On The Diversity of Interdomain Routing in Africa

Rodérick Fanou[1,2], Pierre Francois[2], and Emile Aben[3]

[1] IMDEA Networks Institute, Madrid, Spain
{roderick.fanou,pierre.francois}@imdea.org
[2] Universidad Carlos III de Madrid (UC3M), Madrid, Spain
[3] RIPE NCC, Amsterdam, Netherlands
emile.aben@ripe.net

Abstract. With IP networking booming in Africa, promotion of BGP peering in the region emerge, and changes in the transit behavior of ISPs serving Africa are expected. However, little is known about the IP transit topology currently forming the African Internet. Enhancing the RIPE Atlas infrastructure, we evaluate the topology interconnecting ISPs based on the continent. We reveal a variety of ISP transit habits, depending on a range of factors such as the official language or the business profile of the ISP. We highlight the emergence of IXPs in Africa, evaluating its impact on end-to-end connectivity. Our results however emphasize the remaining dominance of ISPs based outside Africa, for the provision of intra-continental paths. We study the impact of this aspect on AS path length and end-to-end delay. Such results illustrate that performing measurements from a broad, diversified, range of vantage points is necessary to assess interdomain routing on the continent.

Keywords: RIPE Atlas, IP transit, African Internet, IXP

1 Introduction

Despite major investments in submarine and terrestrial cable deployments in Africa, Internet access is still perceived as of low quality, with high latency and low bandwidth [1,2]. According a study of the African Union [3], Africa spends between US $400 millions and $600 millions per year in transit fees for intra-African traffic. Initiatives such as the African Internet eXchange System have thus been launched to promote the creation of IXPs and regional carriers, and improve the fragmented status of the IP infrastructure [4]. It provides an enabling environment for cross-border interconnection to thrive and become more competitive to reduce transit costs paid by Africa for intra-African traffic exchange. Meanwhile, caches and peering points have been deployed by CDNs (Google, Akamai, Cloudfare, etc) in the region proving the capacities to offload traffic from the expensive transit links. However, little is known about the current state of the Internet topology in Africa, due to its low representation in existing measurement projects. Obtaining relevant topological data, especially for access-to-access interconnection is thus essential to understand its current state and observe its foreseen evolution. Recent work focusing on Africa, such as [5,2,6], relied on a very limited set of vantage points, and had different focuses, as explained later.

In this paper, we set as key milestone to obtain an interdomain map that is not biased towards the South African perspective. To this end, we met ISPs in Benin, Burkina Faso, Congo, Ghana, Ivory Coast, Mauritania, Morocco, Niger, Nigeria, Senegal, and Togo, to deploy RIPE Atlas probes within their networks. We complemented our set of probes with those deployed by RIPE Atlas Ambassadors. To obtain insights on the evolution of the peering ecosystem, our measurement campaign covers 6 months, monitoring end-to-end paths among v4 and v6 probes scattered on the continent [7]. We discover a large variety of ISP transit habits, notably correlated with the location, the official language, and the monetary union of the country in which operate the ISP. Our results illustrate that, with the exception of ISPs based in South Africa, the provision of intra-continental paths is dominated by ISPs based outside Africa, while South Africa is being adopted as a hub for East-West African communications. We study the impact of such aspects on the end-to-end delay between ISPs, notably among networks based in the same country. Finally, we illustrate the benefits of new IXPs with respect to end-to-end delay.

The remainder of this paper is structured as follows. Section II discusses related work. Section III presents our methodology, while Section IV exposes our results. Section V concludes the paper and describes our plans for further research.

2 Related work

An extensive amount of research has been carried out on the discovery of the Internet topology, both at the router level and the AS level [8,9,10]. Archipelago has the goal of reducing the efforts needed to develop and deploy sophisticated large-scale measurements [11]. Of its 94 monitors, only 5 are deployed in Africa. For this study, a larger deployed base of vantage points was needed. Similarly, although the PingER project [12] involves 46 African countries, only Burkina Faso and South Africa host a monitoring site, preventing us from doing large scale end-to-end measurements.

Gilmore *et al.* mapped, in [5], both the router level and AS level graphs of intra-African traffic. Traceroutes from South Africa towards all the IPs allocated by AFRINIC were performed for a week. To improve their results, they extracted from the RIBs of routers in the South African Tertiary Education Network, the links among ASes registered in AFRINIC and those towards their direct neighbouring ASes. As a consequence, they obtained one way paths from which they infered a tree, of which South Africa is the root. They acknowledged that the link density might look different if the traceroute probes were sent out from other countries in Africa.

Recently, Gupta *et al.* investigated Internet connectivity between Kenya, Tunisia, and South Africa [6], by performing traceroutes from access networks to sites hosting popular content. They noticed that 66.8% of paths from their vantage points towards Google cache servers located in Africa leave the continent. They generalized, from their results, the nature of intra-domain interconnectivity on the continent. Nevertheless, broadband access networks in those countries are more developed in comparison with most African countries, so that, as acknowledged by the authors, the obtained dataset may not reflect the nature of the paths in other countries.

In contrast, our study presents discoveries of the Internet infrastructure based on measurements performed from access to access networks, as we aim at studying how Africans communicate with one another. We perform these measurements among a large variety of networks, and for a long enough period of time to study the dynamics of the African Internet topology. The paths in our dataset are typically not seen in

RouteViews, RIS and PCH datasets, as these do not host monitors in the studied regions. We also show that ISP transit and peering habits vary throughout the continent. By exhibiting newly established IXPs located in Africa, as well as the use of other ones, we show that ISPs do peer now in Africa, illustrating the first benefits of the initiatives promoting peering. An exploration of how these measurements in Africa compare to measurements in other regions is left to future work.

3 Methodology

3.1 Data Collection

Multiple challenges influenced our choice for the measurement infrastructure. First, whereas network operators are reticent to the intrusion of foreign devices, for legitimate security and privacy reasons, we had to find a relevant number of hosting locations for the measurement devices. Such devices have to be robust, as power outages and surges frequently occur in the studied countries. Finally, we preferred an open measurement infrastructure, as we wanted to provide means for network operators and researchers to further study African networking. We chose the RIPE Atlas measurement platform, which consists of over 7400 deployed worldwide [7]. These devices are free, secure, and require no maintenance.

In June 2013, Africa only hosted a few active RIPE Atlas devices, with almost no deployment in the West. To improve the situation, we deployed 21 RIPE Atlas probes in 15 ISPs networks covering 11 countries, focusing on that region. These devices are hosted by either ISPs, universities, or home networks. None of them are behind a wireless access link, to reduce the impact of last mile latency on our results. Collaborating institutions such as AFRINIC and ISOC also deployed a considerable amount of probes in the Southern and Eastern regions, which we used in this study.

We used paris-traceroute for all our measurements to discover as many paths as possible, and not suffer from inconsistent results caused by load balancing, as happens with classic traceroute [13]. Probes performed traceroutes with 16 different paris_id defaults. We used the UDP-traceroute to reduce the potential bias caused by differentiated traffic handling of ICMP packets [14]. We conducted 3 measurement campaigns. During the first one, we performed traceroutes among all probes located in Africa, with a period of 3 hours, from November 30, 2013 to April 06, 2014. It results in total in 675,421 traceroute outputs. Second, we issued v4 and v6 traceroutes, at the same frequency, focusing on countries hosting IPv6 enabled probes, from June 01, 2014 to August 01, 2014, in order to compare v4 and v6 routing. It results in total in 408,383 v4 and 21,744 v6 traceroute outputs. Finally, to highlight the launch of the Serekunda IXP in GM[4], we performed hourly, during the second week of August 2014, 3,161 traceroutes among all RIPE Atlas probes in GM, publicly available in a Technical Report [15].

An essential step is to undertake an in-depth sanity-check on the raw data to only consider the valid traceroute outputs during our analysis. Before this filtering process, our raw data involved 214 probes hosted in 90 ASes operating in 32 African countries. The geographical and networking spread of the used probes are available in [15]. As

[4] In this paper, we refer to countries using ISO 2-letter country codes, that we list in [15].

for the granularity of our results, the percentage of ASes allocated by AFRINIC [16] covered per country is 21.7% on average [15]. With this dataset, we first map IPs into Country Codes (CCs) in order to infer the set of countries traversed by the packets during each traceroute. Second, we map IPs into ASes to infer the ASes sequences.

3.2 Data Analysis

3.2.1 IP to Country Code Mapping and Validation. Geolocation is said to be of poor quality, especially for IPs located in Africa [17]. To geographically locate the 8,328 v4 and 465 v6 public IPs found in the traceroute data as accurately as possible, we thus analyzed 6 public databases (DBs), that we cross-correlated with delay measurements, as explained in this section. We used OpenIPMap (*OIM*) [18], Reverse DNS lookups (*RDNS*), MaxMind GeoIP2City (*MM*) [19], Team Cymru (*TC*) [20], the AFRINIC DB (*AF*) [16], and Whois (*Whois*).

When all databases providing an entry for an IP returns the same CC, we retain it for that IP. When DBs are inconsistent for an entry, we use a latency-based method to tie-break among them. We ping each IP from up to 10 random RIPE Atlas probes hosted in each country returned by the DBs[5]. We compute the minimum delay recorded per possible country, and assume that the IP is located in the country for which the minimum delay is the lowest. We compare in Table 1 the entries of the selected DBs. The coverage column (Cov.) is the percentage of addresses of our dataset for which the DB provided a valid country field. Trust is the percentage of IPs for which the DB entry is equal to the country that we finally selected for that IP.

	IPv4 entries		IPv6 entries			IPv4 entries		IPv6 entries	
DB	**Cov.**	**Trust**	**Cov.**	**Trust**	**DB**	**Cov.**	**Trust**	**Cov.**	**Trust**
OIM	26%	93.8%	30.1%	92.8%	*TC*	86.7%	71%	99.1%	79.4%
RDNS	56.7%	88.8%	46.7%	78.5%	*AF*	36.2%	93%	56.7%	83.7%
MM	83.9%	74%	99.1%	71.4%	*Whois*	85.6%	68%	43.2%	67.7%

Table 1: *Comparison of Geolocation databases*

5,430 v4 IPs (resp. 292 v6 IPs) out of the 8,328 v4 (resp. 465 v6) have an identical CC mapping among all DBs for which an entry was available. Our delay-based tie-breaking approach was used to geolocate the rest of the IPs that responded to pings (81.2% of IPv4, 92.2% of IPv6). That is for 2,406 v4 IPs (resp. 164 v6 IPs), the ping technique allows us to deduce the country. At the end of this process, 94.1% v4 IPs (98.1% v6 IPs) of our dataset are associated with a location. With the obtained geolocation data, we can compute the country path corresponding to the IP path of each traceroute output [17].

3.2.2 IP to AS Lookup and Raw Data Sanity Check. We map, using *TC*, public IP addresses of our traceroute data into ASes, with the following filtering procedure: We first keep traceroutes for which the obtained AS Sequence contains source and destination ASes corresponding to the ASes which are known to host the probes. If it is not the case, we check if the first AS on the path is a known direct upstream of the source and

[5] The raw data for these delay measurements can be found in [15].

the last AS on the path is a known direct upstream of the destination, as observed in the previous set of traceroutes. If these checks succeed, we keep the traceroute as well. Note that we only use this second set of inferred AS sequences for AS Path analysis, and excluded them from our RTT analysis.

To assess the accuracy of the inferred AS Paths, we keep track of intermediate traceroute hops for which the IP has no entry in *TC*, or for which we did not receive a reply [10]. We respectively refer to them as *unresolved* and *unknown* ASes.

We then compress AS Paths into AS Sequences. *Unresolved* or *unknown* hops found between two resolved hops of the same given AS, are considered as belonging to that AS. Consecutive equal AS numbers are compressed into a single AS hop. We only infer an edge between two ASes if there are no *unresolved* or *unknown* hops in the IP path, and if both ASes are consecutive in the AS sequence.

In our first campaign, we identified 4,648 traceroutes with inferred AS path loops. The top 3 ASes in those paths with loops were AS3356 (Level3, US), AS37282 (MainOne, NG) and AS37054 (Data Telecom Service, MG) at respectively 32%, 15% and 11%. Similar results were found in our second campaign. These paths are a small fraction of the total dataset, so we decided to filter them out. Including these paths in our results is part of our ongoing work.

By the end of this raw data cleaning method, we retained, for the first campaign, 87.81% of v4 traceroutes. For the second one, we retained 97.27% of v4 traceroutes and 90.11% of v6 traceroutes. For the last campaign, we retained 86.93% of v4 traceroutes outputs. The corresponding total numbers are listed in section 3.1. The dataset resulting from this filtering process comprises paris-traceroutes outputs from 181 probes located in 30 African countries, hosted in 90 ASes.

Finally, we estimate the RTT between the source and the destination AS as the difference between the RTT from the source probe to the ingress point of the destination AS, and the RTT from the source probe to the egress point of the source AS.

4 Results

In this section, we first discuss the biases of our dataset, and compare it with the view of the African topology that can be made from public BGP data. We then investigate the dynamics of the observed paths. Next, we discuss the length of the AS Paths, highlighting different trends among the studied regions. We then illustrate the impact of the intercontinental aspect of paths on the RTTs among African ISPs. We finally detect new peering links and IXPs in Africa to shed light on the progress made by some operators towards localizing interdomain routing. We provide a detailed analysis of the coverage of the IP ranges per country, path dynamics, the AS-Centrality of our dataset in [15].

4.1 Dataset limitations and public BGP data

We acknowledge that not all the probes were deployed at the beginning of our initial measurement campaign. The constant evolution of the RIPE Atlas infrastructure on the continent leads us to daily add new probes to the set of probes that we use. As of today, our dataset involves 7.2% of the ASes allocated by AFRINIC. Shortcomings of IP to AS mapping also have to be considered. For instance, 40.6% (resp. 35.9%) of the unique v4 (resp. v6) AS paths either contain one *unknown* or an *unresolved* AS.

We extracted from Routeviews, RIPE RIS, and PCH, all the paths containing one of the 90 ASes hosting a probe (2,258,692 v4 and 840,180 v6 paths), using all available data for 2013 and 2014 [21,22,23]. We compare these paths with our set of (2,529 v4 and 91 v6) paths containing no *unknown* or *unresolved* ASes. As most of the routes collectors are hosted outside the continent, our dataset is more precise when it comes to end-to-end African paths. Among the 960 v4 (resp. 63 v6) AS adjacencies that we inferred from the discovered paths, 733 v4 (resp. 35 v6) are not visible in these public datasets. Note that most of the AS adjacencies found in both datasets are among ASes based outside the continent. Quite intuitively, entire African AS paths - 2,519 v4 (resp. 79 v6) - are not visible in RouteViews, RIS, and PCH either. The only AS paths found in both datasets actually belong to the set of paths measured from AS3741 (Internet Solutions, ZA); this ISP hosts a RouteViews collector.

4.2 Path dynamics

Considering the data of our first two campaigns, we identify all unique AS paths for each pair of monitored ASes, and compute their frequency of observation. In the rest of this paper, we refer to the path among two ASes which has the highest frequency of observation, as the preferred path for that AS pair. About 72.6% (resp. 82%) of the v4 (resp. v6) preferred paths have been used with a frequency higher than 90%. Only 4% of the v4 AS pairs have used their preferred paths at a frequency lower than 50%, whereas no v6 preferred path has a frequency of usage lower than 50%.

Some outliers were found in this analysis. Among them, Isocel Telecom (BJ) was showing 42 different AS Paths towards Onatel/FasoNet-AS (BF), and 22 paths were observed in the opposite direction. Link, node failures, and flapping could be listed as the possible reasons of such changes. However, we validated these results by visiting Isocel Telecom, and discovered that this ISP is constantly performing interdomain traffic engineering in order to offer the best possible QoS to its customers. As Onatel is doing the same, a large number of paths were explored between these two ASes.

4.3 AS path length distribution

Based on our first two measurement campaigns, we study the distribution of the length of AS Sequences among pairs of ASes. We notably take a perspective focused on West Africa and South Africa, to highlight differentiated trends. We also carry out a specific analysis for pairs of ASes located within the same country. Note that we remove *unknown* and *unresolved* ASes from paths used for plotting graphs of Figure 1, except for those corresponding to paths length distribution within countries. Thus, the AS paths in those cases could be even longer than what is presented.

On Figure 1(d), we show the AS path length distribution for all the paths of the dataset. Since ASes in West Africa are based in geographically collocated countries, one could guess that paths would be shorter. However, based on the specific view provided in 1(a), we discover unusually long AS paths in West African communications. Figure 1(b) and 1(f) highlight that short paths tend to be found in Southern Africa, and precisely

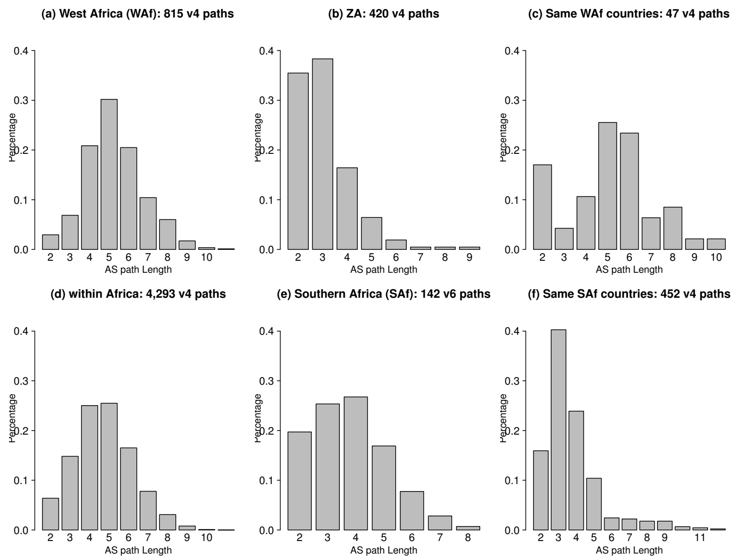

Fig. 1: *Path length distributions for all paths and for some African regions*

in ZA. Paths among ASes operating in the same country (Figure 1(c) and (f)), are much shorter in South Africa than in West Africa. IPv6 AS paths, all observed in *SAf*, tend to be short, reflecting similar peering and localized transit habits as for v4 in the region. These observations confirm that focusing solely on measurements from ZA does not provide a representative sample of Internet paths characteristics for the rest of Africa.

4.4 Trends in African Interdomain Routing

We now study the role of transit played by each ISP found in our data. To this end, we use the concept of AS-centrality of an AS, defined as the percentage of paths containing that AS [17], but for which that AS is neither the source nor the destination. Note that AS-centrality is not equivalent to the betweenness centrality of ASes on the AS graph. We only account for presence within AS paths among pairs of ASes, radically diverging from betweenness centrality in the AS graph. We then define the concept of "joint AS-centrality", which captures the centrality of tuples of ASes present together on AS paths.

We classify the 164 ASes of our dataset into 5 categories, depending on their region of operation: *WAf* (for ASes based in West Africa), *SAf* (Southern Africa), *EAf* (East Africa), *RAf* (ASes operating in Africa but not in any of the previous regions), and *Int* - Intercontinental - (all ASes based outside the continent). Note that we find 61 *Int* Ases.

The 4 most central ASes in our view of the v4 African interdomain topology are all intercontinental ones: Level3 (US) with 23.4% of the 4,293 AS paths, TATA (US) with 22%, Cogent (US), 13.6%, and Orange (FR), 12%. 65.2% of the AS pairs were served using at least one of these 4 ASes. The most central African AS, Internet Solutions (ZA) has an AS-centrality of 11.6%. In contrast, Orange is the dominating ISP when it comes to paths among ASes in the *WAf* category, with an AS-centrality of 37.8%, while

TATA and Level3 respectively own 32% and 26.1% of the market share. We notice that a relevant percentage of paths (18.9%) connecting *WAf* ASes transit via MTN (ZA). The most central local AS is MainOne (NG), found on 17.2% of the paths.

The reliance on *Int* transit providers is considerably lower within the Southern African region; the top 2 ASes remain Level3 with 20.5% and TATA (15.3%), but Internet Solutions (ZA), SAIX-NET, a private IXP owned by Telkom SA (ZA), and MWEB (ZA) follow with respectively 15%, 12.9%, and 11.2%. *SAf* ASes appear to benefit from diversity in their transit offerings, and resort a lot to peering, as no ISP was found to completely dominate transit in the region. Note that the reliance of *SAf* ASes on ISPs based on other African regions is insignificant.

Some ASes which are not relevant for v4 routing, show a high AS-centrality when it comes to v6. The top 3 ASes in v6 are Hurricane Electric (US) with 23.9%, TENET (ZA) with 22.5%, and Liquid Telecom (GB) with 21.1%. They are followed by IXPs AS5459 (LINX-AS, GB) and AS1200 (AMSIX, NL) traversed respectively by 19.7% and 14.7% of the paths while Level3 and TATA are present on only 9.8 % and 9.1%.

One can observe a diversity of transit trends based on technico-economical factors. In Figure 2, we present the centrality of Orange, TATA, and Level3, discussing whether these ASes jointly serve a path, or are lying on the path on their own. The three left-most triplet of barplots are based on all the paths of the dataset, while the last triplet focuses on the *WAf* category.

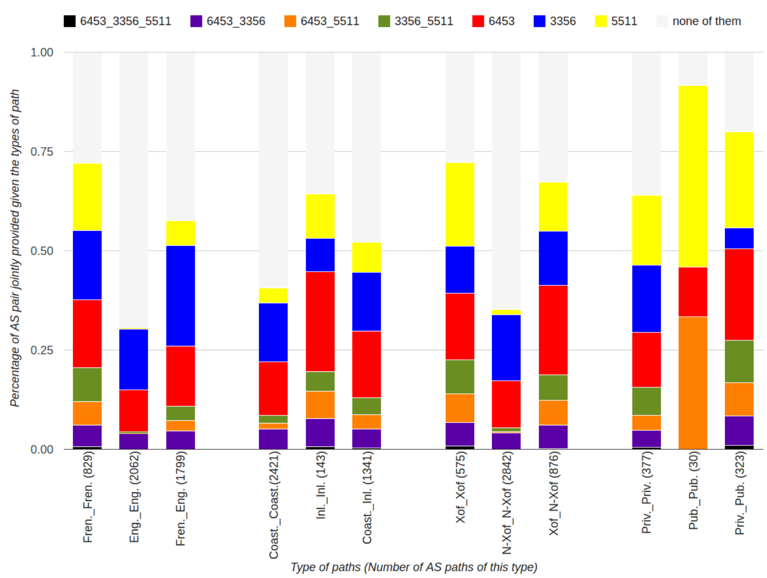

Fig. 2: *Joint AS-centrality of AS3356 (Level3, US), AS6453 (TATA, US), and AS5511 (Orange, FR), for paths among various categories of ASes*

French speaking countries mostly rely on Orange, which serves 17% of the West African AS pairs, without TATA nor Level3. Another 14.5% of AS pairs are also served by Orange, but jointly with TATA or Level3. Orange completely disappears from our internetworking map when it comes to communications among English speaking countries. Such diverse transit habits are also observed when classifying ASes according to the monetary region to which they belong. Within the XAF-XOF (CFA Franc) monetary union, Orange has a centrality of 36.7%, but is barely present in the market of communications among ISPs that are not belonging to this union.

From the same figure, we learn that Orange and TATA are lying together on 33% of the paths among the public owned *WAf* ASes[6], and Orange is lying alone on another 45.8% of these paths. In [15], we discuss the transit behaviour of these ISPs as a function of their ownership by Orange. No public operator seems to get transit from Level3. However, in the same region, a relevant proportion of pairs of ASes involving a private owned AS are served via Level3. Finally, the second triplet of barplot shows that African inland AS pairs rely much more on TATA (54.17%) than on Level3 (29.16%), dominating Orange. Such differences can be explained by the scarcity of Internet transit offerings in inland countries, mostly relying on Satellite transport companies which peer with Level3 and TATA.

4.5 Impact of transit localization on end-to-end delay

We identify, per AS path, the IP path over which the minimum RTT was observed, as well as its corresponding country path. We group AS paths into two categories; continental AS paths (from 1 to 1,073) are those which stay within Africa, whereas inter-continental ones (from 1,074 to 4,082) are via at least one node geolocated outside the continent (i.e. the country path contains at least one country outside Africa).

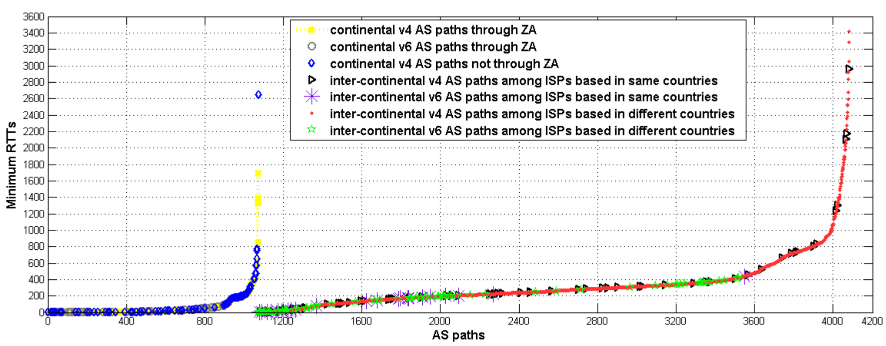

Fig. 3: *Minimum RTT distribution*

Figure 3 shows the distribution of the minimum RTTs among our probes. Continental paths with very low RTTs mostly correspond to paths among pairs of ASes based in the same country, or those passing through collocated regional ISPs. As highlighted by the yellow crosses, many of such paths are through South Africa, acting as a regional hub. Note that all the continental v6 paths traverse ZA.

Slightly longer RTTs (50-150ms) are seen among AS pairs from geographically distant countries. For instance, a path from a KE ISP to a ZA ISP, only served by African

[6] We categorized the *WAf* ASes as public or private, based on gathered private information

transit ISPs, shows a minimum RTT of 80ms. A striking result comes from the presence of very long RTTs in paths that are categorized as continental ones. These v4 paths are typically those between Eastern African ISPs and Western African ISPs, which are served by ZA transit ISPs. However, a v4 path, between GA and NG, categorized as continental, appears with an RTT of 2.6 seconds. Actually, this path is probably mis-categorized, as its IP level traceroutes contains many non answering hops. The following long RTTs are recorded on paths from TZ to ZA via AS37100 (SEACOM, MU), from Internet Solutions (ZA) to Simbanet (TZ) via KE, or from SAIX-NET to TENET in ZA. They are having the same issue of mis-categorization, as per our manual checks, but we have no data allowing us to certify that they leave the continent.

Let us now analyze the paths categorized as intercontinental. Intercontinental paths with a low RTT (i.e. $\prec$100ms) also reveal the weakness of geolocation. These AS paths contain *Int* ASes, as per TC, and have also been consistently geo-located in either GB, NL, FR or US by the geo-location databases. These are cases where all geo-location databases are returning the same Country Code, located outside Africa, although delay-based measurements clearly indicate that the device is located on the continent. Our ongoing work includes the correction of such databases, in order to account for the new measurements performed towards these mislocalized IPs.

Most of the measured RTTs in this category however reflect intercontinental transit of continental traffic, with a RTT around 200ms. 95,4% of the paths with a RTT between 100ms and 400ms are through Europe. Paths with RTTs scattered around 750ms are mostly from and towards ISPs that are served by Satellite providers, routing traffic through another continent. For example, a path in this group is from Connecteo in BF to AFNET in CI, passing through SkyVision, Level3 (in New York), Level3/Global-Crossing (in London), and MTN (ZA). The paths measured with an RTT above 1000ms are mostly those served via 2 satellite links. For instance, one is from Connecteo in BF to Sonitel in NE, going through the US and Europe, but arriving in NE via another satellite, provided by IntelSat. Finally, we highlight the RTTs between ISPs operating in same African countries, exchanging packets over intercontinental AS paths. These are notably observed in BJ, CM, MA, MZ, and MU.

4.6 Emergence of new IXPs

Let us now focus our analysis on paths revealing the use of IXPs to exchange traffic. We collected IXP information from [23,24,25]. We also learned, through word-of-mouth, that new IXPs were being deployed in BJ, SC, and GM. We crossed such information with our traceroute results, and detected IPs used to address interfaces to these IXP in our traceroute data.

We notice some frequently used IXPs, notably JINX, CINX, DINX, and NAP Africa in ZA. Actually, 58,6% of the continental paths which traverse ZA, go through one of these IXPs. We found the new IXPs in SC, BJ and GM. In SC, 4 members of the new IXP were hosting one of our probes, at the beginning of the 2nd campaign. We could observe a delay around 1ms among each pair of this clique, formed by CWS, Intelvision, Telecom Seychelles, and Kokonet Ltd. In the data collected during our third measurement campaign, we find probe hosts connected to SIXP, in GM: QCell, Netpage, and GAMTEL. RTTs are around 1.5ms among QCell and NetPage, while RTTs involving

GAMTEL fluctuate between 1ms and 460ms. Measurements performed between the GAMTEL probe and the IXP platform itself actually revealed unstability of the link from GAMTEL to SIXP, as detailed in [15].

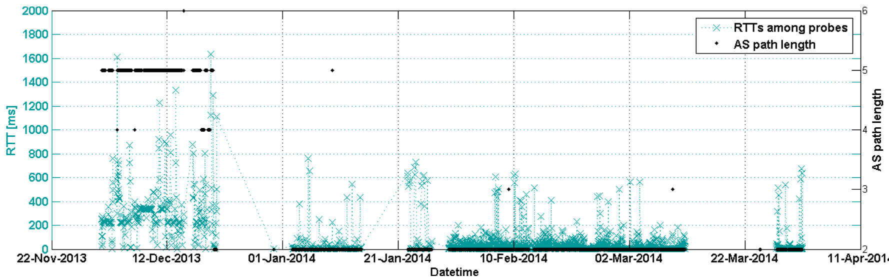

Fig. 4: *RTTs from Benin Telecom to Isocel Telecom during Benin-IX (BJ) establishment*

Let us now analyze the measurements performed among members of Benin-IX, being Benin Telecom, Isocel Telecom, and OTI Telecom. From December 2013 to the end of our first measurement campaign, in April 2014, RTTs measured among those ASes considerably drop from 314ms on average between November 30th and December 20th 2013, to 42ms on average from January 2nd to April 6th 2014. Figure 4 illustrates the benefit brought by this IXP, depicting the RTT among two members of that IXP, as well as the length of the measured AS Sequence. The figure also shows that our probes lost internet connectivity during the establishment of the IXP, as very few traceroutes succeeded during that period.

5 Conclusions And Future Work

In this paper, we assessed the global African interdomain routing topology. To this end, we deployed new RIPE Atlas probes, and carried out active measurements from 214 RIPE Atlas probes located in 90 ASes, covering 32 African countries[7].

We notice striking differences in transit habits of ISPs, notably depending on the official language of the country, the monetary region, or the business profile of the ISP. These illustrate how critical it is to have a large, diversified set of vantage points before drawing conclusions on the state of interdomain routing on the continent. Our results show a lack of interconnection among African ISPs (South Africa being an exception), confirming the interest of initiatives to promote peering on the continent. We highlight the remaining reliance on intercontinental ISPs for the establishment of continental connectivity. We correlate such trombonning paths with long RTTs among our probes. Nevertheless, new IXPs are emerging in Africa, notably in the West. We illustrate their benefits by showing the improvement in terms of RTT observed among their members.

In the future, we plan to measure the connectivity between African ISPs and the rest of the world. We also plan to provide a model to study the opportunities for cost reduction brought by IXP initiatives on the continent.

[7] As of December 10, 2014, the RIPE Atlas platform has evolved to 318 probes hosted in 147 ASes and spread across 44 countries all over Africa [7].

Acknowledgement. We are grateful to whoever deploys or hosts a RIPE Atlas probe. We also thank Michuki Mwangi, Nishal Goburdhan, Vesna Manojlovic, Andra Lutu, Camilo Cardona, Ignacio De Castro, and Pablo Camarillo for their insightful comments.

References

1. R. Les Cottrell, “How Bad Is Africa’s Internet?.” http://spectrum.ieee.org/telecom/internet/how-bad-is-africas-internet, January 2013.
2. M. Chetty, S. Sundaresan, S. Muckaden, N. Feamster, and E. Calandro, “Measuring Broadband Performance in South Africa,” in *Proceedings of the 4th Annual Symposium on Computing for Development*, ACM DEV-4 ’13, (New York, NY, USA), ACM, 2013.
3. African Union, “Study On Harmonisation of Telecommunication, Information and Communication Technologies Policies and Regulation in Africa.” http://www.itu.int/ITU-D/projects/ITU_EC_ACP/hipssa/docs/2_Draft_Report_Study_on_Telecom_ICT_Policy_31_March_08.pdf, March 2008.
4. African Union, “African Internet eXchange System.” http://pages.au.int/axis.
5. J. Gilmore, N. Huysamen, P. Cronje, M. de Klerk, and A. Krzesinski, “Mapping the African internet,” in *Proceedings Southern African Telecommunication Networks and Applications Conference (SATNAC), Mauritius*, September 2007.
6. A. Gupta, M. Calder, N. Feamster, M. Chetty, E. Calandro, and E. Katz-Bassett, “Peering at the Internet’s Frontier: A First Look at ISP Interconnectivity in Africa,” in *PAM*, p. 10, 2014.
7. RIPE NCC, “Global RIPE Atlas Network Coverage.” https://atlas.ripe.net/results/maps/network-coverage/, 2014.
8. Z. M. Mao, J. Rexford, J. Wang, and R. Katz, “Towards an Accurate AS-level Traceroute tool,” in *Proceedings of ACM SIGCOMM*, pp. 365–378, 2003.
9. L. Subramanian, S. Agarwal, J. Rexford, and R. H. Katz, “Characterizing the Internet Hierarchy from Multiple Vantage Points,” in *Proceedings of IEEE INFOCOM*, p. 12, 2002.
10. H. Haddadi, M. Rio, and A. Moore, “Network Topologies: Inference, Modeling, and Generation,” in *IEEE Communications Surveys and Tutorials*, pp. 48–69, IEEE, 2008.
11. Y. Hyun, “Archipelago Infrastructure.” http://www.caida.org/projects/ark.
12. PingER, “PingER Project.” http://www-iepm.slac.stanford.edu/pinger/, 2013.
13. B. Augustin, X. Cuvellier, B. Orgogozo, F. Viger, T. Friedman, M. Latapy, C. Magnien, and R. Teixera, “Avoiding traceroute anomalies with Paris traceroute,” in *IMC ’06 Proceedings of the 6th ACM SIGCOMM conference on Internet measurement*, pp. 153–158, October 2006.
14. P. Cristel, C. Luca, V. Stefano, and R. Bush, “From Paris to Tokyo: On the suitability of ping to Measure Latency,” in *IMC ’13 Proceedings of the 13th IMC*, pp. 427–432, 2013.
15. R. Fanou, P. Francois, and E. Aben, “African Measurement Campaigns: Technical Report.” https://fourier.networks.imdea.org/external/techrep_amc/, September 2014.
16. AFRINIC, “AFRINIC database.” ftp://ftp.afrinic.net/, 2014.
17. J. Karlin, S. Forrest, and J. Rexford, “Nation-State Routing: Censorship, Wiretapping, and BGP,” *CoRR*, vol. abs/0903.3218, no. http://arxiv.org/abs/0903.3218, 2009.
18. RIPE NCC, “OpenIPMap database.” https://labs.ripe.net/Members/emileaben/infrastructure-geolocation-plan-of-action, 2014.
19. MaxMind, “GeoIP.” http://www.maxmind.com/en/geolocation_landing, 2014.
20. Team Cymru, “Team Cymru Services.” https://www.team-cymru.com/.
21. RouteViews Project. http://routeviews.org, 2005.
22. RIPE, RIS, “Raw data.”
23. Packet Clearing House (PCH), “PCH IXP directory.” http://prefix.pch.net/images/applications/ixpdir/ip_asn_mapping.txt, 2014.
24. CAIDA, “AS Relationships.” http://www.caida.org/data/as-relationships.
25. PeeringDB. http://www.peeringdb.com/private/exchange_list.php.