

Unintended consequences: Effects of submarine cable deployment on Internet routing

Rod  rick Fanou, Bradley Huffaker, Ricky Mok, KC Claffy

CAIDA/UC San Diego

Abstract. We use traceroute and BGP data from globally distributed Internet measurement infrastructures to study the impact of a noteworthy submarine cable launch connecting Africa to South America. We leverage archived data from RIPE Atlas and CAIDA Ark platforms, as well as custom measurements from strategic vantage points, to quantify the differences in end-to-end latency and path lengths before and after deployment of this new South-Atlantic cable. We find that ASes operating in South America significantly benefit from this new cable, with reduced latency to all measured African countries. More surprising is that end-to-end latency to/from some regions of the world, including intra-African paths towards Angola, increased after switching to the cable. We track these unintended consequences to suboptimally circuitous IP paths that traveled from Africa to Europe, possibly North America, and South America before traveling back to Africa over the cable. Although some suboptimalities are expected given the lack of peering among neighboring ASes in the developing world, we found two other causes: (i) problematic intra-domain routing within a single Angolese network, and (ii) suboptimal routing/traffic engineering by its BGP neighbors. After notifying the operating AS of our results, we found that most of these suboptimalities were subsequently resolved. We designed our method to generalize to the study of other cable deployments or outages and share our code to promote reproducibility and extension of our work.

1 Introduction

The underlying physical infrastructure of the Internet includes a mesh of submarine cables, generally shared by network operators who purchase capacity from the cable owners [8,48]. Little academic research has tried to isolate performance changes induced by the deployment of new submarine cables, although a few studies have investigated the end-to-end performance impacts of disruptions to existing cable operations [21,23]. Recently Bischof *et al.* [8,9] made a case for a new research agenda focused on characterizing the fundamental role these cables play in inter-continental connections. We agree with this aspiration and undertook a study that represents a step toward it.

In 2018, Angola Cables, Inc. (AC) deployed the first trans-Atlantic under-sea cables (SACS) crossing the South Hemisphere [56], linking Fortaleza, Brazil to Sangano, Angola [24,26,56,57]. We developed a methodology to analyze the impact of a specific cable launch on observed end-to-end round-trip latencies and paths across different regions of the world, and applied it to the case of the new SACS cable. The initial challenge in such a task is to identify the cable of interest using IP-layer traceroute measurements. Transit providers often

do not publicly disclose cable details, *e.g.*, IP addresses, and existing measurement techniques cannot easily distinguish multiple co-terminating (or nearly co-terminating) parallel cable systems [9]. The unique landing points of SACS created an opportunity to identify it in large-scale traceroute datasets: besides the fact that SACS is the first South-Atlantic cable system linking Africa to the Americas, only two cable systems (WACS and SACS) anchor at Sangano post-SACS, versus 18 to Fortaleza, the second landing point of the new cable.

Our high-level approach was to analyze traceroutes paths that crossed SACS from mid-Sept 2018 to late Jan 2019, to the paths those same endpoints traversed before the cable activation. This comparison revealed significantly reduced latency from ASes operating in South America toward Africa. However, we were surprised to find 21.3% of observed paths, with sources in Europe and Asia, as well as intra-African paths, experienced worse performance – in terms of higher RTTs across the corresponding endpoints – after SACS. Even more surprising, the median RTT of intra-African paths towards Angola doubled. We analyzed the root causes of these unintended consequences – suboptimal circuitous paths that unnecessarily crossed continents *e.g.*, from Africa/Europe toward Angola.

This work makes three contributions. First, we introduce a methodology to investigate submarine cable-related events, and second, we applied it to the case of the first operational South-Atlantic submarine cable to Africa. Finally, we suggest ways operators can avoid/mitigate suboptimal routing post-cable activation during future deployments. We emphasize that as third-party observers, we do not have access to traffic data: the observed suboptimality may occur on paths traversed by little to no traffic. In other words, this analysis does not necessarily reflect the performance of most traffic actually using that link. That said, these circuitous paths lasted the whole period considered in our analysis (*i.e.*, 3.5 months post-SACS) until we notified the provider.

We believe this work is the first attempt to evaluate the macroscopic impact of a new submarine cable on end-to-end paths and performance, and our results reveal how lack of diagnostic tools and exercises can amplify the existing routing inefficiencies involving the developing world, that derive from investment decisions, peering strategies [7,10,29,32,34], or traffic engineering [55].

2 Methodology

Our method requires first identifying the link of interest and its terminating IP addresses. We use these IP addresses to extract relevant paths from historical traceroute archives and then use this subset of paths to study the effects of the event on AS topology and performance (§2.1). We assume there is, and the method requires identifying, a cable of interest [1,2], and its IP addresses, which we call **link IPs**. The method also requires some meta-data about the event of interest, including date, duration, and the AS(es) operating the cable. We believe our method generalizes to the study of other cable deployments, and cable failures/outages. If the cable supports the use of layer-2 tunnels or wavelengths by different operators, identifying these link IPs is more complex and requires further study.

Step 1: Collect candidate IP paths that could have crossed the cable. We conduct traceroute measurements from vantage points (VPs) near the two cable endpoints toward each other; these are **candidate IP paths** that possibly traversed or rerouted through the link/cable after the event. Researchers can use public sea cable databases/maps [1,2,41,58] to inform the scheduling and execution of targeted traceroutes on existing measurement platforms [11,18,51].

Step 2: Identify router IP interfaces on both sides of the cable. This task requires disambiguating the IP addresses terminating the cable of interest from those terminating other cable systems. We combine two approaches: an RTT-threshold based on speed-of-light constraints and IP geolocation. We analyze only the traceroute hops inferred (using bdrmapIT [42]) to be owned by the AS of interest. For these hops, we look for an RTT difference gap between consecutive hops in traceroute, using a threshold of $t = \frac{2 \times l}{(2/3) \times c}$, where l is the physical length of the cable, and $\frac{2}{3}c$ is the speed of light traveling in fiber optics.

At this point, we can narrow down the set of consecutive hops to the ones that match the landing sites of the cable of interest. We use IP geolocation databases (*e.g.*, NetAcuity [25], MaxMind [43]) to map IP addresses to countries. Given the low accuracy of such geolocation databases for router infrastructure [33,35,49], we also apply hostname-based geolocation. We validate the inferred location of IP addresses adjacent to these IPs by measuring the RTTs from VPs located in the inferred country. We consider the geolocation correct if the minimum RTT is less than 10ms. We then resolve the router aliases of the selected IPs using CAIDA’s MIDAR [16], Vela *aliasq* [19], and ITDK [20]. We obtain two lists of IP addresses of the router interfaces at the two ends of the cable denoted by \mathcal{R}_A and \mathcal{R}_B . We call these two lists **link IPs**.

Step 3: Search for comparable historical traceroutes. We use \mathcal{P} to denote all source IP/destination prefix pairs, $\langle s, d \rangle$, where s is the VP’s source IP address and d is the longest-match prefix for the destination in the BGP routing table. We use longest-match because existing measurement platforms (Ark and Atlas) randomly probe within prefixes [17,52], and thus probing an exact destination IP address twice within a short period of time is unlikely. Furthermore, in many cases, only some IPs within a prefix respond to measurement probes [45]. We first look for a set of traceroutes, $\mathcal{T}_{\langle s, d \rangle}, \forall \langle s, d \rangle \in \mathcal{P}$, that contain either $\mathcal{R}_A \rightarrow \mathcal{R}_B$ or $\mathcal{R}_B \rightarrow \mathcal{R}_A$ *after* the occurrence of the event. With this list of prefix pairs, we search for pre-event traceroutes from the same $\langle s, d \rangle$ pairs, $\mathcal{T}'_{\langle s, d \rangle}$, for comparison.

Step 4: Annotate collected paths. For every hop in the traceroute sets $\mathcal{T}_{\langle s, d \rangle}$ and $\mathcal{T}'_{\langle s, d \rangle}$, we resolve the hostname and AS number, perform country-level IP geolocation, and compute the difference in RTT from that of the previous hop. To accurately map IP addresses appearing in traceroutes to AS numbers, we run bdrmapIT [42] on the traceroutes collected on each day from both Ark and Atlas, using as inputs *daily RIB from Routeviews and RIPE RIS* [44,53], CAIDA’s *AS relationship file* [13] from the first day of the month, a *daily dump of IXP prefixes from peeringDB* [39,46], and *WHOIS delegation files* collected in the middle of the period of the study. To resolve IP addresses to hostnames, we

use `zdns` [28] and `qr` [37]. Next, we collect a combined list of Internet eXchange Points (IXPs) prefixes from CAIDA’s IXP Dataset [14], compare them to the prefix corresponding to every hop in the traceroute sets $\mathcal{T}_{\langle s,d \rangle}$ and $\mathcal{T}'_{\langle s,d \rangle}$, and single out traces for which an IXP prefix matches the prefix of the IP hop. By doing so, we identify the IXPs through which the cable operator received/routed the packets pre and post-event.

We then group the traceroutes of each set by $\langle s, d \rangle$ pair and, based on their corresponding timestamps of execution, we cluster them per week. For every IP hop of each traceroute, we include its inferred annotations. These annotated traceroutes enable us to compare the AS paths and latency before and after the event using metrics described in  2.1.

2.1 Metrics for quantifying the impact of the event.

We compare the performance and AS paths between $\mathcal{T}_{\langle s,d \rangle}$ and $\mathcal{T}'_{\langle s,d \rangle}$ using three metrics.

1: RTTs to the common IP hops closest to the traceroute destinations.

This metric compares RTT values to reveal the change in latency across the network paths before and after cable deployment. Fig. 1 illustrates the identification of traceroutes in sets $\mathcal{T}_{\langle s,d \rangle}$ and $\mathcal{T}'_{\langle s,d \rangle}$ between the same $\langle s, d \rangle$ pair that share at least one IP address. Among all traceroutes run toward a destination prefix, we locate the common IP hop, h_c , closest to the destination IP and extract the RTTs from s to h_c in $\mathcal{T}_{\langle s,d \rangle}$ and $\mathcal{T}'_{\langle s,d \rangle}$, denoted as d_c and d'_c , respectively. We only consider the subset $\hat{\mathcal{P}}$ of $\langle s, d \rangle$, such that $\hat{\mathcal{P}} \subseteq \mathcal{P}$ and $h_c \neq \emptyset$ (*i.e.*, that contains non-empty h_c) in our analysis. For each $\langle s, d \rangle$ pair, we then compute the medians of all d_c and d'_c per week and choose their respective minimum values over the periods pre&post-event to mitigate noise.

2: AS-centrality of transit ASes in paths. We use `bdrmapIT` [42] to infer AS paths from the IP paths and compute from $\mathcal{T}_{\langle s,d \rangle}$ and $\mathcal{T}'_{\langle s,d \rangle}$ the AS-centrality of each observed transit AS. This metric is defined as the percentage of $\langle s, d \rangle$ pairs for which the AS path with the minimum observed RTT d_c (or d'_c for pre-event) contains the considered AS, and where that AS is neither the source nor the destination [29]. A higher AS-centrality of an AS post-event indicates increased transit importance, *i.e.*, more ASes use that AS for transit.

3: Length of AS paths crossing cable operator’s network pre&post event. We analyze the length of AS paths between source AS/destination prefix pairs observed to cross the cable operator’s network in RouteViews and RIPE RIS [44,53] data pre&post event. Similar to previous work [40], we consider paths collected on the first 5 days of the month before and after the event.




Fig. 1: Pre&post-event path comparison. Orange circles indicate common IP hops in $\mathcal{T}_{\langle s,d \rangle}$ and $\mathcal{T}'_{\langle s,d \rangle}$. s is the source IP address, and h_c is the common IP hop closest to the destination IP. The red circles $r_A \in \mathcal{R}_A$ and $r_B \in \mathcal{R}_B$ are the router interfaces of the two ends of the submarine cable (Step 2).

3 Data Collection: Case study of SACS cable deployment

We collected candidate IP paths that crossed SACS (§3.1) on Mar 25–26, 2019. We identified the link IPs from those candidate IP paths and ITDK [20] (§3.2). We used those link IPs to search in Ark and RIPE Atlas historical data for matching traceroutes post-SACS (Jan–mid-Sept 2018) and the traceroutes with the same $\langle s, d \rangle$ pairs pre-SACS (mid-Sept 2018–Jan 2019) (§3.3). Next, we annotated these traceroutes with supplementary information for its analysis (§3.4).

3.1 Collecting candidate IP paths crossing SACS

At the beginning of this study (Mar 2019) there were 8 active Ark VPs in South America, but none in Angola. AC hosted a looking-glass (LG) server [5] connected to the Sangano landing point [26,27,36]: An LG server allows BGP and traceroute queries by third-parties. Using both CAIDA’s *Vela* interface [18] to execute measurements on the Ark infrastructure, and the AC LG server [5], we collected traceroutes from VPs located in South America toward the AC LG server (and in the reverse direction) to obtain IP paths that possibly crossed SACS, *i.e.*, candidate IP paths.

3.2 Identifying link IPs

Based on the length of the cable, we estimate the round-trip time to cross SACS to be about $t_{SACS} = \frac{6,165km \times 2}{(2/3) \times c} = 62ms$. By inspecting the candidate IP paths, we found a pair of AC IP addresses (170.238.232.146 and 170.238.232.145) in the same /30, which had RTT differences with preceding and subsequent IPs that matched our latency heuristics. We could not resolve their hostnames, but the hostnames of their adjacent hops contained geolocation hints `ao.sgn` and `br.ftz`. Because of the small differential RTTs between the two IPs and their adjacent hops, we inferred that 170.238.232.146 and 170.238.232.145 were in Sangano, Angola and Fortaleza, Brazil, respectively. We leveraged VPs in Angola and Brazil to conduct latency measurements toward these two IPs to confirm our inference. Using the two IPs, we obtained a set of aliases of SACS routers in Angola (\mathcal{R}_A) and Brazil (\mathcal{R}_B) from ITDK [20] of Jan 2019. We found that \mathcal{R}_A and \mathcal{R}_B contained respectively 29 and 18 MIDAR-observed IP addresses aliases of the same router.

3.3 Fetching matching traceroutes paths

We analyzed CAIDA’s Ark [11] and RIPE Atlas data [51]. We considered the *on-going IPv4 Routed /24 Topology measurements* [17] from 178 Ark VPs that execute ICMP Paris-traceroute [6] toward a random destination in every routed /24 prefix. Using CAIDA’s *Henry* [15] interface to search Ark traceroute data, we split historical Ark traceroutes into two sets. ARK-AFTER includes traceroutes going through SACS from mid-Sept 2018 to late Jan 2019 (after SACS) and which had an IP of \mathcal{R}_A followed by an IP of \mathcal{R}_B or vice-versa; and ARK-BEFORE includes traceroutes from early Jan 2018 to mid-Sept 2018 between the same $\langle s, d \rangle$ pairs as those measured in ARK-AFTER. Of the 8,035 $\langle s, d \rangle$ pairs common to both ARK-BEFORE and ARK-AFTER, we enumerate 6,778 (84.3%) $\langle s, d \rangle$ pairs that contained a common IP hop.

RIPE Atlas (Atlas) had more VPs (10,196 vs 178) than CAIDA’s Ark project, but far fewer usable $\langle s, d \rangle$ pairs (823 vs. 6,778). Although both platforms probe the full set of routed prefixes, Atlas divides its prefix list across 10,196 VPs [52], while Ark divides /24 prefixes across its 178 VPs. Thus, an Ark probe has a larger probability of probing the same prefix. The set of common pairs did not change despite our attempts to augment our dataset with targeted traceroutes between and toward Atlas VPs in Angola and Brazil post-SACS.

3.4 Adding supplementary datasets

We annotated each IP address with its operating AS, router hostname, and geographic information. Using bdrmapIT [42], we mapped 95% of our IPs into ASes. We used `zdns` [28] and `qr` [37] to resolve 35% of those IPs to hostnames. We geolocated IP addresses using the methodology described in §2. We mapped IP hops to their corresponding AS’s country if either: (i) the AS had no customers and NetAcuity [25] geolocated more than 50% of its IP addresses (*i.e.*, those it originates into BGP) to the country, or (ii) 50% of its AS customers geolocated to the same country (by the same process as in (i)). We marked all IP addresses whose hostnames contained geographic hints and updated the city and country they refer to. For cases where we found suboptimal routing (§4.2), we manually cross-checked the geographic hints and the RTT difference to validate the inferred locations. We then identified IXPs at which AC peered pre&post-event, using IXP prefixes in CAIDA’s IXPs dataset [14] as described in §2.

The cable deployment, although entirely within AC’s network, could have triggered a substantial change in the number of BGP paths traversing this AS, since other ASes would have incentive to leverage it, especially those who route traffic between the connected countries/continents. To explore this hypothesis, we analyzed BGP-observed AS paths traversing AC pre and post-SACS. For computation and evaluation of the AS path length, we gathered AS paths (without loops or private ASes) collected from Routeviews [44] and RIS [53] during the first 5 days of Aug and Oct 2018 and included AC (AS37468). To check the post-SACS path stability, we collected new IP paths using Ark and LG servers in AC transit providers and customers between mid-May and end-June 2019.

4 Results and Validation

4.1 Effects on Performance

We quantified the observed RTT changes for packets sent from ASes hosting Ark and Atlas VPs that crossed the cable. We discovered cases of both performance improvements and degradations on paths used pre vs. post-SACS (Fig. 2 and 3). Our results confirm Prior’s claim [50] that the new cable “reduced latency to the Americas substantially, including a reduction from 338ms to 163ms between Cape Town and Miami”. VPs in South America also experienced lower latencies to Africa, with a median RTT decrease of 38% toward all measured African countries. Our findings confirm the drop of latencies from Europe/Africa toward Brazil and those from Brazil to Angola as claimed in [27,36], except for VPs in

North America and Asia, which experienced higher latencies to Brazil (Fig. 3). However, our data does not confirm the claim that latencies to Angola generally experienced an improvement [27,36,57] – on the contrary, paths from VPs in Africa, Asia, and Europe had median latency increases!

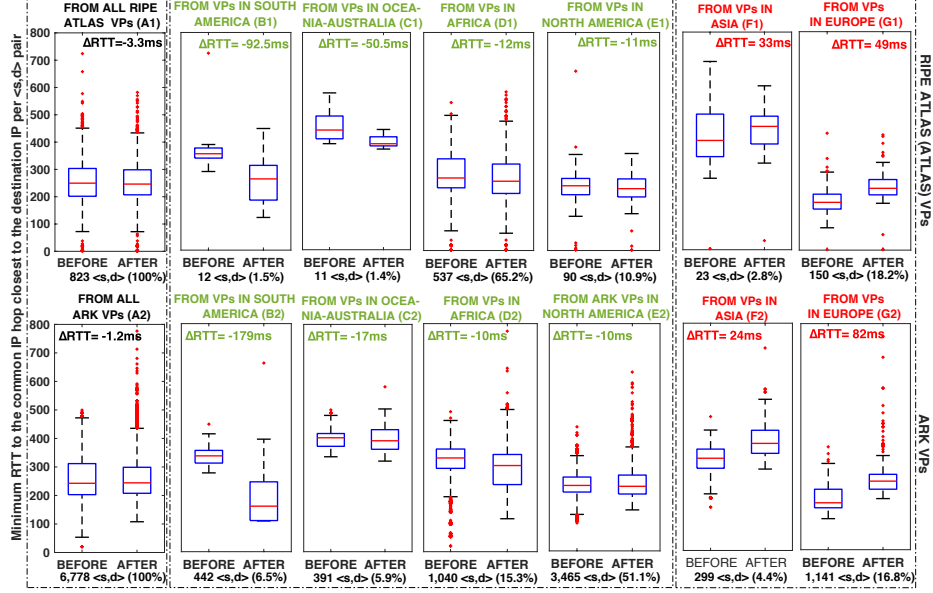


Fig. 2: Boxplots of minimum RTTs from Ark and Atlas VPs to the common IP hops closest to the destination IPs. Sets BEFORE or AFTER are defined in §2. We present $\Delta RTT_{AFTER-BEFORE}$ per sub-figure. RTT changes are similar across platforms. Paths from South America experienced a median RTT decrease of 38%, those from Oceania-Australia, a smaller decrease of 8%, while those from Africa and North-America, roughly 3%. Conversely, paths from Europe and Asia that crossed SACS after its deployment experienced an average RTT increase of 40% and 9%, respectively.

Fig. 2 shows a boxplot of minimum RTT values observed between Ark/Atlas source IP/destination prefix ($\langle s, d \rangle$) pairs. After fetching matching traceroutes (§3.3), half of the 6,778 Ark $\langle s, d \rangle$ pairs were sourced from North America, while most (65.2%) of the 823 Atlas ones were sourced from Africa. For both measurements platforms, at least 16% of the $\langle s, d \rangle$ pairs were sourced from Europe. Fig. 3 presents a heatmap of RTT differences pre vs. post-SACS, for continent/destination country pairs. For statistical significance, we considered only such pairs for which we had at least 20 IP paths. Each box contains the number of observed $\langle s, d \rangle$ pairs (§2.1). The x -axis shows the VP locations, while the y -axis shows the destination prefix countries. The countries on the y -axis are all direct customers of AC. None of Angola’s direct geographic neighbors (Zambia, Zimbabwe, Botswana, Namibia, or Democratic Republic of Congo) are represented on the y -axis. Neither are those neighbors in the 1,034 ASes of AC’s AS customer cone [12,54].

Fig. 2 shows that the Ark and Atlas platforms show similar trends in RTT performance pre to post-SACS per region, as one would expect. In fact, 64% of countries and 89% of $\langle s, d \rangle$ pairs represented in Fig. 3 are already present in the same matrix inferred only from Ark data. Overall, RTT values on IP paths observed by Atlas VPs as crossing SACS are statistically stable (from 249ms to 246ms) with a decrease of the interquartile range (IQR) of 10% (from 102ms to 92ms). The trend for Ark VPs is similar: median RTT drops from 245ms to 243ms, and the IQR drops 18%.

One would expect the greatest performance improvements for VPs in Africa and South America, *i.e.*, close to the cable. Fig. 2B1 and 2B2 show that this is the case for communications from South America crossing SACS. For example, before SACS launch, traffic from Brazil to Angola via AC visited S  o Paulo, London/Lisbon, and Sangano via the WACS cable [59], traversing double the great-circle distance between Brazil and Angola, before reaching Luanda (AO) with an RTT of at least 279ms. The use of SACS dropped this RTT to a low of 108ms. These statistics are consistent with those AC presented in [36].

In contrast, Fig. 2D1 and 2D2 reveal only a slight RTT decrease (10ms *i.e.*, 3%) for VPs in Africa, comparable to that of VPs in North America (Fig. 2E1 and 2E2). While Fig. 3 shows that the most significant RTT drops are on paths from South America to Angola (226ms a 67% drop), South Africa (199ms, a 55% drop), and Nigeria (138ms, a 46% drop), it shows that these are all at least twice the percent drop observed on paths from Africa to Brazil (73ms, a 21% drop). In fact, IP paths from, for instance, Dar-es-Salam (TZ) traversed Mombassa (KE), London (UK), Paris (FR), Amsterdam (NL), Miami (US) to reach Brazil before SACS deployment, and switched to Mombassa (KE), Marseille (FR), Madrid (ES), Lisbon (PT), Sangano (AO), and Brazil after SACS. We inspect these circuitous paths and their causes in  4.2.

Our dataset confirms that, for $\langle s, d \rangle$ pairs from South Africa toward Brazil that benefited from SACS, observed minimum RTTs decreased from 298ms to 116ms (highlighted in [60]). Minimum RTTs decreased 44% for $\langle s, d \rangle$ pairs


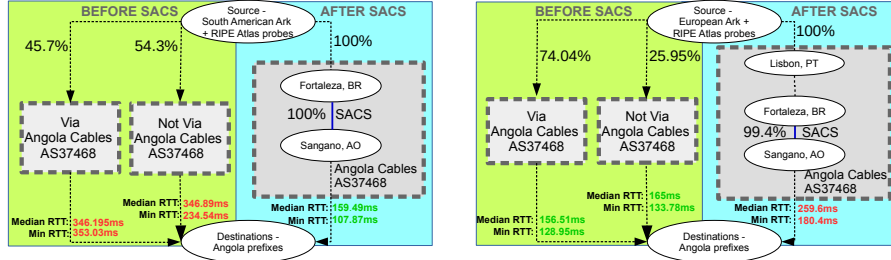


Fig. 3: $\Delta RTT_{AFTER-BEFORE}$ of the medians of minimum RTTs per week pre&post SACS for observed $\langle s, d \rangle$ pairs. We sort the x -axis by the average change per region and the y -axis by ΔRTT for all VPs. **Each cell contains the number observed $\langle s, d \rangle$ pairs**, and is colored according to the corresponding ΔRTT ; a grey cell means data non-available. The highest performance improvements are observed from South America to Angola or South Africa, while the worst degradations are from Africa to Angola or North-America to Brazil.

from Zambia, 35% for those from Nigeria and 3.5% from Ghana toward prefixes in Brazil. The dataset also reveals performance degradations *e.g.*, for RTTs from most VPs in Europe and Asia (Fig. 2G and 2F). From the inspection of performance per continents/countries destination, we learned that the biggest RTT increase occurred for $\langle s, d \rangle$ pairs sourced from Africa to Angola (241ms *i.e.*, 161%), which surprisingly crossed SACS after its launch (Fig. 3). This is followed by cases of paths from North America to Brazil (189ms increase *i.e.*, 123%), Europe to Angola (102ms – 69%), and Africa to China (24%).

4.2 Effects on country paths and transit ASes serving forward paths

We investigated the change in forward paths from South America, Africa, and Europe to Angola. Before using SACS, packets from South America to Angola first traveled to Europe, and then went through the existing WACS cable [59] to Angola (inferred via hostnames that indicate WACS landing points). AC served 46% of $\langle s, d \rangle$ pairs observed by both Ark and Atlas VPs. After SACS, paths for all observed $\langle s, d \rangle$ pairs transited through AC, leveraging SACS for lower latency (Fig. 4A). Fig. 4B shows paths from Europe to Angola, where the forward paths crossed SACS instead of the existing WACS. In this case, the use of SACS increased latency due to higher propagation delay and an increase in the number of transited routers (Fig. 2G1 and 2G2).



(A) Partial AS paths from South America to Angola. Before using SACS, paths between 46% of $\langle s, d \rangle$ pairs crossed Europe and then Angola via AC. SACS provided to all measured $\langle s, d \rangle$ pairs a more direct path between these two continents and improved performance.

(B) Partial AS paths from Europe to Angola. AC was the major transit provider for traffic from Europe to Angola throughout the entire period of study. However, the use of SACS within AC significantly lengthened the physical path, and thus the latency of the forward path.

Fig. 4: Impact of SACS deployment on the set of transit ASes on observed paths going South America to Angola (RTT improvement) and from Europe to Angola (RTT degradation). The white ovals inside AC are part of traceroutes post-SACS we manually geolocated using hints in hostnames.

Fig. 5 illustrates how, after SACS, a high proportion of observed paths for certain continent/destination country pairs followed circuitous paths within AC’s network, crossing the sea multiple times.

We computed the AS-centrality (§2.1) of ASes within the forward paths and inferred the top 3 transit ASes that serve most $\langle s, d \rangle$ pairs (Table 1). After SACS, the same top 2 ASes remained, although the AS-centrality of AC shifted to 90%. However, observed packets routed within AC took a suboptimal route: for 27.2% of $\langle s, d \rangle$ pairs, packets routed within AC via Cape

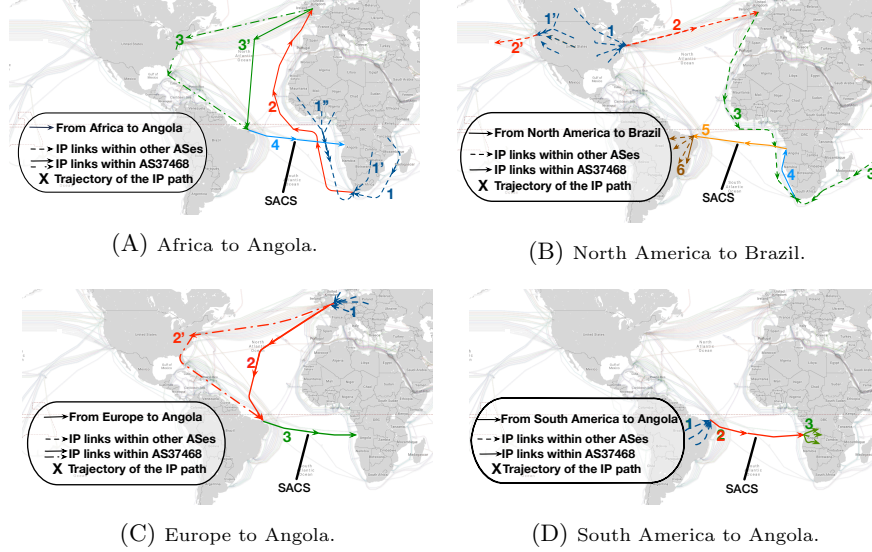


Fig. 5: Examples of suboptimal trajectories followed post-SACS by most paths from Africa to Angola (at least 55%), North America to Brazil (25%), and Europe to Angola (99.3%) *within AC's network (AS37468)* or within other ASes in the paths vs. straightforward trajectory *within AC* or other ASes of most paths from South America to Angola ($\simeq 100\%$), explaining the values of ΔRTT_{A-B} in Fig. 2. We use the same colors to code stages (1, 2, 3, 4, 5, and 6) regardless of the subfigure.

Town/Johannesburg (ZA) traveled a great-circle distance of 13,502km more than before SACS, while for another 55% of $\langle s, d \rangle$ pairs, packets entering AC through London traveled 7,081km more than before SACS. Suboptimal paths from Africa (through Europe, possibly North America, and Brazil) to Angola inducing the RTT increase of Fig. 3 (241ms) post-SACS were either due to suboptimality within AC itself or to neighbors that were routing packets towards AC even though going through SACS was not the shortest route anymore. Fig. 5A depicts how 55% of paths originating in different African countries entered AC either through South Africa, via Europe down to Brazil, and crossed SACS before landing in Angola.

The next largest median RTT increase was for paths from North America to Brazil, which rose 187ms (123%) for observed $\langle s, d \rangle$ pairs of this category. Fig. 5B shows 2 trajectories used by 25% of these paths: from North America, packets crossed Europe or Asia, enter AC PoPs at IXPs in South Africa, then all went to Angola before crossing SACS to Brazil: this proves the existence of a direct link from South Africa to Angola (via WACS), making the suboptimal African paths previously mentioned even more curious. All 3 most-central ASes for the same pairs changed after SACS launch, with a higher AS-centrality and 100% of $\langle s, d \rangle$ pairs were served by AC post-SACS (Table 1).

Paths from Europe to Angola showed a median increase of 102ms (69%). Fig. 5C shows the trajectory of such paths sourcing from Europe and entering

Table 1: Top 3 transit ASes serving $\langle s, d \rangle$ pairs from continents to destination countries. The categories for which we noticed suboptimal routing and RTT increase post-SACS are in italic. Although all our pre-selected paths post-SACS cross SACS, AC may still have an AS-centrality lower than 100%, since the AS-centrality does not account for cases where the AS is either the source or the destination of the AS path.

Category (# $\langle s, d \rangle$)	Before			After		
	CC	AS-cen- trality	Transit AS		AS-cen- trality	CC
<i>From Africa to Angola (201)</i>	AO	66.7%	Angola Cables (AS37468)		90.1%	AO
	ZA	32.3%	Internet Solutions (AS3741)		22.4%	ZA
	BG	20.9%	Sofia Connect (AS47872)	WIOCC-AS (AS37662) IPPLANET (AS12491)	16.4% 16.4%	MU IL
<i>From North America to Brazil (122)</i>	US	44.4%	ATT-Internet4 (AS7018)		100%	AO
	BR	30.1%	NipBr (AS27693)		60.2%	CN
	US	23%	Nitel (AS53828)	Abilene (AS11537)	58.3%	US
<i>From Euro- pe to Angola (705)</i>	AO	62.9%	Angola Cables (AS37468)		78.1%	AO
	BG	18.6%	Sofia-Connect (AS47872)		17.6%	EU
	EU	14.2%	Telianet (AS1299)	TWTC (AS4323)	9.9%	US
<i>From Asia to Brazil (141)</i>	AO	50.3%	Angola Cables (AS37468)		90.1%	AO
	US	28.4%	TATA (AS6453)		31.9%	US
	JP	24.1%	KDDI (AS2516)		26.2%	JP
<i>From South America to Angola (212)</i>	AO	45.7%	Angola Cables (AS37468)		96.2%	AO
	BR	36.8%	Terremark do Brasil (AS28625)	Cilnet (AS28580)	18.4%	BR
	US	36.3%	Cogent (AS174)	CO.PA.CO. (AS27768)	11.8%	PY

AC in Europe before going to Brazil and crossing SACS, on their way to their destinations in Angola. We learned from our dataset that after SACS, 99% of paths went through Fortaleza within AC’s network vs. none before. Since using the WACS cable was an option for AC post-SACS, there was suboptimal routing within AC for this category. Packets routed this way traveled roughly 6,435km more than when they went from London (UK) to Luanda (AO) through WACS. Conversely, the largest median RTT decrease (38%) corresponds to paths from South America to Angola: 99% of observed paths directly traversed SACS when routed within AC, enabling packets to travel a great-circle distance of 6,641km less than before. This case shows that optimal routing within AC’s network can indeed substantially improve end-to-end performance for AS paths it serves.

We saw only a third of such improvement from Africa to Brazil (a drop of 73ms *i.e.*, 21%). Further investigation revealed cases of suboptimal interdomain routing for paths going notably from Mauritius, Ghana, Tanzania, South Africa, or Zambia to Brazil via cities on other continents, which result from the persistent lack of peering among neighboring ASes [29,32,34].

We then used Fig. 3 and Table 1 to check whether SACS introduced new backup IP paths between the regions AC connected. No observed $\langle s, d \rangle$ pairs hinted the existence of paths from South America to Europe/Asia via SACS and Africa. Instead, paths from North America toward destinations in Africa via SACS benefit from an RTT decrease of at least 20ms; SACS could thus play the role of a valid backup path for North American ASes to reach African countries or could be used for load balancing purposes. We also checked whether AC received/routed packets post-SACS through new IXPs. Before the SACS launch, AC was present at public peering points spanning 5 continents [3,4,47]. We observed AC peering at 5 additional IXPs (in UK, US, BR, and RU) post-

SACS for the same set of $\langle s, d \rangle$ pairs, *i.e.*, and expanded interconnection footprint.

4.3 Impact on AS path lengths

From Routeviews and RIPE RIS BGP snapshots of Aug 1st–5th and Oct 1st–5th, 2018 (the months before and after SACS launch), we extracted all AS paths through AC post-SACS (Set AFTER), and all AS paths between the same source AS/destination prefix routed pre-SACS (Set BEFORE). We found 2,115,761 unique AS paths that crossed AC in both snapshots. Since the number of observed AS paths differed in each set and the measurements windows are not strictly identical, we computed the average AS path length per source AS/destination prefix pairs: the percentage of outliers *i.e.*, paths of lengths 10–13 (max) was $\approx 1\%$. We noticed the AS path length distribution shifted, with AS paths of length 2–7 generally increasing, reflecting the fact that more neighboring ASes preferred AS paths via AC after the SACS launch (Fig. 6). Interestingly, AC apparently announced many paths to prefixes owned by multiple ASes 2–3 months before the SACS launch [22], perhaps preparing for the launch.




Fig. 6: Distribution of the length of AS paths between same source AS/destination prefix pairs served via AC (AS37468) pre&post SACS, showing the increase of paths of lengths 2–7.

4.4 Validation with the ISP

In Jul 2019, we successfully contacted AC and were able to validate the inferred set of SACS link IPs and their respective locations. AC distinguished cases where the anomalous routes occurred outside their network, and tromboning occurred due to lack of local peering, or where neighbor ASes were circuitously routing traffic toward AC after SACS. During our exchange with them, we took subsequent measurements that showed some AC neighbors had modified their routing configurations in ways that improved performance. Although AC did not validate cases of suboptimal routing within their network, most observed IP paths (from North America/Asia to Brazil or Europe/Asia to Angola) switched to more optimal paths after our conversation. AC also explained that internal link failures could account for the performance degradations. For example, if the MONET cable [3,59] (which AC’s router in Miami crosses to reach Fortaleza) becomes unavailable, the router may re-route traffic through London. They also noted that no customers had complained, so if there were any suboptimal routing, it was unlikely to be affecting any routes that carried any traffic. That said, we found that a few ($\approx 4\%$) $\langle s, d \rangle$ pairs used remarkably suboptimal paths as late as Jul 12, 2019, *e.g.*, from Africa to prefixes in Angola served via Europe and Brazil or those from North America to Angola routed by AC via SACS and Lisbon. Finally, AC informed us that most traffic crossing SACS through AC goes from either South America to Angola or South Africa to Brazil, cases where our results show a pronounced decrease post-SACS (§4.1).

4.5 Potential root causes of suboptimal routing

We confirmed the occurrence of the routing suboptimalities described in this paper using two measurements platforms that revealed similar trends per region. We tried to obtain insights from the ISP operating the cable (AC) into potential causes, without success. We conjecture that these suboptimalities derived from multiple causes (potentially concurrent): *(i)* misconfigurations of either the Internal or External Gateway Protocol (IGP/EGP), due to typos, errors, etc, [10] *(ii)* slow IGP or EGP convergence [38], *(iii)* some ASes routing packets through AC although it is not the optimal path to the destination, *(iv)* the persistent lack of peering among local ASes in Africa (despite ongoing efforts for more local interconnections) [30] and frequent use of default routes via international transit providers in developing regions.

5 Conclusion

It is generally assumed that deployment of undersea cables between continents improves performance, at least for paths that cross a cable once it is deployed. We present a reproducible scientific method by third-parties to investigate the performance impact of a cable deployment, and we apply it to the case of the first South-Atlantic cable, connecting Brazil to Angola. We used traceroute and BGP data from global measurement platforms, and geolocation data and inferences to find that this new cable had at least initially reduced RTTs asymmetrically: the median RTT decrease from Africa to Brazil was roughly a third of that from South America to Angola (226ms). More surprising is that latency statistics to/from other regions of the world, including paths within Africa, increased post-activation due to circuitous IP paths that suboptimally crossed continents and the cable. We uncovered other potential sources of suboptimality: slow BGP route updates/lack of traffic engineering after a major event occurring in a neighboring AS, and problematic intra-domain routing within a single network. Our results suggest ways operators can avoid suboptimal routing post-activation of cables in the future: *(i)* informing BGP neighbors of the launch to allow time for appropriate changes in advance; *(ii)* ensuring optimal iBGP configurations post-activation, not only for pairs of ASes/countries expected to route most traffic through the cable, but also for served intra-regional and cross-regional traffic; and *(iii)* collaborate with measurements platforms or research institutions to verify path optimality. Our methodology is general enough to apply to other cable deployments, as well as cable failures, and contributes to a toolbox to support further scientific study of the global submarine cable network [8,9]. We share our code [31] to promote reproducibility and extension of our work.

Acknowledgment

We thank the anonymous reviewers and our shepherd, Fabian Bustamante, for their insightful comments. We also thank Angola Cables, especially their IP services department, for their cooperation, despite their tight schedule and Stephen

Strowes for the introductions. This research was supported by the National Science Foundation (NSF) grant OAC-1724853.

Bibliography

- [1] Submarine Cable Networks. <https://www.submarinenetworks.com/en/> (Jan 2020)
- [2] Subsea World News. <https://subseaworldnews.com> (Jan 2020)
- [3] Angola Cables: Angola Cables Network. <https://www.angolacables.co.ao/en/network/> (Jan 2020)
- [4] Angola Cables: SACS. <https://sacs.angolacables.co.ao> (Jan 2020)
- [5] Angonix: Angonix – BIRD Looking Glass. <http://lg.angonix.net/> (Jan 2019)
- [6] Augustin, B., Cuvellier, X., Orgogozo, B., Viger, F., Friedman, T., Latapy, M., Magnien, C., Teixeira, R.: Avoiding Traceroute Anomalies with Paris Traceroute. In: Proceedings of the ACM SIGCOMM Internet Measurement Conference (IMC). pp. 153–158. ACM (Oct 2006)
- [7] Berenguer, S.S., Carisimo, E., Alvarez-Hamelin, J.I., Pintor, F.V.: Hidden internet topologies info: Truth or myth? In: Proceedings of the 2016 workshop on Fostering Latin-American Research in Data Communication Networks. pp. 4–6 (2016)
- [8] Bischof, Z.S., Fontugne, R., Bustamante, F.E.: Submarine Cables and Internet Resiliency. https://www.iiij.ad.jp/en/dev/iir/pdf/iir_vol41_focus2_EN.pdf (2018)
- [9] Bischof, Z.S., Fontugne, R., Bustamante, F.E.: Untangling the world-wide mesh of undersea cables. In: Proceedings of the 17th ACM Workshop on Hot Topics in Networks (HotNets). pp. 78–84. ACM (2018)
- [10] Bischof, Z.S., Rula, J.P., Bustamante, F.E.: In and out of Cuba: Characterizing Cuba’s connectivity. In: Proceedings of the 2015 Internet Measurement Conference. pp. 487–493 (2015)
- [11] CAIDA: Archipelago (Ark) Measurement Infrastructure. <https://www.caida.org/projects/ark/> (Jan 2020)
- [12] CAIDA: AS Rank. <http://as-rank.caida.org> (Jan 2020)
- [13] CAIDA: AS Relationships. <http://www.caida.org/data/as-relationships> (Jan 2020)
- [14] CAIDA: CAIDA Internet eXchange Point (IXP) Dataset. <http://www.caida.org/data/ixps/> (Jan 2020)
- [15] CAIDA: Henya. <https://www.caida.org/tools/utilities/henya/> (Jan 2020)
- [16] CAIDA: MIDAR. <http://www.caida.org/tools/measurement/midar/> (Jan 2020)

- [17] CAIDA: The IPv4 Routed /24 Topology Dataset. http://www.caida.org/data/active/ipv4_routed_24_topology_dataset.xml (Jan 2020)
- [18] CAIDA: Vela. <https://vela.caida.org> (Jan 2020)
- [19] CAIDA: Velasq. <https://www.caida.org/projects/ark/vela/aliasq-api/> (Jan 2020)
- [20] Center for Applied Internet Data Analysis (CAIDA): Macroscopic Internet Topology Data Kit (ITDK). <http://www.caida.org/data/internet-topology-data-kit/> (Jan 2019)
- [21] Chang, R., Chan, E., Li, W., Fok, W., Luo, X.: Could Ash Cloud or Deep-Sea Current Overwhelm the Internet? In: Proceedings of USENIX Workshop on Hot Topics in System Dependability (HotDep) (Oct 2010)
- [22] Chung, T., Aben, E., Bruijnzeels, T., Chandrasekaran, B., Choffnes, D., Levin, D., Maggs, B.M., Mislove, A., Rijswijk-Deij, R.v., Rula, J., et al.: RPKI is Coming of Age: A Longitudinal Study of RPKI Deployment and Invalid Route Origins. In: Proceedings of the Internet Measurement Conference. pp. 406–419 (2019)
- [23] Clark, B.: Undersea Cables and the Future of Submarine Competition. Bulletin of the Atomic Scientists **72**(4), 234–237 (2016)
- [24] Dawn-Hiscox, T.: Angola Cables lights up World’s First Submarine Cable linking Africa to the Americas. <https://www.datacenterdynamics.com/news/angola-cables-lights-up-worlds-first-submarine-cable-linking-africa-to-the-americas/> (Sep 2019)
- [25] Digital Element: Netacuity. http://www.digital-element.net/ip_intelligence/ip_intelligence.html (Jan 2020)
- [26] Doug, M.: First Subsea Cable Across South Atlantic Activated. <https://internetintel.oracle.com/blog-single.html?id=First+Subsea+Cable+Across+South+Atlantic+Activated> (Sep 2018)
- [27] Doug, M., Darwin, C., Humberto, G.: South Atlantic Cable System The Impact on the Internet LACNIC 30 - Lightning Talk. https://www.lacnic.net/innovaportal/file/3209/1/sacs_lightning_talk_lacnic30.pdf (Sep 2018)
- [28] Durumeric, Z., Pearce, P.: Fast CLI DNS Lookup Tool. <https://github.com/zmap/zdns> (Sep 2019)
- [29] Fanou, R., Francois, P., Aben, E.: On the Diversity of Interdomain Routing in Africa. In: International Conference on Passive and Active Network Measurement (PAM) (Mar 2015)
- [30] Fanou, R., Francois, P., Aben, E., Mwangi, M., Goburdhan, N., Valera, F.: Four Years Tracking Unrevealed Topological Changes in the African Interdomain. Computer Communications **106**, 117–135 (Jul 2017)
- [31] Fanou, R., Huffaker, B., Mok, R.K., Claffy, K.: Submarine Cable Impact Analysis: Public Codebase. <https://github.com/CAIDA/submarine-cable-impact-analysis-public> (Jan 2020)
- [32] Formoso, A., Chavula, J., Phokeer, A., Sathiaselalan, A., Tyson, G.: Deep Diving into Africa’s Inter-Country Latencies. In: IEEE Conference on Computer Communications (INFOCOM) (2018)

- [33] Gharaibeh, M., Shah, A., Huffaker, B., Zhang, H., Ensafi, R., Papadopoulos, C.: A Look at Router Geolocation in Public and Commercial Databases. In: Proceedings of the ACM Internet Measurement Conference (IMC) (2017)
- [34] Gupta, A., Calder, M., Feamster, N., Chetty, M., Calandro, E., Katz-Bassett, E.: Peering at the Internet’s Frontier: A First Look at ISP Interconnectivity in Africa. In: International Conference on Passive and Active Network Measurement (PAM) (Mar 2014)
- [35] Huffaker, B., Fomenkov, M., Claffy, K.: Geocompare: A Comparison of Public and Commercial Geolocation Databases. Proc. NMMC pp. 1–12 (2011)
- [36] Humberto, G.: South Atlantic Cable System - SACS The Impact on the Internet WTR POP-BA/RNP 2018 - Lightning Talk. <https://wtr.pop-ba.rnp.br/2018/files/apresentacoes/10-WTR2018-LT02-AngolaCables-HumbertoGaliza.pdf> (2018)
- [37] Hyun, Y.: Dolphin: Bulk DNS Resolution Tool. http://www.caida.org/publications/presentations/2014/dolphin_dhs/dolphin_dhs.pdf (June 2014)
- [38] Labovitz, C., Ahuja, A., Bose, A., Jahanian, F.: Delayed Internet routing convergence. ACM SIGCOMM Computer Communication Review **30**(4), 175–187 (2000)
- [39] Lodhi, A., Larson, N., Dhamdhere, A., Dovrolis, C., Claffy, K.: Using PeeringDB to Understand the Peering Ecosystem. ACM SIGCOMM Computer Communication Review **44**(2), 20–27 (2014)
- [40] Luckie, M., Huffaker, B., Dhamdhere, A., Giotsas, V., Claffy, K.: As relationships, customer cones, and validation. In: Proceedings of the 2013 conference on Internet measurement conference. pp. 243–256 (2013)
- [41] Mahlknecht, G.: Greg Mahlknecht’s Cable Map. https://cablemap.info/_default.aspx (Jan 2020)
- [42] Marder, A., Luckie, M., Dhamdhere, A., Huffaker, B., Smith, J.M., et al.: Pushing the Boundaries with bdrmapIT: Mapping Router Ownership at Internet Scale. In: Proceedings of the ACM Internet Measurement Conference (IMC) (2018)
- [43] MaxMind: GeoIP. <https://dev.maxmind.com/geoip/geoip2/geoip2-city-country-csv-databases/> (Aug 2019)
- [44] Mayer, D.: University of Oregon RouteViews Archive Project. routeviews.org (Jan 2020)
- [45] Padmanabhan, R., Schulman, A., Levin, D., Spring, N.: Residential Links Under the Weather. In: ACM (ed.) In Proceedings of the ACM Special Interest Group on Data Communication (SIGCOMM). pp. 145–158 (2019)
- [46] PeeringDB: <https://www.peeringdb.com/> (Apr 2019)
- [47] PeeringDB: PeeringDB: Angola Cables. <https://www.peeringdb.com/net/4894> (Jan 2020)
- [48] Phil, E.: A Map of all the Underwater Cables that Connect the Internet. <https://www.vox.com/2015/3/13/8204655/submarine-cables-internet> (Nov 2015)
- [49] Poese, I., Uhlig, S., Kaafar, M.A., Donnet, B., Gueye, B.: IP Geolocation Databases: Unreliable? ACM SIGCOMM Computer Communication Review **4**(2), 53–56 (2011)

- [50] Prior, B.: Teraco data centres will benefit from SACS cable. <https://mybroadband.co.za/news/cloud-hosting/284682-teraco-data-centres-will-benefit-from-sacs-cable.html> (Nov 2018)
- [51] RIPE NCC: Global RIPE Atlas Network Coverage. <https://atlas.ripe.net/results/maps/network-coverage/> (Jan 2020)
- [52] RIPE NCC: RIPE Atlas: Built-in Measurements. <https://atlas.ripe.net/docs/built-in/> (Jan 2020)
- [53] RIPE NCC: RIPE RIS. <https://www.ripe.net/analyse/internet-measurements/routing-information-service-ris/> (Jan 2020)
- [54] RIPE NCC: RIPE Stats. <https://stat.ripe.net/> (Jan 2020)
- [55] Shah, A., Fontugne, R., Papadopoulos, C.: Towards Characterizing International Routing Detours. In: Proceedings of the ACM Asian Internet Engineering Conference (AINTEC) (2016)
- [56] Submarine Cable Networks: SACS. <https://www.submarinenetworks.com/systems/brazil-africa/sacs> (Mar 2018)
- [57] Subsea World News: South Atlantic Cable System Launched in Angola. <https://subseaworldnews.com/2017/08/09/south-atlantic-cable-system-launched-in-angola/> (Aug 2019)
- [58] Telegeography: Submarine cable frequently asked questions. <https://www2.telegeography.com/submarine-cable-faqs-frequently-asked-questions> (Jan 2020)
- [59] Telegeography: Telegeography Submarine Cable Map. <https://www.submarinecablemap.com/> (Jan 2020)
- [60] Vermeulen, J.: From Brazil to South Africa and back in 98ms. <https://mybroadband.co.za/news/broadband/303574-from-brazil-to-south-africa-and-back-in-98ms.html> (Apr 2019)