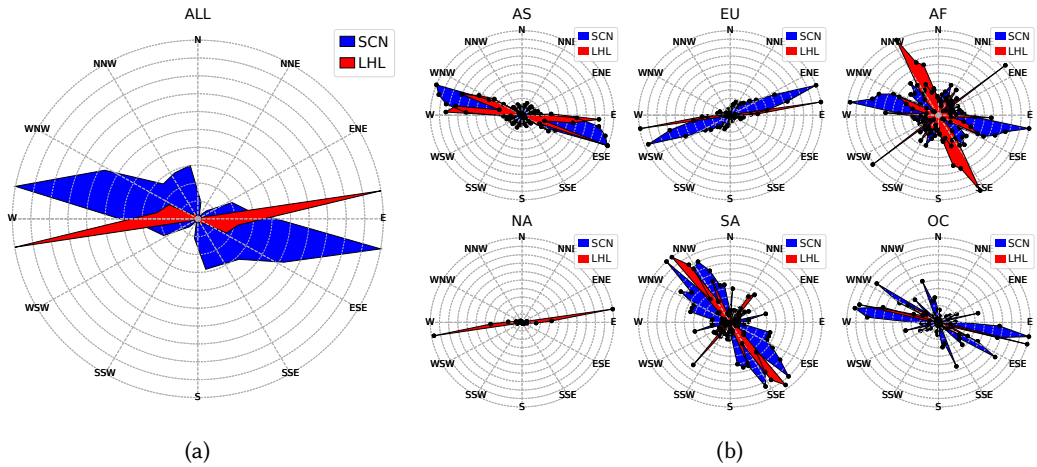

Fig. 6. Prevalence of inter-router connections directions in the 12-wind compass rose.

To visually compare the orientations (*e.g.*, North-Northeast to South-Southwest) of both LHLs and the underlying intercontinental submarine cable segments, we rely on polar histograms. A polar histogram is a bar chart that uses a polar coordinate system, with the length of each bar capturing the normalized frequency of a specific coordinate in a dataset. Figure 6 shows seven polar histograms over wind roses comparing the orientations of both LHLs and submarine cable segments[4] when looked at all together (Fig. 6a) and separated by continent (Fig. 6b). Note that the

[4]The orientation of an intercontinental submarine cable segment is given by the orientation of the straight line connecting the two landing points of the segment.

maximum diameter of each dataset, in this and the by-continent histograms, is normalized (i.e., the total number of LHLs is not the same as, but larger than the number of cable segments).

The left histogram, aggregating all LHLs and SCNs, shows that the majority of the LHLs (78%) and cable segments (37%) connects routers in the East-West direction, though LHLs show a wider spread in that direction. These global patterns align with our expectations and can be explained by the prominence of inter-country links and submarine cable infrastructure following this orientation in the North Atlantic, mostly connectivity the US and key Western European countries, including Germany, France, the UK and the Netherlands (§5.4).

The by-continent histograms reveal some interesting trends. For starters, they reveal a certain degree of independence of the LHLs from the underlying submarine infrastructure. While some intercontinental LHLs rely on a single cable segment laid out in the same general orientation, many are concatenation of several cable segments bridging end points (routers) not directly connected at the physical layer. Examples of the former case can be seen in the histograms of Asia, Europe and, particularly, North America. On the other hand, Africa, and to a lesser extent Oceania and South America, shows a clear decoupling between the LHLs and the underlying cable infrastructure. These latter patterns can be explained by the number of LHLs bridging regions with no direct submarine cable connectivity (e.g., Africa and North America) and the presence of regions serving as "stepping stones" (e.g., Western Europe linking Africa and North America, or Africa linking Oceania and Asia with Europe).

5.3 Visible MPLS in LHLs

Are these MPLS tunnels? While it is not possible to characterize the adoption of all virtualization mechanisms, in the following paragraph we investigate the fraction of visible MPLS tunnels among the LHLs in our traceroute datasets. While prior work has examined the length of MPLS in the network topology [22, 63], to the best of our knowledge, no prior work has focused on the geographic span of tunnels nor their implementation to serve intercontinental routes.

The visibility of MPLS tunnels depends on the combination of RFC4950-compliant nodes and ingress nodes enabling the `ttl-propagate` option [22]. RFC4950-compliant nodes append the MPLS label stack of the time-exceeded message to the ICMP packet providing traceroute visibility to Label Switching Routers (LSR) of the Label Switched Path (LSP). If the first MPLS router of an LSP copies the IP-TTL value to the LSE-TTL field rather than setting the LSE-TTL to an arbitrary value such as 255, LSRs along the LSP will reveal themselves via ICMP messages even if they do not implement RFC4950. If MPLS nodes implement both RFC4950 and `ttl-propagation`, MPLS tunnels are called *explicit tunnels* while if they implement RFC4950 but not `ttl-propagation`, they are referred to as *opaque tunnels* where only the last hop of the LSP is visible.

We can identify the presence of MPLS tunnels in LHLs when the ICMP time-exceeded message of far-side node contains MPLS labels in the payload. More sophisticated inferring techniques can improve visibility of *invisible tunnels* [67]. In this work we investigate the most elemental implementation of MPLS tunnels leaving a more exhaustive explorations for future work.

We find evidence of the use of MPLS in ≈ 2.54% of all LHLs (2,181 LHLs) connecting 2,161 routers (≈7%). We look at the prevalence of LHLs over visible MPLS tunnels from a country-, continent- and AS-level perspective, and find low prevalence of this type of tunnels with weighted average values of ≈3%, ≈3%, ≈8% at the country, continent and AS levels, respectively. While the prevalence of these tunnels is far from being homogeneous across ASes, we find a notable high adoption rate in some specific networks, such Claro Brazil (AS4230), NTT (AS2914), Uruguay's ANTEL (AS6057), and Vodafone (AS1273), ranging between 60% and 95% (60%, 66.1%, 90%, and 95.8%, respectively).

5.4 Preferred long-haul destinations

What are the preferred destinations of LHLs? We focus on the preferred destinations of long-haul connections and whether preferences vary across continents. Different factors may determine preferred destinations, including technical issues such as content availability and peering opportunities, geographic features challenging long-haul deployments (*e.g.,* being able to anchor landing points), cultural affinity (*e.g.,* a common language) and economic cooperation (*e.g.,* the European Union partially financed the EASSy cable in East Africa [26]).

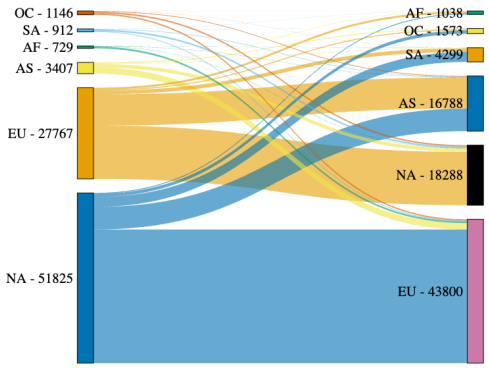

Fig. 7. Intercontinental long-haul connectivity.

We first explore preferred long-haul destinations at continent-level granularity. Figure 7 shows a Sankey diagram of the prevalence of LHLs between continents. The large majority of LHLs have the near-side router in North America, with far-side routers in all regions, but most commonly in Europe. North America is the major contributor to LHLs terminating in Europe and South America and a remarkable actor in the rest of the regions. We find a correlation between these preferences and the number of submarine cables connecting continents. For instance the number of submarine cables that connects Asia with Europe (28) doubles the number that connects Asia with North America. The North Atlantic routes dominate intercontinental LHL connectivity with 67.69% of all LHLs in this snapshot.

Focusing on the preferred long-haul destination, Fig. 8 shows the top 10 preferred countries for long-haul connectivity across different regions. The vertical axis on the left side of each figure shows the number of LHLs while the horizontal axis lists the top 10 countries, in order, among the preferred destination for a given region. We observe the US as the preferred destination for all regions, except Africa, where the fraction of *all* inter-router long-haul links ending in the US drops to 0.15.

The preference of African countries for major European hubs over the US may be motivated by proximity and other factors such as cultural affinity. Previous studies have shown, for instance, the prevalence of French operators across French-speaking countries in Africa [28] and the presence of circuitous paths in African connectivity including detours to London and Amsterdam [33]. A similar cultural influence is visible in South America where the second preferred destination is Spain. Another observation is that despite efforts, a large number of countries rely on content that is hosted in or routed through North America [24, 40, 48]. Other countries common to all regions top-10 listing includes Singapore, Great Britain, Germany and France. Singapore is a major hub with a large, state-owned transit ISP SingTel (AS7473) [15] while Great Britain, Germany and France host some of the largest IXPs in the world (LINX, DE-CIX and France-IXP).

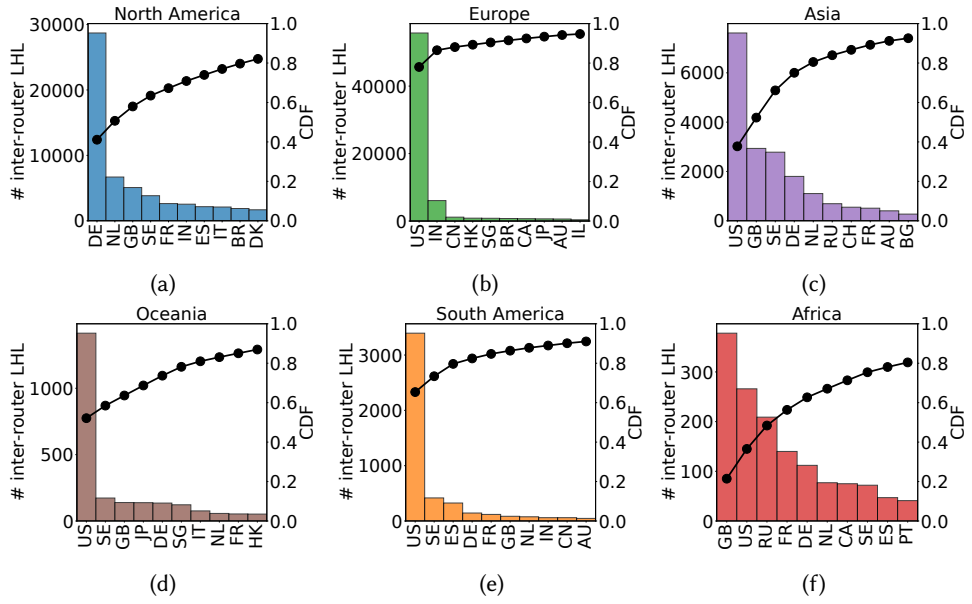

Fig. 8. Preferred destinations for long-haul connectivity across regions.

5.5 Super routers

Where are the most common termination points? Figure 9a shows the node degree distribution of LHL termination points. As the figure shows, node degrees range widely from a handful to as many as 1,326, and a clear long-tail distribution with the top 5^{th} percentile of vertices having node degrees larger than 13. Changing perspectives to the country-level connectivity of vertices, we find the top 5^{th} percentile connecting between 1 and 24 countries!

We call this popular LHLs destinations *super routers*, given their large router and country-level reach. Specifically, we define a *super router*[5] as those routers with directly connected (i.e., next hop) with large number of routers scattered across several countries. Figure 10 illustrates the idea of *super routers* with an example of a router operated by GTT (AS3527) in Seattle, Washington connecting neighboring routers scattered in 25 countries.

We investigate the geographic distribution of *super routers*, focusing on their location and country-level footprint.

Given the preferred orientation of LHLs (East-West), some of the popular destinations at either end of LHLs, as well as the distribution of their length, one would imagine LHLs to be the direct, *network-level view* of submarine cables links and thus to have end-points, primarily, in coastal areas, more or less evenly distributed near submarine cables' landing points. While there are instances of LHLs matching this intuition, our analysis shows that many popular routers are found in-land – as far in-land as Chicago (US), over 700 km (as the crow flies) from the closest landing point (Tuckerton, US)!

The country-level footprint of routers associated with LHLs is shown in Fig. 9b. The graph plots the CCDF of the number of countries that each router is connecting to via LHLs. While 90% of super-routers connect at most 3 countries, the top 1%, 2% and 5% super routers connect at least to 11, 8 and 4 countries, respectively. Among ASes, we note the vast majority operating routers

[5]Appendix B shows raw traceroute sequences traversing a *super router*.

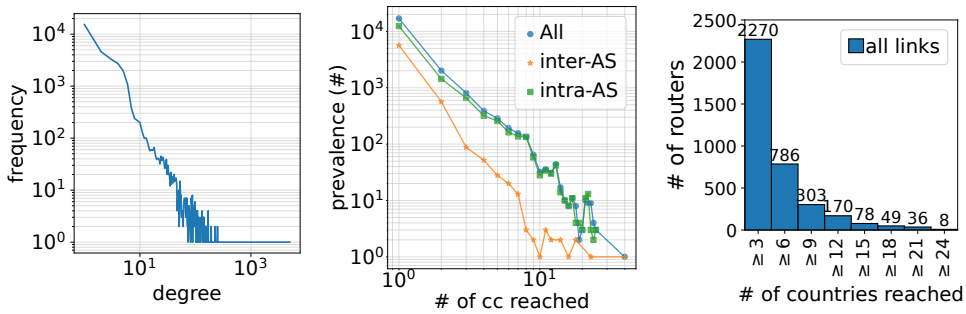

(a) (b) (c)

Fig. 9. Node degree distribution (Fig. 9a) and country-level footprint of routers shown as a CCDF (Fig. 9b) and bar graph (Fig. 9c) .

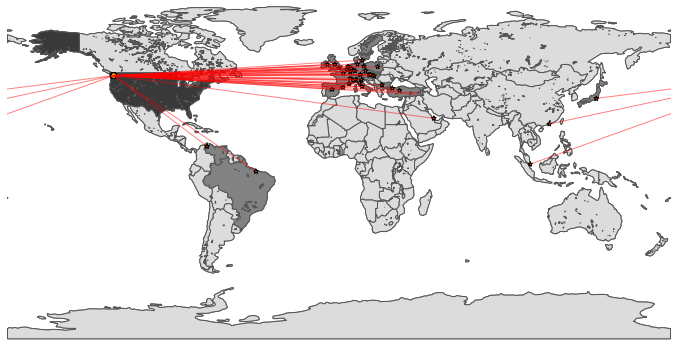

Fig. 10. An example of a super router (GTT, AS3257).

connecting to 3 or 6 countries, with a selective group of 11 ASes (mainly Tier-1 transit providers) operating super routers as international gateways connecting 9 or more countries.

Figure 11 shows the top 8 ASes (top) and countries (bottom) holding the largest number of super-routers connecting 5 countries or more via LHLs. We observe that international Tier-1 transit providers lead in the use of super routers with AS3356 (Lumen, formerly Level3), AS9894 (Bharti Airtel), AS3257 (GTT), AS6762 (Telecom Italia) at the top of this list.

From a country-level perspective, the US hosts by far the largest number of super router, followed by India, Sweden and Germany. The presence of super routers in Germany can be explained, at least in part, by the scale of the IXPs it hosts. Other popular countries are headquarters of large transit networks (Telecom Italia-AS6762 and Arelion-AS1299 in Italy and Sweden, respectively). The dominance of the US is perhaps not surprising given that the country hosts over a quarter of locations for many of the top cloud computing services such as Google Cloud Platform, Microsoft's Azure and IBM (25.7%, 25%, 27.3%, respectively), and over half of Amazon's locations (52.6%), and is also relatively far from the rest of the world.

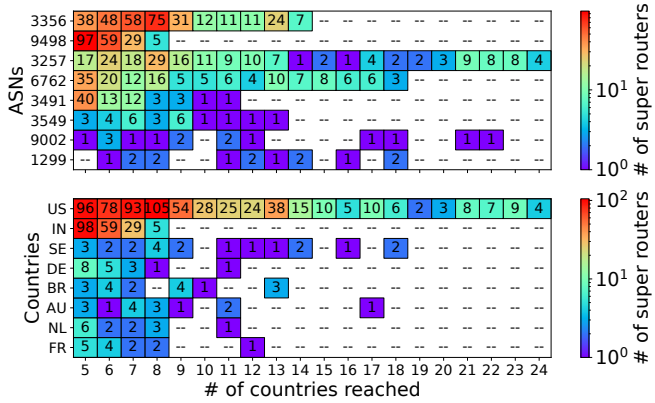

Fig. 11. Top 8 ASes and countries with *super routers* connecting at least 5 countries

5.6 Router-Landing Point decoupling

How far are these common termination points from submarine cable landing points?

Based on the previous observation that LHLs terminate in routers that are hundreds of kilometers away from the coast, we conclude our analysis by exploring the distance between common termination points and submarine cable landings.

We build upon recent efforts to map physical paths in network infrastructure by extending the approach to the router level. We use iGBD [4], a physical infrastructure database that compiles various data sources to identify the underlying physical paths. To enhance this approach, we modify iGBD's path selection algorithm to prioritize submarine links and adopt a weighted path selection method. We apply the updated version of iGDB to the set of LHLs and compute the distance between link termination points and submarine landing points at both ends.

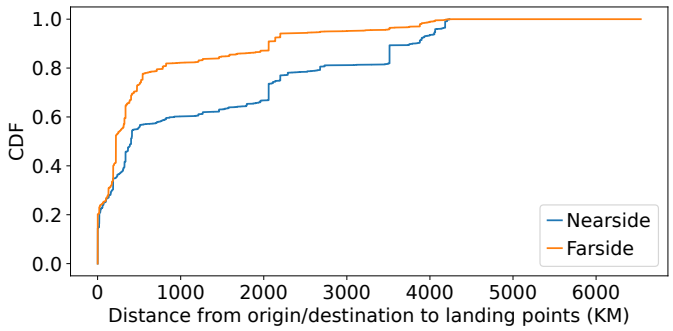

Fig. 12. Distance between LHLs termination points and submarine landing points.

Figure 12 shows the CDF of these distances. We observe that that only 26% of the routers are located within 100 km of the nearest submarine cable landing point, while a significant fraction are at much greater distances: 64% at a range of 500 km or less and 10% at distances exceeding 3,513 km. The geography of the countries in our dataset partially explains this; LHLs that terminate in regions like Hong Kong and Singapore tend to be close to landing points, while those terminating with large regions such as the United States, Russia, and Australia contribute to the larger distances.

These observations further highlights the decoupling between network links and the underlying physical infrastructure that challenges most infrastructure mapping efforts.

5.7 Takeaways

The previous paragraphs focus on answering some key questions about intercontinental long-haul connectivity: How long are most LHLs and how do they compare with submarine cables?, Are LHLs more common in certain parts of the world? Can we infer the network virtualization techniques in use? Where are their most common end points and how far are they from cable landing points?

We find that a quarter of LHLs have at least 155 ms delay and run mostly East-West, connecting locations in the North Atlantic and the US with the far East. We find a very skew distribution of *visible* MPLS tunnels across LHLs with average adoption of values, at the country, continent and AS levels, of ≈3%, ≈3%, ≈8%, respectively, but with over 90% adoption by some operators. We found that at a country-level, the US is the preferred long-haul destinations for most regions, except for Africa. We introduce the concept of *super routers* – nodes that aggregate multiple LHLs and connect to several countries simultaneously – and showing them to be most commonly hosted by Tier-1 ASes and found in the US. We show that a significant fraction of these routers can be found further than 3,000 km from the nearest submarine cable landing point.

6 A LONGITUDINAL PERSPECTIVE

How has the long-haul link infrastructure changed over time? In the following paragraphs we explore topological changes in LHL connectivity over a period of 7 years, starting in 2016. We control for sampling bias introduced by CAIDA Ark's expansion over this period, by focusing our analysis on the 53 vantage points (≈ %21.7 of all active probes in that period) with snapshots in at least 5 years of our dataset.

We explore changes in the long-haul network (*LHnet*) at router, Autonomous Systems (AS) and country levels, looking at variations in node degree distribution (§6.1) and in the degree of each node (§6.2). We also explore changes in the composition of the LHnet over time, focusing on the most densely connected nodes (§6.5) and studying characteristics of the connected components in these graphs (§6.6). We conclude our analysis looking at changes in the per-network inter-hop latency (§6.3) and the composition of business relationships between networks in both ends of the LHLs (§6.4).

6.1 Changes in the long-haul network

Figure 13 shows the node degree distribution for the router (Fig. 13a), AS (Fig. 13b) and country-level (Fig. 13c) graphs derived from the LHnet for snapshots over a 7-year period. In the three cases, the log-log plots show characteristics of *heavy-tailed* distributions suggesting that these graphs can be explained by the *preferential attachment* model [2]. We apply Clauset *et al.* [18] methods to discern whether these distributions fit better to lognormal or power law distributions. Our results show that the power law fits well for the AS- and country-level distributions, with minor changes over time and a subtle shift towards a slower tail decay.

6.2 Node degree changes over time

We examine changes in the connectivity of each individual node in the AS- and country-level graphs, focused on determining whether individual ASes and countries have developed a more dense intercontinental connectivity during this 7-year period.

Figure 14 shows the cumulative distribution of the year-to-year and 2016-to-2022 variations in the node degree of nodes in the AS- (Fig. 14a) and country-level (Fig. 14b) graphs. We observe that both graph undergo major changes over time, at the AS-level, the 2016-distribution shows

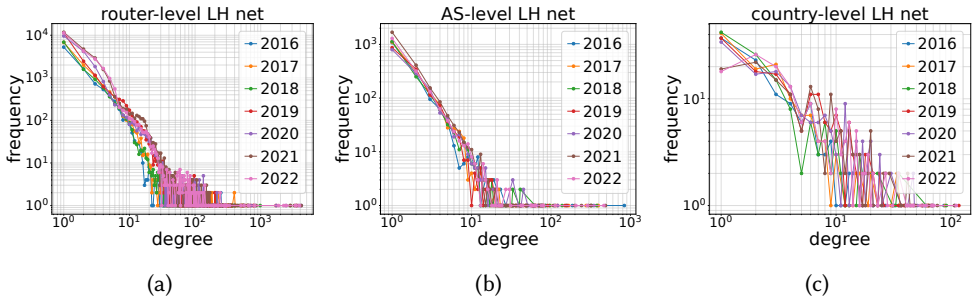

Fig. 13. Node degree distribution for router, AS and country-level graphs derived from the LHnet.

a symmetric distribution meaning neutral changes overall, while at the country-level is skewed towards positive values with a mean node degree variation of 7.17. These changes indicate that the LHnet has been reshaped at the AS-level but new nodes replace connectivity of declining nodes, on the other hand, changes at the country level show a that countries are getting more densely connected over time.

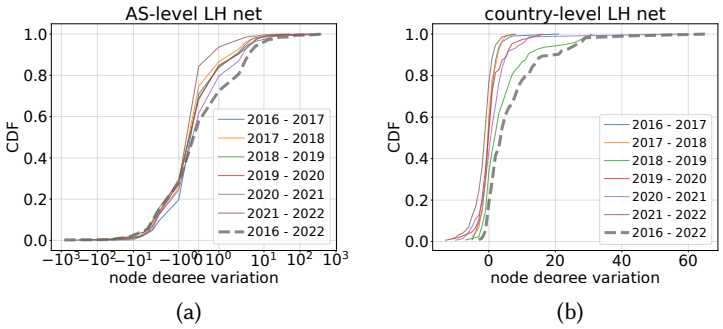

Fig. 14. Distribution of YoY and 2016-2022 variations of node degrees in the AS and country-level graphs.

6.3 The LHL length variations

Based on our previous observations that showed some changes in the characteristics of the LHnet, we now focus on the inter-router latency. We use our multi-year dataset to investigate changes in the AS-level inter-router latency that could describe modifications in the structure of the network that were not visible from previous macroscopic granularities.

To investigate changes in the composition of long-haul connectivity at AS level, we compare the mean LHL inter-router latency for networks present in the 2016 snapshot with the networks in the 2019 and 2022 ones. Figure 15 shows a scatter plot for each comparison. The axes show the mean LHL inter-router latency for (15a) the 2016 and 2019 snapshots, and (15b) the 2016 and 2022 snapsots. In both, the size of the circles is a function of the variation in the number of per-AS LHL, with colors indicate positive (green) or negative (red) variations. The red-dashed diagonal line in each figure indicate whether networks increased or decreased their mean inter-router latency if their dots are above or below the line. Comparing both figures, the 2022 plots show a continuation of the trends observed in 2019 including the growth in the number of LHLs of DTAC-3320 and Bharti Airtel-9848. We focus on networks with significant variations in the number of LHLs (>100) in Fig. 15b where we observe a wide variety of trends in this 6-year period. There is a group of

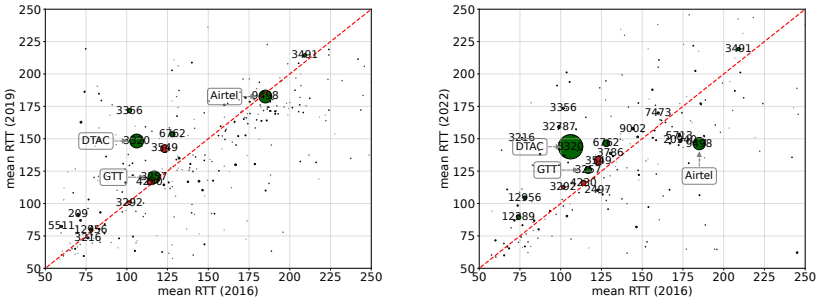

(a) (b)

Fig. 15. Sizes and colors show per-AS link variations and sign.

large providers with neutral changes in the mean LHL inter-router latency despite having a growth (Telecom Italia-6762, SingTel-7473, GTT-3257, IIJ-2497) or a decay (Level3-3549, Claro Brasil-4230, TDC-3292) in the number of LHLs. Another example of stability over time is PCCW with the largest mean LHL inter-router latency (despite a significant growth) in our analysis which is given by a large fraction of LHLs connecting far distant locations such as Miami and Frankfurt or Tokyo and Frankfurt. This figure show some networks with clear changes in the number of LHLs and the mean LHL inter-router latency, DTAC-3320 (+17,317 LHLs, +17,317 ms), Telxius-12956 (+395 LHLs, +26.48 ms), Bharti Airtel-9848 (+4,447 LHLs, -38.78 ms). There are several possible explanations for these changes such as deployments of new submarine cables, organization-level reconfigurations (Level3), mergers and acquisitions (DTAC purchase of Sprint [66]), and technological upgrades (adoption of MPLS).

6.4 Commercial relationships behind LHLs

We shift our attention to the AS relationship between both ends of long-haul links. For this part of the analysis we use CAIDA's AS relationship files [11, 32], curated from both BGP and traceroute-derived sources from RouteViews and RIPE RIS collectors.

Table 2. Intra-/inter-domain LHLs (and unmapped) over time.

Year	Intra-domain	Inter-domain	Unmapped	# LHL
2016	5101 (0.79)	1354 (0.21)	17 (0.00)	6472
2017	6173 (0.80)	1538 (0.20)	40 (0.01)	7751
2018	7086 (0.79)	1888 (0.21)	0 (0.00)	8974
2019	12406 (0.81)	2746 (0.18)	224 (0.01)	15376
2020	14632 (0.89)	1750 (0.11)	2 (0.00)	16384
2021	20677 (0.84)	3876 (0.16)	2 (0.00)	24555
2022	18398 (0.85)	3162 (0.15)	0 (0.00)	21560

We first look at the fraction of long-haul links that are intra- or inter-domain links. Table 2 list both, over time, and their respective fractions. We observe a rapid growth in the number of LHLs – which doubles over the observation period – with the fractions of intra-domain LHLs growing at fastest pace. Throughout the period of analysis, the vast majority LHLs in our dataset are intra-domain (between 79 and 89%).

For the remaining LHLs, we list the fraction of links based on the categories that can be inferred in Table 3. We find that roughly one-fifth of the LHLs correspond to inter-domain links, with customer-to-provider (c2p) links responsible for two-thirds of the inter-domain LHLs. In c2p and p2p scenarios, it is difficult to identify which of the two sides provides the physical connectivity by either allocating its own resources or purchasing capacity. In some cases, we can not find any inferred AS relationship between reported ASes at either end of the LHL. A small fraction of these non-inferred relationships corresponds to incorrectly mapped peering relationships at IXPs.

Table 3. Commercial relationships between ASes at both ends of Inter-domain LHLs.

year	Inter-domain		non-inferred ToR	
	p2p	p2c	IXP	Unknown
2016	893 (0.13)	461 (0.07)	131 (0.02)	197 (0.03)
2017	1044 (0.13)	494 (0.06)	130 (0.02)	310 (0.04)
2018	1286 (0.14)	602 (0.06)	170 (0.02)	211 (0.02)
2019	1873 (0.12)	873 (0.06)	187 (0.01)	403 (0.03)
2020	1285 (0.08)	465 (0.03)	90 (0.01)	212 (0.01)
2021	2961 (0.12)	915 (0.04)	100 (0.00)	548 (0.02)
2022	2510 (0.11)	652 (0.03)	106 (0.00)	637 (0.03)

6.5 A stable LHnet core

The previous analyses looked at the LHnet as a whole, finding a general trend towards a more connected network. In the following paragraphs, we focus on the most prominent nodes of the network, the group of most densely connected nodes, to understand whether this core shapes the structure of the LHnet.

	TOPcore members
AS level	Cogent-174, Airtel-9498, NTT-2914, Telecom Italia-6762, Telxius-12956, Arelion-1299, GTT-3257, LUMEN-3356, HE-6939, TATA-6453, China Unicom-4837
Country level	South Africa, India, Germany, France, Great Britain, United States, China, Japan, Hong Kong, the Netherlands, Singapore, Australia, Italy, Turkey, Canada and Brazil

Table 4. TOPcore members at AS and county levels present at least in 6 different snapshots

We apply k-core decomposition [3, 16] to identify the group of most densely connected nodes in each snapshot. We divide a graph into subsets (equivalent to tiers) where each one is generated recursively removing all nodes with degree lower than (k, *shell-index*) until the degree of all nodes is larger or equal to the k threshold. We then identify the TOPcore (tier with the highest shell-index) and look for ASes and countries with long-term presence at the TOPcore.

Table 4 shows ASes and countries that were in the TOPcore in 6 different years. The TOPcore in the AS-level graph has a stable group os ASes composed of are large international carriers (*e.g.,* Cogent-AS174, Hurricane Electric-AS6939, LUMEN-AS3356) typically inferred to be part of the *transit-free clique* [51]. At a country level, the stable TOPcore is primarily composed of countries known as international or regional hubs, with vast presence of submarine cable networks (*e.g.,* the US, Brazil, India and Singapore) or hosting large IXPs (DECIX, AMS-IX, IX.br).

	router-level graph					AS-level graph					country-level graph				
	LHnet (G)			max(C_G)		LHnet (G)			max(C_G)		LHnet (G)			max(C_G)	
year	$\|C_G\|$	$\|V\|$	$\|E\|$	$\|V\|$	$\|E\|$	$\|C_G\|$	$\|V\|$	$\|E\|$	$\|V\|$	$\|E\|$	$\|C_G\|$	$\|V\|$	$\|E\|$	$\|V\|$	$\|E\|$
2016	898	9560	18026	5443	13723	179	1633	2326	1406	2123	1	122	374	122	374
2017	1133	12355	23751	7618	19871	177	1753	2613	1528	2411	1	134	413	134	413
2018	1084	11478	20615	5880	15478	164	1698	2458	1492	2277	1	135	433	135	433
2019	1511	18410	38243	12185	32309	325	1543	2297	1135	1925	1	167	775	167	775
2020	1298	18597	40452	13324	35536	253	1446	2235	1104	1931	1	164	828	164	828
2021	1704	7970	8697	3310	5548	454	1410	1834	833	1313	1	143	963	143	963
2022	1605	23267	52066	15078	40518	269	1965	2805	1611	2488	1	158	863	158	863

Table 5. Graph dimension of the *LHnet* and the largest connected component.

6.6 The long-haul network

Table 5 shows the evolution of the number of vertices, edges and connected components and the size of the sub-graph containing the largest connected component (nodes and vertices) for the router-, AS- and country-level graphs in a 7-year timeframe. Notably, the three graphs show that the largest connected component comprises a large fraction (or even all) nodes creating a contiguous network that represents *the intercontinental long-haul backbone of the public Internet*. At level that, despite the LHnet being composed of hundreds of connected components, the largest connected component includes between 41.5% to 71.6% of nodes and 63.8% to 63.8% of links. For AS- and country-level perspective, the connected component comprises up to 92.6% and 100% of the nodes and 87.7% and 100% of the edges, respectively. These observations indicate links and routers visible in traceroute campaigns create a contiguos intercontinental network containing 1,611 ASes (≈ 2.5% of all active ASes in October 2022) and 158 countries (≈ 81% of all countries).

6.7 Takeaways

This section focuses on understanding how the collection of LHLs and termination points has changed over time. We found minor fluctuations in node degree distribution at the router, AS and country levels, though, with changes towards powerlaw distributions with lower α values. At a country level, we also discovered that the network is more densely connected meaning that new LHLs were created interconnecting new pairs of countries. We showed that the LHnet is composed by a stable group of networks and countries in its core and the largest connected component of the long-haul network contains up to 71.6% of all nodes linked by LHLs making it *the intercontinental long-haul backbone of the public Internet*. We also found a wide range of changes in the inter-router latency at the AS level, with some networks creating new and longer LHLs.

7 PRACTICAL IMPLICATIONS AND LIMITATIONS

In this section, we discuss practical implications of our observations and limitations of our study. We focus primarily on implications on studies of the submarine cable network and its resilience, but touch briefly on other topics, from network management and operations to cyber sovereignty.

Submarine cable networks: Private conversations with operators affiliated to the SubOptic community, confirm the adoption of network virtualization and other techniques to *move landing points inland* to the proximity of datacenters, suggesting that the list of threats to submarine cable connectivity should include those of the infrastructure connecting the cable's landing point with its final destination. This is well illustrated by the city-wide network outage that the Colombian city of Cali experienced in 2021 (visible from IODA [39]) due to a cable cut in the 100km cable segment that connects the city with the submarine landing point [61].

Consequences of sea-level rise for coastal network infrastructure is a growing concern across multiple entities including the UN [57], researchers [35], the submarine cable organizations [38] and other communities. Being unable to identify physical infrastructure underneath these network-layer links limits our ability to identify LHL exposed to natural disasters and climate change threats. Technologies such MPLS [67] or SD-WAN [36, 41], increase the opaqueness of the network, limits our understanding of underlying physical characteristics of the network, and challenges a thorough assessment of network resilience.

Management and operations: The opaqueness of a tunneled network structure reduces the ability to debug the network through ICMP-based tools (*i.e.,* pings and traceroutes). The lack of basic path-discovery tools capable of mapping elements within tunnels impedes the ability to easily detect path changes caused by traffic engineering (*e.g.,* load balance) or reconfigurations (*e.g.,* rerouting, outages). In absence of these tools, debugging is only available for a group of network operators with privileged access to devices in the network (*e.g.,* MPLS switches) using sophisticated tools.

A previous study conducted in Microsoft's WAN network showed persistent latency inflation caused by MPLS *autobandwidth* algorithms [59]. Other research studies observed path inflation and speculate that presence of MPLS tunnels could explain this behavior [8, 45]. In the context where LHLs are the result of MPLS tunnels, they could experience similar path and latency inflation.

Security and cyber sovereignty: The growing opacity of the network underlying infrastructure presents new challenges in terms of security and government regulations. In recent years, governments' cyber sovereignty concerns brought consequent updates to regulatory frameworks, such as GDPR [25] and LPDG [65]. These concerns sometimes include the routing system which may be also subject to regulations [34, 58]. Following those steps, research efforts evaluated whether end-to-end paths applied forwarding rules skipping specific countries. The vast adoption of tunneling mechanisms, however, challenges such assessments and attempts to validate regulation compliance or precisely identify potential vulnerable choke points.

Limitations. We acknowledge several limitations of this study, to help put our findings in perspective. For starters, the presence of routers that ignore ICMP messages and middleboxes [62] (*e.g.,* NAT, firewalls, *etc.*), may impact our latency-derived estimations. We believe, however, that our main observations such as on the scale and rapid growth of the long-haul infrastructure should hold. Our reliance on measurement collected from a volunteers platform may introduce biases on our observations resulting from its uneven adoption, while inaccuracies in topological datasets (*e.g.,* geolocation databases, alias resolution, router-to-AS mappings) may affect the topologies we generate. Our primary approach for addressing these issues was to be conservative in all of our design and implementation choices to the limit the impact on our results.

Appendix C includes reults from a sensitivity analysis including timeframe, dataset size, threshold variations and filters' contributions.

8 RELATED WORK

The router- and AS-level Internet topologies have been widely studied from multiple angles, including graph-theoretical models [19, 27], commercial relationships [30, 51], flattening and rewiring [21], among others.

Several research efforts have focused on documenting structural changes of the Internet to accommodate the rise of new technologies (*e.g.,* video streaming, smartphones). The transition from a hierarchical network into a flat structure [21] in the early 2000s is well documented. The irruption of CDNs gained a large deal of attention with studies focusing on changes in topological and traffic characteristics [46], the widespread of direct peer-to-peer connectivity of CDNs [17], and *off-net* cache deployments [31], among others. This transition also included the consolidation

of IXPs as key pieces of the Internet topology, creating peering fabrics with volumes of traffic comparable to those of Tier-1 Transit providers [1], and expanding to all regions [14, 28].

A handful of previous studies have focused on long-haul connectivity. Durairajan *et al.* [23] investigated the domestic long-haul infrastructure of the United States using information extracted from optical cable deployments along pre-existing transport infrastructure. Bishof et al. [7] puts forward a research agenda focused on the criticality of the submarine cable network. Fanou *et al.* [29] studied the impact of new submarine cable deployments in developing regions and the drop in latency to cross the Atlantic. Liu *et al.* [48] found that submarine cable infrastructure enables fetching resources contained in most popular websites.

International connectivity is at times closely related to geopolitics and foreign affairs. Levin *et al.* [47] investigated traffic censorship of intermediary countries along the path of international routes and routing avoidance of specific censorship regions. Future research directions could combine long-haul connectivity and its implications on geopolitics.

The resulting LHL graph of intercontinental LHLs its high-degree vertices are a reminder of previously reported Tier-1-network graphs [69] and may similarly hide important details about the underlying physical paths, relevant to resilience and cyber-sovereignty discussions [34, 47, 58], and challenge existing approaches from root cause analysis of failures to congestion control.

9 CONCLUSIONS AND FUTURE WORK

This paper presented the first in-depth, longitudinal study of intercontinental long-haul links (LHLs) and their preferred destinations, as the network-layer manifestation of critical transoceanic undersea cables. The study contributes to an ongoing community effort to create consistent maps across layers of the Internet, from AS-level, to logical and physical connectivity, critical to a range of important analysis including performance, robustness to security. We presented a methodology for identifying LHLs in traceroute measurements, and reported on our analysis of the long-haul infrastructure using a large corpus of traceroute data collected at the edge of the network. We found a vast and rapidly growing network with links spanning over 10,000 km and nodes that connect as many as 45 countries. Despite its rapid growth, we found a graph with key characteristics and a core component that remain stable over time. This new perspective opens a wide range of promising directions for future research, from alternative views of the long haul infrastructure to an exploration of that infrastructure's key properties and temporal stability.

10 ACKNOWLEDGEMENTS

This work is supported by the National Science Foundation grants CNS-2107392, CNS-1703592, CNS-2039146, and CNS-2106517. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation. The authors would also like to thank Erick W. Contag of SubOptic for the valuable conversations maintained in the context of this work.

REFERENCES

[1] Bernhard Ager, Nikolas Chatzis adn Anja Feldmann, Nadi Sarrar, Steve Uhlig, and Walter Willinger. 2012. Anatomy of a large European IXP. In *Proc. of ACM SIGCOMM*.

[2] Réka Albert and Albert-László Barabási. 2002. Statistical mechanics of complex networks. *Reviews of modern physics* 74, 1 (2002), 47.

[3] José Ignacio Alvarez-Hamelin, Luca Dall'Asta, Alain Barrat, and Alessandro Vespignani. 2008. K-core decomposition of Internet graphs: hierarchies, self-similarity and measurement biases. *Networks and Heterogeneous Media* 3, 2 (2008), 371–393.

[4] Scott Anderson, Loqman Slamanatian, Zachary S. Bischof, Alberto Dainotti, and Paul Barford. 2022. iGDB: Connecting the Physical and Logical Layers of the Internet. In *Proc. of IMC*.

[5] ARUNO. 2022. ADTK Detectors. (2022).

[6] Robert Beverly, Arthur Berger, and Geoffrey Xie. 2010. Primitives for active Internet topology mapping: Toward high-frequency characterization. In *Proc. of IMC*.

[7] Zachary S. Bischof, Romain Fontugne, and Fabián E. Bustamante. 2018. Untangling the world-wide mesh of undersea cables. In *Proc. of HotNets*.

[8] Ilker Nadi Bozkurt, Waqar Aqeel, Debopam Bhattacherjee, Balakrishnan Chandrasekaran, Philip Brighten Godfrey, Gregory Laughlin, Bruce M Maggs, and Ankit Singla. 2018. Dissecting Latency in the Internet's Fiber Infrastructure. *arXiv preprint arXiv:1811.10737* (2018).

[9] CAIDA. 2019. Macroscopic Internet Topology Data Kit (ITDK). http://www.caida.org/data/internet-topology-data-kit/. (2019).

[10] CAIDA. 2023. Archipelago Measurement Infrastructure Updates. https://catalog.caida.org/details/media/2011_archipelago. (2023). Accessed: 2021-9-30.

[11] CAIDA. 2023. AS Relationships (serial-1). https://catalog.caida.org/details/dataset/as_relationships_serial_1. (2023). Accessed: 2022-1-15.

[12] CAIDA. 2023. CAIDA Ark IPv4 prefix-probing data. https://catalog.caida.org/details/dataset/ark_ipv4_prefix_probing. (2023). Accessed: 2021-12-22.

[13] CAIDA. 2023. Internet eXchange Points Dataset. https://catalog.caida.org/details/dataset/ixps/. (2023). Accessed: 2022-1-17.

[14] Esteban Carisimo, Julián M Del Fiore, Diego Dujovne, Cristel Pelsser, and J Ignacio Alvarez-Hamelin. 2020. A first look at the Latin American IXPs. *ACM SIGCOMM Computer Communication Review* 50, 1 (2020), 18–24.

[15] Esteban Carisimo, Alexander Gamero-Garrido, Alex C Snoeren, and Alberto Dainotti. 2021. Identifying ASes of state-owned internet operators. In *Proc. of IMC*.

[16] Esteban Carisimo, Carlos Selmo, J Ignacio Alvarez-Hamelin, and Amogh Dhamdhere. 2019. Studying the evolution of content providers in IPv4 and IPv6 internet cores. *The International Journal for the Computer and Telecommunications Industry* 145 (2019), 54–65.

[17] Yi-Ching Chiu, Brandon Schlinker, Abhishek Balaji Radhakrishnan, Ethan Katz-Bassett, and Ramesh Govindan. 2015. Are we one hop away from a better internet?. In *Proc. of IMC*.

[18] Aaron Clauset, Cosma Rohilla Shalizi, and Mark EJ Newman. 2009. Power-law distributions in empirical data. *arXiV* (2009).

[19] Reuven Cohen, Keren Erez, Daniel ben Avraham, and Shlomo Havlin. 2000. Resilience of the Internet to Random Breakdowns. *Phys. Rev. Lett.* 85 (Nov 2000), 4626–4628. Issue 21. https://doi.org/10.1103/PhysRevLett.85.4626

[20] Amogh Dhamdhere, David D Clark, Alexander Gamero-Garrido, Matthew Luckie, Ricky KP Mok, Gautam Akiwate, Kabir Gogia, Vaibhav Bajpai, Alex C Snoeren, and Kc Claffy. 2018. Inferring persistent interdomain congestion. In *Proc. of ACM SIGCOMM*.

[21] Amogh Dhamdhere and Constantine Dovrolis. 2010. The Internet is flat: modeling the transition from a transit hierarchy to a peering mesh. In *Proc. of CoNEXT*.

[22] Benoit Donnet, Matthew Luckie, Pascal Mérindol, and Jean-Jacques Pansiot. 2012. Revealing MPLS Tunnels Obscured from Traceroute. *ACM SIGCOMM Computer Communication Review* 42, 2 (mar 2012), 87âĂŞ93.

[23] Ramakrishnan Durairajan, Paul Barford, Joel Sommers, and Walter Willinger. 2015. InterTubes: A Study of the US Long-haul Fiber-optic Infrastructure. In *Proc. of ACM SIGCOMM*.

[24] Anne Edmundson, Roya Ensafi, Nick Feamster, and Jennifer Rexford. 2018. Nation-state hegemony in internet routing. In *Proc. of the ACM SIGCAS Conference on Computing and Sustainable Societies*.

[25] European Comission. 2022. General Data Protection Regulation. https://www.gdpreu.org. (2022).

[26] European Investment Bank. 2009. EU-Africa Infrastructure Trust Fund: Annual report 2009. https://www.eib.org/attachments/country/eu_africa_infrastructure_trust_fund_annual_report_2009_en.pdf. (2009).

[27] Michalis Faloutsos, Petros Faloutsos, and Christos Faloutsos. 1999. On Power-Law Relationships of the Internet Topology. (1999).

[28] Rodérick Fanou, Pierre Francois, and Emile Aben. 2015. On the diversity of interdomain routing in africa. In *Proc. of PAM*.

[29] Rodérick Fanou, Bradley Huffaker, Ricky Mok, and Kimberly C Claffy. 2020. Unintended consequences: Effects of submarine cable deployment on Internet routing. In *Proc. of PAM*. Springer, 211–227.

[30] Lixin Gao and Jennifer Rexford. 2001. Stable Internet routing without global coordination. *IEEE/ACM Transactions on networking* 9, 6 (2001), 681–692.

[31] Petros Gigis, Matt Calder, Lefteris Manassakis, George Nomikos, Vasileios Kotronis, Xenofontas Dimitropoulos, Ethan Katz-Bassett, and Georgios Smaragdakis. 2021. Seven years in the life of Hypergiants' off-nets. In *Proc. of ACM SIGCOMM*.

[32] Vasileios Giotsas, Matthew Luckie, Bradley Huffaker, and KC Claffy. 2014. Inferring complex AS relationships. In *Proc. of IMC*.
[33] Arpit Gupta, Matt Calder, Nick Feamster, Marshini Chetty, Enrico Calandro, and Ethan Katz-Bassett. 2014. Peering at the internetâĂŹs frontier: A first look at isp interconnectivity in africa. In *Proc. of PAM*.
[34] Melissa E Hathaway. 2014. Connected choices: how the Internet is challenging sovereign decisions. *American Foreign Policy Interests* 36, 5 (2014), 300–313.
[35] Rebecca Hersher. 2018. Rising Seas Could Cause Problems For Internet Infrastructure. https://www.npr.org/2018/07/16/627254166/rising-seas-could-cause-problems-for-internet-infrastructure. (2018).
[36] Chi-Yao Hong, Srikanth Kandula, Ratul Mahajan, Ming Zhang, Vijay Gill, Mohan Nanduri, and Roger Wattenhofer. 2013. Achieving high utilization with software-driven WAN. In *Proc. of ACM SIGCOMM*.
[37] Bradley Huffaker, Marina Fomenkov, and K Claffy. 2011. Geocompare: a comparison of public and commercial geolocation databases. In *Proc. of NMMC*.
[38] The International Cable Protection Committee (ICPC). 2021. Statement at the 21st Meeting of Informal Consultative Process on Seal-level Rise and its Impacts. https://www.un.org/depts/los/consultative_process/icp21/ICP21_item%203_ICPC_English.pdf. (2021).
[39] IODA. 2021. IODA Signals for Valle del Cauca between April 1 2021 and April 29 2021. https://bit.ly/3sGUtU0. (2021).
[40] Quentin Jacquemart, Clément Pigout, and Guillaume Urvoy-Keller. 2019. Inferring the Deployment of Top Domains over Public Clouds using DNS Data. In *Proc. of TMA*.
[41] Sushant Jain, Alok Kumar, Subhasree Mandal, Joon Ong, Leon Poutievski, Arjun Singh, Subbaiah Venkata, Jim Wanderer, Junlan Zhou, Min Zhu, et al. 2013. B4: Experience with a globally-deployed software defined WAN. In *Proc. of ACM SIGCOMM*.
[42] Ethan Katz-Bassett, John P John, Arvind Krishnamurthy, David Wetherall, Thomas Anderson, and Yatin Chawathe. 2006. Towards IP geolocation using delay and topology measurements. In *Proc. of IMC*.
[43] Ken Keys, Young Hyun, Matthew Luckie, and Kim Claffy. 2012. Internet-scale IPv4 alias resolution with MIDAR. *IEEE/ACM TONS* 21, 2 (2012), 383–399.
[44] Rowan Klöti, Bernhard Ager, Vasileios Kotronis, George Nomikos, and Xenofontas Dimitropoulos. 2016. A comparative look into public IXP datasets. *ACM SIGCOMM Computer Communication Review* 46, 1 (2016), 21–29.
[45] Rupa Krishnan, Harsha V Madhyastha, Sridhar Srinivasan, Sushant Jain, Arvind Krishnamurthy, Thomas Anderson, and Jie Gao. 2009. Moving beyond end-to-end path information to optimize CDN performance. In *Proc. of ACM SIGCOMM*.
[46] Craig Labovitz, Scott Iekel-Johnson, Danny McPherson, Jon Oberheide, and Farnam Jahanian. 2010. Internet inter-domain traffic. In *Proc. of ACM SIGCOMM*.
[47] Dave Levin, Youndo Lee, Luke Valenta, Zhihao Li, Victoria Lai, Cristian Lumezanu, Neil Spring, and Bobby Bhattacharjee. 2015. Alibi Routing. In *Proc. of ACM SIGCOMM*.
[48] Shucheng Liu, Zachary S Bischof, Ishaan Madan, Peter K Chan, and Fabián E Bustamante. 2020. Out of Sight, Not Out of Mind: A User-View on the Criticality of the Submarine Cable Network. In *Proc. of IMC*.
[49] Ioana Livadariu, Thomas Dreibholz, Anas Saeed Al-Selwi, Haakon Bryhni, Olav Lysne, Steinar Bjørnstad, and Ahmed Elmokashfi. 2020. On the Accuracy of Country-Level IP Geolocation. In *Proceedings of the Applied Networking Research Workshop*.
[50] Aemen Lodhi, Natalie Larson, Amogh Dhamdhere, Constantine Dovrolis, and Kc Claffy. 2014. Using peeringDB to understand the peering ecosystem. *ACM SIGCOMM Computer Communication Review* 44, 2 (2014), 20–27.
[51] Matthew Luckie, Bradley Huffaker, Amogh Dhamdhere, Vasileios Giotsas, and KC Claffy. 2013. AS relationships, customer cones, and validation. In *Proc. of IMC*.
[52] M. Luckie, B. Huffaker, A. Marder, Z. Bischof, M. Fletcher, and k. claffy. 2021. Learning to Extract Geographic Information from Internet Router Hostnames. In *Proc. of CoNEXT*.
[53] M. Luckie, A. Marder, B. Huffaker, and k. claffy. 2021-12. Learning Regexes to Extract Network Names from Hostnames. In *Asian Internet Engineering Conference (AINTEC)*.
[54] Alexander Marder, Matthew Luckie, Amogh Dhamdhere, Bradley Huffaker, kc Claffy, and Jonathan M Smith. 2018. Pushing the boundaries with bdrmapit: Mapping router ownership at Internet scale. In *Proc. of IMC*.
[55] Maxmind. 2023. Maxmind Geolocation Data. https://www.maxmind.com/en/geoip2-services-and-databases. (2023).
[56] Thierry Mayer and Soledad Zignago. 2011. Notes on CEPIIâĂŹs distances measures: The GeoDist database. (2011).
[57] The United Nations. 2016. Summary of the First Global Integrated Marine Assessment. https://www.un.org/depts/los/global_reporting/WOA_RPROC/Summary.pdf. (2016).
[58] Jonathan A Obar and Andew Clement. 2012. Internet surveillance and boomerang routing: A call for Canadian network sovereignty. In *TEM 2013: Proceedings of the Technology & Emerging Media Track-Annual Conference of the Canadian Communication Association (Victoria.*

[59] Abhinav Pathak, Mingn Zhang, Y Charlie Hu, Ratul Mahajan, and Dave Maltz. 2011. Latency inflation with MPLS-based traffic engineering. In *Proc. of IMC*.
[60] Ingmar Poese, Steve Uhlig, Mohamed Ali Kaafar, Benoit Donnet, and Bamba Gueye. 2011. IP geolocation databases: Unreliable? *ACM SIGCOMM Computer Communication Review* 41, 2 (2011), 53–56.
[61] RCN Radio. 2021. Falla en cable del Pacífico afecta a miles de usuarios de Internet en todo el país. https://www.rcnradio.com/colombia/falla-en-cable-del-pacifico-afecta-miles-de-usuarios-de-internet-en-todo-el-pais. (2021).
[62] Justine Sherry, Peter Xiang Gao, Soumya Basu, Aurojit Panda, Arvind Krishnamurthy, Christian Maciocco, Maziar Manesh, Jo ao Martins, Sylvia Ratnasamy, Luigi Rizzo, and Scott Shenker. 2015. Rollback-Recovery for Middleboxes. In *Proc. of ACM SIGCOMM*.
[63] Joel Sommers, Paul Barford, and Brian Eriksson. 2011. On the prevalence and characteristics of MPLS deployments in the open Internet. In *Proc. of IMC*.
[64] Richard A. Steenbergen. 2009. A practical guide to (correctly) troubleshooting with traceroute. NANOG.
[65] the Federative Republic of Brazil. 2022. Brazilian General Data Protection Law. https://www.lgpdbrasil.com.br. (2022).
[66] T-Mobile USA. 2020. T-Mobile Completes Merger with Sprint to Create the New T-Mobile. https://www.t-mobile.com/news/un-carrier/t-mobile-sprint-one-company. (2020).
[67] Yves Vanaubel, Pascal Mérindol, Jean-Jacques Pansiot, and Benoit Donnet. 2017. Through the wormhole: Tracking invisible MPLS tunnels. In *Proc. of IMC*.
[68] Colin Wall and Pierre Morcos. 2021. Invisible and Vital: Undersea cables and transatlantic security. *Center for Strategic and International Studies* June (2021).
[69] Walter Willinger, David Alderson, and John C. Doyle. 2009. Mathematics and the internet: A source of enormous confusion and great potential. *Notices of the American Mathematical Society* 56, 5 (2009), 586–599.

A ETHICS

This work does not raise any ethical issues.

B TRACEROUTES THROUGH A *SUPER ROUTER* – TWO EXAMPLES

Listings 1 and 2 show the results of two traceroute, collected by CAIDA's Ark monitors `jfk-us` (Lst. 1) and `ord-us` (Lst. 2), traversing Telia's (AS1299) *super router* in Chicago during the measurement cycle 8820 in October 2020. These results show that the Chicago router is the ingress point to a global link-layer backbone with egress points in different major cities across the US (Seattle, WA) and the world (Paris, France).

```
# traceroute from 216.66.30.102 (Probe hosted in NYC, NY, US. No AS info found)
  to 223.114.235.32 (MAXMIXD: Turpan, CN)
1  216.66.30.101  0.365 ms
2  62.115.49.173  3.182 ms
3  *
4  62.115.137.59  17.453 ms [x]  (chi-b23-link.ip.twelve99.net., CAIDA-GEOLOC
  -> Chicago, IL, US)
5  62.115.117.48  59.921 ms [x]  (sea-b2-link.ip.twelve99.net., RIPE-IPMAP ->
  Seattle, WA, US)
6  62.115.171.221  69.993 ms
7  223.120.6.53 69.378 ms
8  223.120.12.34  226.225 ms
9  221.183.55.110 237.475 ms
10  221.183.25.201  238.697 ms
11  221.176.16.213  242.296 ms
12  221.183.36.62 352.695 ms
13  221.183.39.2  300.166 ms
14  117.191.8.118 316.270 ms
15  *
16  *
17  *
18  *
```

```
19  *
```

Listing 1. Traceroute #1 traversing Telia's super-router in Chicago. FROM jfk-us (jfk-us.team-probing.c008820.20201002.warts.gz)

```
# traceroute from 140.192.218.138 (Probe hosted in Chicago, IL, US at Depaul
  University-AS20130) to 109.25.215.237 (237.215.25.109.rev.sfr.net., MAXMIXD:
  La Crau, FR)
1  140.192.218.129  0.795 ms
2  140.192.9.124  0.603 ms
3  64.124.44.158  1.099 ms
4  64.125.31.172  3.047 ms
5  *
6  64.125.15.65  1.895 ms      [x] (zayo.telia.ter1.ord7.us.zip.zayo.com., CAIDA
  -GEOLOC -> Chicago, IL, US)
7  62.115.118.59  99.242 ms    [x] (prs-b3-link.ip.twelve99.net., CAIDA-GEOLOC
  -> Paris, FR)
8  62.115.154.23  105.214 ms
9  77.136.10.6  119.021 ms
10  77.136.10.6  118.830 ms
11  80.118.89.202  118.690 ms
12  80.118.89.234  118.986 ms
13  109.24.108.66  119.159 ms
14  109.25.215.237  126.085 ms
```

Listing 2. Traceroute #2 traversing Telia's super-router in Chicago. FROM ord-us (ord-us.team-probing.c008820.20201002.warts.gz)

C SENSITIVITY ANALYSIS

We explored the sensitivity of our results to timeframe, dataset size (§4) and threshold variations, and the impact of removing geolocation and latency filters from our methodology (§3).

We analyzed fluctuations over a one-week period – a timeframe in which the network is expected to be stable – and found no changes in the characteristics of long-haul connectivity.

We also evaluated the impact of sampling the dataset and found that downsamples of 1:2 and 1:4 still captures a fraction of 0.97 and 0.82 of the nodes and 0.96 and 0.76 of the links, respectively.

An analysis of the sensitivity of our results to changes of the LHL threshold, varying it from 57ms to 20ms, shows stable results with only a 2.4% variations in the number of links (and 3.4% variation on the number of nodes) identified.

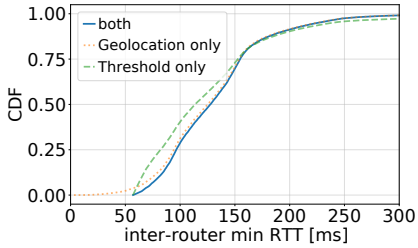

Fig. 16. Cumulative distribution of long-haul inter-router latency. ≈75% of the intercontinental inter-router latency difference ranges between 60 and 155ms.

In addition to evaluating the sensitivity of our findings to different LHL thresholds, we also analyzed the impact on our results of applying only the geolocation-based or latency-based filters as part of our methodology (§3). As in Sec.5, we use a recent dataset from Ark with cycle 2022 (10019-10020-10021). Figure 16 shows a CDF of inter-router latency differences (in milliseconds) resulting from applying our methodology, i.e., filtering out potential long-haul links based on geolocation data and our LHL threshold, as well as two additional CDF curves corresponding to the use of only a geolocation-based or a latency-based filter. The first CDF (*both*) correspond to the one included in Fig. 16. Note that the three curves show similar profiles with Jensen-Shannon scores of 0.10 and 0.15 between our method and geolocation-based and latency-based, respectively.

Received August 2023; revised October 2023; accepted October 2023