

Deep Diving into Africa’s Inter-Country Latencies

Agustin Formoso*, Josiah Chavula†, Amreesh Phokeer† ‡, Arjuna Sathiseelan§, Gareth Tyson¶

*LACNIC †AFRINIC ‡University of Cape Town

§University of Cambridge ¶Queen Mary University of London

Emails: agustin@lacnic.net, josiah.chavula@afrinic.net, aphokeer@cs.uct.ac.za

arjuna.sathiseelan@cl.cam.ac.uk, g.tyson@qmul.ac.uk

Abstract—The Internet in Africa is evolving rapidly, yet remains significantly behind other regions in terms of performance and ubiquity of access. This clearly has negative consequences for the residents of Africa, but also has implications for organisations designing future networked technologies that might see deployment in the region. This paper presents a measurement campaign methodology to explore the current state of the African Internet. Using vantage points across the continent, we perform the first large-scale mapping of inter-country delays in Africa. Our analysis reveals a number of clusters, where countries have built up low delay interconnectivity, dispelling the myth that intra-communications in Africa are universally poor. Unfortunately, this does not extend to the remainder of the continent, which typically suffers from excessively high delays, often exceeding 300ms. We find that in many cases it is faster to reach European or North American networks than those in other regions of Africa. By mapping the internetwork topology, we identify a number of shortcomings in the infrastructure, most notably an excessive reliance on intercontinental transit providers.

I. INTRODUCTION

Africa currently has the lowest rate of Internet penetration in the world [1], with many unable to afford access [23]. This is set to change with Africa predicted to be a major driving force in expanding global uptake. Despite this, recent studies have observed generally poor performance on the continent, *e.g.*, slow page load times [11]. Although the exact causality is yet to be seen, there are a number of general trends that can be highlighted, including the use of high-delay access technologies and suboptimal country-level topologies [10]. We argue that understanding and quantifying these issues is critical for not only their short-term amelioration, but also for informing future design and deployment strategies, *e.g.*, for CDNs in the region. Importantly, the underdeveloped (but rapidly expanding) nature of Africa’s Internet ecosystem means that this must be done *now*.

Although there have been a number of seminal studies that have explored global Internet performance and topology [5], [9], [25], the diversity of networking infrastructure across Africa makes them largely inapplicable. Of particular interest is the means by which the countries of Africa are interconnected; only by improving this can a local Africa-wide Internet ecosystem flourish. For example,

the deployment of much-needed African data centres depends on underlying connectivity to make them available to the wider region [16]. Thus, in this paper we ask a simple question: *What is the inter-country delay in Africa, and how is this impacted by topology and interconnection strategies?* Our long term goal is to exploit this newfound understanding for informing future protocol and application developments in the region.

Answering the above question, however, requires vantage points across Africa — a challenge which has prevented many studies from focussing on Africa. Hence, we exploit a new platform, *Speedchecker* (a commercial measurement service), which has around 850 probes in Africa. This allows us to cover 322 networks and 52 countries (§III). Our measurements reveal a highly uneven delay distribution, with some countries exhibiting European-like delays (*e.g.*, South Africa → Botswana takes 25ms), whilst others suffer from delays exceeding 300ms (going up to 900ms). For context, typical latency in North America is <45ms and <30ms for Europe [2]. This leads us to explore patterns and cluster countries into groups of high connectivity (§IV). This reveals distinct geographical patterns, as well as a number of corner cases, where more distant countries actually have lower delay than nearer countries (§V). We find that some countries and regions have built up relatively low delay infrastructure, although many others have not. To explore this, we inspect the continent’s topology to identify key issues in the region (§VI). We find that the use of intercontinental transit (rather than local interdomain peering) plays a key role in inflating delays. This sees Africa→Africa packets leaving the continent via international transit, simply to re-enter again in a circuitous manner. This dramatically increases network operator costs, due to the high prices charged for international transit [7]. It also makes certain common infrastructure deployment practices unworkable, *e.g.*, it makes little sense to deploy content servers at regional exchange points if networks do not peer there [11]. Whereas we quantify the delay impact these decisions have, we also observe cases in which using European or American upstream providers actually results in *better* performance than using African upstream networks. Such observations best highlight the immediate challenges in the region, and help explain the difficulty in accessing and deploying services on a continent-wide basis (as exemplified by several

studies [17], [20]). Our findings offer insight into how these problems can be addressed (§VII).

II. RELATED WORK

Recent studies have begun to recognise the traffic engineering problems in Africa’s Internet topology [6], [10], [13], [14]. Primarily, these studies have highlighted Internet performance issues that are attributed to a lack of peering amongst Africa’s ISPs. Due to this, models that aim to predict global latencies have consistently modelled Africa as the slowest in the world [15]. Studies have also looked at inefficient DNS configurations, a lack of local content caching servers, as well as a lack of cross-border cable systems [11], [16], [27].

Gilmore *et al.* [13] performed a logical mapping of Africa’s Internet topology, highlighting the router level and Autonomous System (AS) level paths followed by intra-Africa traffic. Their analysis was based on traceroute data obtained from measurements conducted from a single vantage point in South Africa towards all AFRINIC allocated IP addresses. The key limitation of this work was that it only contained one-way paths from South Africa.

Gupta *et al.* [14] increased the number of vantage points, although they still launched probes from a small set of countries. Similarly, Chavula *et al.* [6] used five nodes from the CAIDA Archipelago platform to conduct logical topology mapping for Africa’s national research and education networks. They found that over 75% of Africa’s inter-university traffic followed intercontinental routes. Although more extensive than [13], these restricted studies only provide insights into a small number of countries and networks. To the best of our knowledge, the first study to take a wide-area perspective was by Fanou *et al.* [10]. This work launched traceroute measurements from 90 ASes. Their results too showed a lack of direct interconnection amongst African ISPs.

Our research differs from these past works in three key ways. *First*, our focus is not on enumerating traceroute paths leaving the continent. Instead, we strive to characterise the delay of the *inter-country* interactions. Our rationale is that the future success of local services depends on low delay underlying connectivity within the whole region (not leaving the region). *Second*, we do not simply observe network performance — we explore the causality behind high delays, and the implications of the topology configurations observed. Further, we utilise this data to identify the key clusters of connectivity in the region. *Third*, our study achieves the above goals on a scale not seen before, covering 52 countries and 319 networks across Africa.

III. DATA COLLECTION

A. Measurement platform

Due to the deficit of research infrastructure in Africa, there are only two feasible platforms for launching our

measurements: (i) RIPE Atlas,¹ which is known for providing a worldwide network of physical probes to their members; and (ii) Speedchecker,² a platform consisting of software agents installed on desktop clients. The Speedchecker platform offers Internet performance monitoring through ICMP ping, DNS and traceroute. Both platforms have probes deployed in Africa. At the moment of writing, RIPE Atlas had 229 active probes in Africa, covering 36 African countries, whilst the Speedchecker platform has nearly 850 installations covering 52 countries. Unfortunately, RIPE Atlas also has a strong bias towards university networks, as well as around half of all probes hosted in South Africa. In contrast, Speedchecker covers 91% of African countries and is not biased towards university networks. Hence, for this study, we select Speedchecker.

B. Data collection

We have used Speedchecker to collect two core datasets based on latency and topology measurements. We choose to focus on these low-level metrics as they can be used to shed insight on how various applications might perform, *e.g.*, web, gaming. Latency data was collected by launching pings from *all* Africa-based Speedchecker probes to randomly selected Speedtest servers located in African countries. There are 213 Speedtest servers in Africa, covering 42 countries (from 54). Note that this means we have countries with sources (Speedchecker clients) but not destinations (Speedtest servers). In these cases, we cannot compute the intra-country delays, and therefore exclude them from later analysis. Speedtest servers are generally hosted by ISPs and are therefore perfect end-points for network to network latency profiling. Full details about the service and its locations can be found online.³ The measurements were launched four times a day, at 00:00, 06:00, 12:00, and 18:00 probe time.⁴ In each case, we randomly selected up to 20 probes from all countries in that time zone. These probes were then instructed to launch 10 consecutive pings (one second apart) to their randomly chosen Speedtest server. Following this, the Speedchecker API returns the minimum ping delay observed, giving us the “best” observed delay at that time period.

By repeating this each day for 3 months, we garnered delay measurements across the continent, consisting of 42.2k ping samples. To quantify the coverage, Figure 1 presents the percentage of networks that the Speedchecker probes covered across each country (we take the overall count from the AFRINIC allocation files). In total, our data covers 319 networks across 52 African countries. As shown in Figure 1 (top CDF), 50% of the countries had at least 20% of their networks probed.

To complement the raw delay measurements, we also launched a parallel traceroute campaign using the same

¹<http://www.atlas.ripe.net/>

²<http://www.speedchecker.xyz/>

³<http://www.speedtest.net/>

⁴Note that the times were based on the local time zone

Network coverage by country




Fig. 1: Percentage of Autonomous Systems covered per country by our measurements.

setup. At 00:00, 06:00, 12:00, and 18:00 local time, we launched traceroutes from up to 20 random probes in that timezone, targeting random Speedtest servers across Africa. This campaign covered 49 countries, and consisted of 31.5k traceroute measurements from 207 distinct networks. For each router hop within the traceroute data, we attached the Autonomous System (AS) using the RIPE Routing Information Service.⁵ We also attach the location of each router using MaxMind GeoLite2-City. We restrict ourselves to country-level analysis, as this has been found to have relatively high accuracy [24].

C. Limitations

It is important to highlight the limitations of our data. First, we have limited insight into the devices launching the measurements, *e.g.*, their network access technology. This means we cannot provide causal insight into the performance of *individual* measurement samples. That said, we have gained a sufficiently large number of samples from each country to allow us to extract regional trends. Second, all measurements are launched against Speedtest

servers. Although they have a wide geographical distribution, they are not necessarily representative of future server deployments. Hence, future services may observe different performances based on how and where they deploy their servers. Third, due to the nature of the Speedchecker service, it was not possible to launch all pings and traceroutes in parallel from the same locations. Hence, we have sampled different locations across the period of our campaign. Due to this, we later perform data aggregation to underpin our clustering and analysis. Although not ideal, this allows more tractable exploration and sheds insight for “typical” users. Finally, we also highlight known limitations with geolocation databases — it is likely that discrepancies within MaxMind introduce noise to the data. However, as stated above, we limit ourselves to country-level analysis [24] (to minimise impact), and note that this limitation is applicable to any study reliant on geolocation.

IV. CLUSTERING COMMUNITIES OF CONNECTIVITY

One of our goals is to detect the strengths and weaknesses of the connectivity in the African Internet. As a precursor to this, we analyse the relationships between countries by clustering them based on their latencies.

A. Clustering methodology

We convert the latency measurements into a graph structure (G) that represents connectivity between countries. Each node in the graph is a country, whilst a link represents a set of latency measurements between two countries. The links are weighted by the *median* latency observed. We select median as it provides insight into the “typical” delay seen between two countries, although clearly this removes a portion of data — particularly outlier networks within a country. Our future work will involve inspecting the full breadth of delays seen, and clustering on a per-network basis..

Once the weighted undirected graph G is computed, we cluster countries based on their latency-defined distance using the Louvain algorithm [3]. The Louvain community detection algorithm is based on the modularity function, which performs clustering based on the measure of partition between communities found in a network. Let $G = (V, E)$ be the graph of vertices V representing countries and E , a set of undirected edges representing latencies. Suppose $u, v \in V$, $e(u, v) \in E$ has a weight $w_{u, v}$, which is the median latencies from $u \rightarrow v$ and $v \rightarrow u$. The community detection algorithm partitions the graph into communities, C , as expressed by Equations (1) and (2).

$$\cup c_i = V, \forall c_i \in C \quad (1)$$

$$c_i \cap c_j = \emptyset, \forall c_i, c_j \in C \quad (2)$$

The quality of the partitioning is measured using *modularity* Q [19], where $-1 < Q < 1$. We define *modularity* as the difference between the number of intra-cluster

⁵<http://ris.ripe.net/>

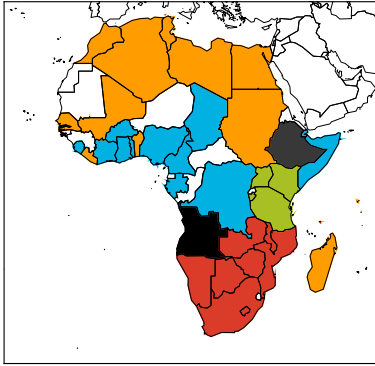


Fig. 2: Country clusters according to the Louvain clustering algorithm. Yellow is Northern; Red is Southern; Blue is (mostly) Western; Green is Eastern. Ethiopia (grey north-east) and Angola (black south-west) showed no clear cluster membership.

communities and the expected number of edges. Executing the algorithm outputs the set of strongly connected communities. It is expressed as follows:

$$Q = \sum_{c \in C} \left[\frac{\sum_{in}^c}{2m} - \frac{(\sum_{tot}^c)^2}{4m^2} \right] \quad (3)$$

where \sum_{in}^c is the sum of all weights (latencies) of the all internal edges of a community c and \sum_{tot}^c is the sum of weights (latencies) from edges incident to any vertex in community c . $m = \sum_{e(u,v) \in E} w_{u,v}$ is used to normalise the modularity and is obtained by adding the latencies across the entire graph. Once we know how to calculate *modularity*, we run the Louvain algorithm [21] to greedily maximise the modularity gain when moving a vertex u to community c .

B. Clustering results

The algorithm returns countries grouped into *four* different clusters, which correspond to the regions of Northern, Southern, Eastern, and Western Africa. Figure 2 presents a map of the clusters. Unsurprisingly, the clusters follow clear geographical properties. However, there are a number of unusual trends: most noticeably, Guinea, Liberia and Benin on the West coast, with neighbouring countries from a different cluster. Similarly, Madagascar, Seychelles and the islands of the Indian Ocean, are clustered alongside countries in the North. Somalia, on the East coast, is clustered with countries on the West coast. This suggests that geography is not the sole factor in defining delay. We explore this in §V and §VI.

The clustering algorithm also returned two special cases: Angola and Ethiopia, which were placed in separate clusters on their own. To understand this, we take a closer look at their latency profiles. Figure 3 depicts the distribution of RTTs between these countries and all other countries




Fig. 3: CDFs of RTTs (in ms) between the 4 clusters and some special cases: Angola and Ethiopia. Libya is added as a reference of a normal case. RTTs to each country's neighbours is also shown.

in the four clusters. We show Libya as an example of a country that shows typical trends. It can be seen that Libya exhibits very different delays across the different clusters. It has low delay to countries in the Northern cluster, but high delay to all others. In contrast, Angola and Ethiopia have roughly equivalent delays to countries within all clusters. For example, the median delay from Angola to all clusters is consistently above 200 ms. This explains why the algorithm could not allocate them to any clusters. To allocate them to an appropriate cluster, we manually inspect the data. The median delay from Angola \leftrightarrow Western cluster is 273 ms, which is 1.33x the median intra-cluster delay. The other options considered resulted in 8x, 2.7x, and 6.8x to the Eastern, Northern, and Southern clusters. Hence, we allocate Angola to the Western cluster. We computed the same ratios for Ethiopia, which is also allocated to the Western cluster. We discuss the implications of this in §VI.

V. QUANTIFYING DELAY IN AFRICA

Before exploring the topology of Africa, we quantify the delay attained between its countries. Note that delay is often the most prominent bottleneck of various protocols and applications [4], [8].

A. Exploring inter-country delays

Figure 5 presents a heatmap of the inter-country median delays observed (where the country originating the measurement is displayed vertically on the left and the target country is displayed horizontally on top). As expected, the lowest delays can mostly be observed within *intra*-country delays. There are also a number of inter-country delays that exhibit similarly low delay characteristics. These are primarily countries with close geographical proximity. For example, the delay between Tunisia and Algeria is just 48ms. This can be compared against the intra-country delays within these countries, which are 25ms and 44ms respectively. Examples of non-neighbouring countries with delays below 90ms include Kenya and Zimbabwe at 85ms, and Mauritius and Tanzania at 80ms.

To generalise this, Figure 4 presents a CDF of the inter and intra country delays. For context, we include the same data from the Latin America and the Caribbean (LAC) region using the same Speedchecker methodology (taken from [12]). Across the entire Africa dataset, *intra*-country latencies average at 78ms. This is significantly higher than that seen in more developed regions; the

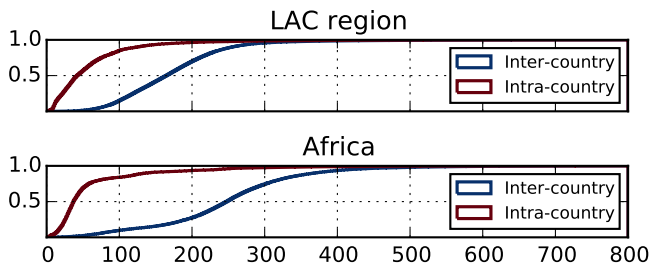


Fig. 4: CDFs showing intra-country and inter-country RTTs for the LAC region (taken from [12]) and Africa.

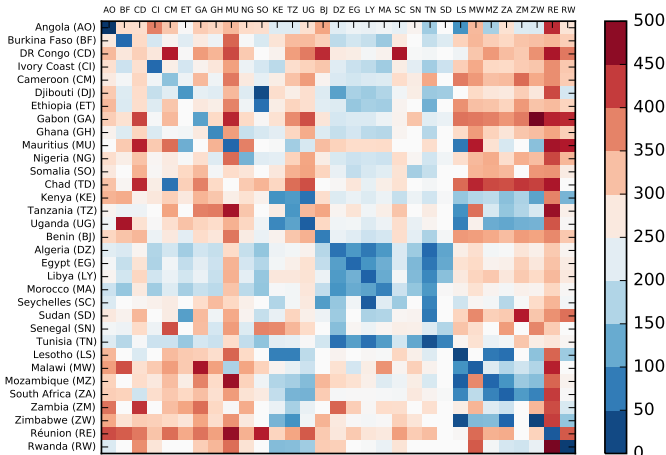


Fig. 5: Country-level (median) latency heatmap. Countries are ordered by latency proximity. White colour indicates no meaningful samples could be gathered between that pair of countries. Countries excluded that had more than 20 white cells.

average monthly latency in North America is $<45\text{ms}$ and $<30\text{ms}$ for Europe [2]. Our results are, however, close to the 76ms measured in the LAC region [12] (which is in a similar process of Internet development). However, the story is significantly different when comparing *inter-country* delays, also shown in Figure 4. Africa has a mean of 280ms , whilst the LAC region has only 154ms : a factor of $1.8\times$. 9% of inter-country delays exceed 400ms , and 2% exceed 500ms (shown in red). Again, this can be compared against the measurements performed in the LAC region, where less than 1% of country pairs have a latency greater than 500ms . In fact, these African delays are so poor that they go well beyond the sensitivity analysis ranges used by past studies that inspect the impact of network delay on things like web page load times [26] and video streaming performance [18]. For example, ranges of only $0\text{--}100\text{ms}$ were tested in [18]. In other words, past application performance studies would need to be entirely repeated with vastly higher delay parameters to understand their behaviour in Africa.

B. Exploring inter-cluster delays

The above presents the delays observed between individual countries. However, our findings in §IV have already identified that some groups of countries are better interconnected than others. Hence, we next inspect the delays observed between the clusters from §IV.

Figure 7 presents the CDFs of all delay samples split into clusters. It can be seen that, although there is a significant amount of overlap between the clusters, the Western cluster performs worse on an intra-cluster basis. The Southern cluster has the lowest delay. In contrast, the inter-cluster delays are consistently worse across all clusters, with over half of all samples exceeding 200ms . To complement the CDF, Figure 6a shows a heatmap of the median delays seen *between* clusters. Again, it can obviously be seen that intra-cluster delays are generally lowest. For example, the Eastern cluster has the lowest intra-delay (median of 44ms), which places it in a similar position to Europe. The one outlier is the Western cluster, which has a high median delay of 215ms . In terms of inter-cluster delays, the Southern and Eastern clusters have the lowest delay between them (median of 92ms).

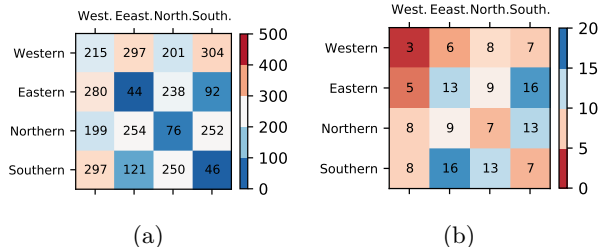


Fig. 6: Heatmap (a) shows median RTT, aggregated at cluster level. Heatmap (b) shows delay normalised by distance (based on equation 4).

There is, however, a key limitation in the above analysis. Small countries in close proximity will naturally have lower propagation delays (assuming direct links). Thus, in some cases, low RTT may simply be a property of geography. To address this, we next normalise the delays based on the geodesic path between the source and destination; this captures the *delay stretch*. We compute the stretch, for each ping sample, p , as:

$$\text{stretch}^p = \frac{d/(RTT^p)}{c \times 0.66} \quad (4)$$

where d is distance to the destination and c is the speed of light. We reduce c by a factor of $2/3$ to approximate propagation time through optical fibre [22]. RTT^p is taken as the minimal RTT (in seconds) measured from a given ping sample, p . The stretch therefore captures the ratio between the optimal observed RTT and the theoretical minimum RTT.

Figure 6b shows a heatmap containing the stretch value between each cluster. High values indicate strong con-




Fig. 7: CDFs showing intra-cluster RTTs (left) and inter-cluster RTTs (right).

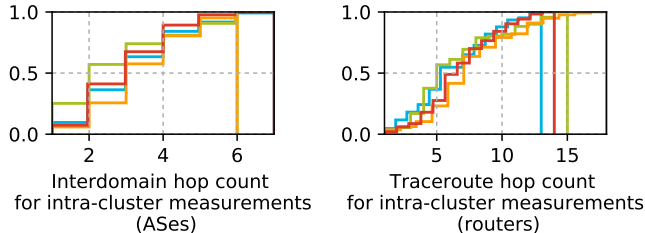


Fig. 8: Number of AS and router-level hops traversed by the traceroute.

nectivity; for example, a value of 20 indicates that the speed of the packet is 20% of the maximum theoretical speed. By comparing figures 6a and 6b we immediately identify differences. In figure 6b, some *inter*-cluster delays are actually lower than *intra*-cluster delays when measured using this normalised metric. In other words, some low observed RTTs are a property of geography — the clusters are still highly suboptimal. A good example is the Southern \leftrightarrow Eastern cluster, which has a high median delay of 92ms, yet performs far better when normalised by distance. Whereas Southern \leftrightarrow Southern only attains 7% of the optimal speed (46ms), Southern \leftrightarrow Eastern gains 16%. That said, there are some results that are consistent between the two heatmaps, namely the poor performance of the Western cluster. To explore why this might be, we initially checked Somalia, as it is actually geographically located on the Eastern coast (*cf.* Figure 2). However, curiously, Somalia actually has the fastest intra-cluster ping measurements in the Western cluster: located at around 20% of $0.66c$. The ping speeds in the remaining countries, instead, attained around just 5% of $0.66c$.

VI. DISSECTING PATHS ACROSS AFRICA

The previous section has highlighted the high network delays suffered when traversing countries in Africa. Next we inspect the reasons behind this using the topology maps obtained via our traceroute campaign.

A. Exploring topological traits

To begin, Figure 8 presents the distribution of hop counts across all traceroutes within the same clusters (*i.e.*, where the Speedchecker client and Speedtest server are in the same cluster). It can be seen that the traceroutes originating from the Eastern cluster have marginally fewer




Fig. 9: Map showing the first AS hop in the traceroute dataset. Countries with multiple networks have multiple links shown. Nodes colours correspond with their cluster’s colour, node sizes correspond with their in-degree.

hops than the others (median 5 router hops). Curiously, however, the Western cluster (which is the worst performing) also has fewer AS (median 3) and router-level (median 6) hops than the Northern and Southern clusters. This suggests that the higher delays are *not* simply driven by hop counts. Figure 9 presents a geographical map showing the upstream providers⁶ serving all networks sampled in each country. This offers immediate insight into the reasons behind the high delays previous seen. We find that a significant number of networks rely on upstream providers outside of the continent. Considering all traceroutes, we find that 37.8% of all routers in the data are geolocated outside of Africa. The remainder inside Africa are heavily biased towards a few prominent countries, namely South Africa and Mauritius via the West Indian Ocean Cable Company (WIOCC). 6.6% and 4.5% of traces upstream through them, respectively. Hence, although the hop counts do not differ significantly, the locations of the networks do. Figure 10 presents the most popular countries for hosting upstream networks. It can be seen that South Africa offers the most upstream provision, followed by the UK and US. Interestingly, the use of these upstreams differs substantially based on cluster, with regional hubs emerging, *e.g.*, Uganda for the Eastern cluster and South Africa for the Southern cluster.

Next, we inspect the specific upstream networks involved in these AS hops. Table I presents a list of the top upstream providers ranked by the number of edge networks connected to them by their first AS hop. The Top 10, alone, provide services to nearly half of all sampled networks. Rather than observing local tier-2 operators, the list is dominated by international tier-1 operators. Anecdotally, many African network operators prefer to

⁶We define upstream as the first AS hop after the origin AS hosting the Speedchecker client.

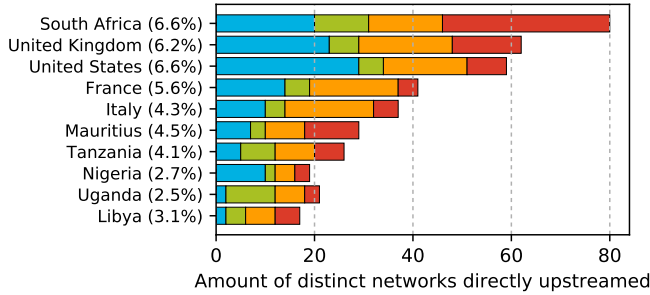


Fig. 10: Number of networks upstreaming through each country. Network samples separated into their clusters.

use such services due to their perceived reliability and international reputation. It is also worth briefly noting that there is a significant presence of the French operator, Orange. 17% of networks in French speaking countries utilise Orange, which add up to almost 40% of Orange’s downstreams (potentially driven by historical ties).

The above has shown that many countries rely on international and intercontinental upstream providers. An obvious follow-up is what *proportion* of paths from each country follow these intercontinental routes. To answer this, we inspect the hops taken through each AS on a per-sample basis. Figure 11 shows the percentage of traceroute hops that traverse each region. For each traceroute we elaborated the sequence of distinct ASes that the traceroute went through. Each stage in the figure indicates the AS path traversed so far. The first column corresponds with 0 hop count (origin), the second one with the first hop count, *etc.* The first observation is that a large fraction of traceroutes have their first hop in an overseas regions (50%), in particular through Europe. Our data shows that Europe acts as a major Internet provider to networks across Africa (35%), and to the Northern cluster in particular (transiting $\approx 40\%$ of the cluster’s traceroutes). Considering the geographic distance, it is surprising that 12% of African connections are routed through N. America. Arabia and Asia only account for 3% of outgoing paths. This also disproves the theory that only Northern African countries rely on Europe and Arabia.

Figure 11 further shows that intercontinental hops are not limited to the immediate (first hop) upstream provider. We find that even edge ASes that utilise African upstream providers see their traffic leaving the continent (potentially without their prior knowledge). This is because their upstream providers, in turn, utilise international transit rather than local peering. For example, 14% are traceroutes go via overseas parties for as many as 4 interdomain hops before returning to African operators. A particular case of this phenomena is the traffic exchange between N. America and Europe: 3.5% of the traceroute paths are routed between these two locations, even though all source and destination locations are within Africa.

Rank	ASN	Network info.	Perc.	Centrality
1	174	Cogent Communications	10.3%	0.095
2	3356	Level 3 Communications	7.4%	0.087
3	37100	SEACOM	7.1%	0.065
4	6762	Sparkle (TIM Group)	6.6%	0.071
5	30844	Liquid Telecom	5.9%	0.137
6	5511	France Telecom (Orange)	3.9%	0.044
7	57023	Oranlink	2.3%	0.003
8	6453	TATA COMMS. - US	2.2%	0.013
9	16637	MTN	2.1%	0.029
10	5713	Telkom SA Ltd	2.0%	0.019
Sum			49.7%	

TABLE I: Top 10 networks providing direct upstream access to African networks (first AS hop considered). Percentage is based on the fraction of traceroutes the ASN appears as a direct upstream.

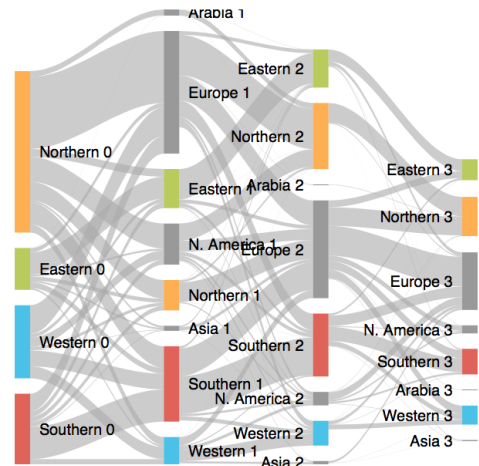


Fig. 11: Hop-by-hop analysis of traceroute data. Each step indicates the fraction of hops that fall into each cluster along the traceroute hops.

Naturally, this becomes a challenging problem to address as it is outside of the control of African edge networks, which largely depend on the routing decisions made by their upstream operator.

B. What are the delay implications?

The above shows that a large number of operators choose to utilise upstream providers outside of their own country (or even continent). To explore the implications of this, Figure 13 presents the minimum delay from African ASes to their upstream providers, *i.e.*, the minimum RTT to the first hop outside the edge AS. We use the minimum to capture the best possible scenario.

It can be seen that networks using upstream providers within their own cluster gain the best first-hop delay (median 40ms). Curiously, the Western cluster is amongst the best performing when using this metric. Note, however, that the use of a local upstream provider does not guarantee high end-to-end performance. For example, in the Western cluster, 3.15% of traceroutes exiting the cluster




Fig. 12: CDF of all RTTs from samples classified into three categories based on the first-AS hop upstream provider used by the origin network.




Fig. 13: Heatmap showing min. RTT to the first AS outside the origin AS network. This heatmap differs from the one in Figure 6a by showing min. RTT to the first AS.

are still intercontinental despite the use of a local first-hop upstream.

As a consequence of this, it can also be seen that regions that utilise intercontinental upstream providers have much higher delays. For instance, Southern networks that use Southern upstream providers witness a minimum of 47ms delay; this can be compared against 227ms when using N. American providers. Strangely, though, intercontinental delays are lower than some inter-cluster delays. Southern networks that use upstream providers in the Northern cluster see a minimum delay of 299ms. In fact, on average, clusters are marginally closer to overseas upstreams than to upstreams in other African clusters (180 vs. 195ms). That said, *intra*-cluster upstream consistently outperform both of these scenarios.

We can also investigate the overall impact these decisions have across our entire set of latency measurements. We split *all* ping samples into (i) networks that use an up-stream provider within the same cluster; (ii) networks that use an African upstream in another cluster; (iii) networks that use an overseas provider for over 50% of traceroutes observed. Figure 12 presents CDFs of RTTs seen within each group. It can be seen that, indeed, the lowest delay networks are those that upstream through a network within their own cluster (203 ms median). In-line with earlier discussions, the delta between networks using international providers and those based in other African clusters is limited (268 vs. 243 ms median).

C. Revisiting Angola and Ethiopia

We have already highlighted the problems with allocating Angola and Ethiopia to clusters (§IV). This motivates us to explore the topological reasons that caused the result. Note these were allocated to the Western cluster which consistently has performed the worst. To explain these results, we start by inspecting the upstream provision. Angola and Ethiopia have very different upstream selections: Ethiopia upstreams entirely through overseas providers, whilst Angola does it for 34% of the paths observed. In the case of Ethiopia, paths go through Europe (70%) and North America (30%). Considering the latency penalty of long haul links (Ethiopia → Europe at 354 ms, and Ethiopia → North America at 144 ms), the high RTT values are primarily driven by these circuitous routes. In contrast, Angola uses African upstream providers in 66% of traceroute samples. This makes it a more interesting case as, theoretically, it should therefore avoid connectivity through long haul links. However, when including the second AS hop, we find that an additional 16% of traceroutes go through the Southern cluster and are subsequently routed through Europe, adding to a total of 50% overseas paths (routed either direct or indirectly). In other words, despite networks in Angola not *directly* using international transit, their upstream providers do so. When combined, this distorts delays from both Ethiopia and Angola and pushes them away from the centre on any of the existing clusters, explaining the observations made in §IV.

These patterns also explain the poor delay performance exhibited by the Western cluster in general. For example, whereas the median intra-cluster delay of the Southern cluster is 46ms, it is 215 for the Western. When looking into the topology of the Western cluster, we further noticed that just 38% of hops occur between ASes within the cluster. Those intra-cluster hops benefit from shorter distances and have a median delay of 24 ms. The rest of hops are mostly Western → Europe hops (9.4%) which account for an added 132 ms (each way), and Europe ↔ Europe (21%), with marginal penalties. Thus, whereas other regions have built up strong intra-cluster connectivity, the Western region still lacks this.

VII. CONCLUSIONS AND FUTURE WORK

This paper has evaluated inter-country latencies across Africa. We have quantified latencies across 91% of African countries, and have identified a number of failings in the regional topology that require urgent attention. To the best of our knowledge, this is the largest study of its kind in Africa (§III).

There are several conclusions to draw. We have found that many intra-country delays in Africa have reached relatively developed levels (§V-A). We observed a large set of intra-country samples below 40ms (*e.g.*, Benin, Egypt, South Africa), and below 30ms (Ivory Coast, Réunion, Mauritius). Furthermore, we identify a series of country clusters, which have also built up strong interconnectivity.

Universally characterising the African Internet as poor is therefore misplaced. However, by corollary, this means that *inter*-cluster delays are significantly worse (§IV). Our clustering has shown that performance varies heavily based on region: Whereas some clusters have relatively low levels of delay (*e.g.*, the median intra-cluster delay in the South is just 46ms), other areas have consistently high delay (§V-B). For example, the Western cluster suffers from intra-cluster delays that are similar to its inter-cluster delays. Our analysis confirms that this is largely driven by the use of transit providers that route traffic through Europe and N. America (§VI). Our work therefore offers an effective means for automatically extracting regions and networks that critically require more local peering and interconnection. Only by addressing these issues will it become possible for high performance service hosting and interaction across the entire African continent.

There are many remaining topics to explore, such as the relationship between physical infrastructure and delay. For those countries with a direct physical connection, it would be expected that RTT would be lower. In order to find a relationship between infrastructure and RTT, there is need for further analysis of the submarine or terrestrial cables and in some cases satellite links. It is also worth noting that inter-country delays are only one part of the problem, and it is necessary to investigate delays between countries and popular web/content infrastructure (which the Speedtest servers may not be representative of). It is also important to expand our analysis to go beyond inspection of aggregated data and, instead, inspect the variance seen within networks, countries and clusters. Indeed, this work only presents one side of a latency story by aggregating the data at country level. Future work should include a more granular analysis *e.g.*, from a city to city or AS to AS perspective. Further, we believe that linking the findings to regional strategies (*e.g.*, deployment of IXPs) would reveal a more complex evolving picture. Finally, we wish to revisit a number of other studies (*e.g.*, web performance evaluations) to understand how they perform with more realistic African-level delays.

ACKNOWLEDGEMENTS

We thank the EPSRC for supporting this work via the African Internet Measurement Observatory (AIMO) project.

REFERENCES

- [1] Internet Penetration in Africa, 31 March 2017 internet world stats, 2017. <http://www.internetworldstats.com/stats1.htm>. Accessed: 26-Jul-2017.
- [2] IP latency statistics. <http://www.verizonenterprise.com/about/network/latency/>, 2017.
- [3] BLONDEL, V. D., GUILLAUME, J.-L., LAMBIOTTE, R., AND LEFEBVRE, E. Fast unfolding of communities in large networks. *Journal of Statistical Mechanics: Theory and Experiment*, 10 (2008).
- [4] BRISCOE, B., BRUNSTROM, A., ROS, D., HAYES, D., PETLUND, A., TSANG, I., GJESSING, S., AND FAIRHURST, G. A survey of latency reducing techniques and their merits. In *ISOC Workshop on Reducing Internet Latency, Sep* (2013).

- [5] BROIDO, A., ET AL. Internet topology: Connectivity of IP graphs. In *ITCom 2001: International Symposium on the Convergence of IT and Communications* (2001), pp. 172–187.
- [6] CHAVULA, J., FEAMSTER, N., BAGULA, A., AND SULEMAN, H. *Quantifying the Effects of Circuitous Routes on the Latency of Intra-Africa Internet Traffic: A Study of Research and Education Networks*. 2015, pp. 64–73.
- [7] CHRISTIAN, P. International bandwidth trends in africa. In *African Peering Forum (AfPIF)* (2015).
- [8] ELKHATIB, Y., TYSON, G., AND WELZL, M. Can spdy really make the web faster? In *IFIP Networking Conference* (2014).
- [9] FALOUTSOS, M., FALOUTSOS, P., AND FALOUTSOS, C. On power-law relationships of the internet topology. In *ACM SIGCOMM CCR* (1999), vol. 29, pp. 251–262.
- [10] FANOU, R., FRANCOIS, P., AND ABEN, E. *On the Diversity of Interdomain Routing in Africa*. 2015, pp. 41–54.
- [11] FANOU, R., TYSON, G., FRANCOIS, P., AND SATHIASEELAN, A. Pushing the frontier: Exploring the african web ecosystem. In *World Wide Web Conference (WWW)* (2016).
- [12] FORMOSO, A., AND CASAS, P. Looking for network latency clusters in the lac region. In *Workshop on Fostering Latin-American Research in Data Communication Networks* (2016).
- [13] GILMORE, J., HUYSAMEN, N., AND KRZESINSKI, A. Mapping the african internet. In *Proceedings Southern African Telecommunication Networks and Applications Conference (SATNAC), Mauritius* (2007).
- [14] GUPTA, A., CALDER, M., FEAMSTER, N., CHETTY, M., CALANDRO, E., AND KATZ-BASSETT, E. Peering at the internet’s frontier: A first look at isp interconnectivity in Africa. *Passive Active Measurement Conference (PAM)* (2014), 204–213.
- [15] KAUNE, S., PUSSEP, K., LENG, C., KOVACEVIC, A., TYSON, G., AND STEINMETZ, R. Modelling the internet delay space based on geographical locations. In *17th International Conference on Parallel, Distributed and Network-based Processing* (2009).
- [16] KENDE, M., AND QUAST, B. Promoting content in africa. *ISOC Report* (2016).
- [17] MATHUR, A., SCHLOTFELDT, B., AND CHETTY, M. A mixed-methods study of mobile users’ data usage practices in south africa. In *Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing* (2015), ACM, pp. 1209–1220.
- [18] MOK, R. K., CHAN, E. W., AND CHANG, R. K. Measuring the quality of experience of HTTP video streaming. In *Integrated Network Management (IM), 2011 IFIP/IEEE International Symposium on* (2011), pp. 485–492.
- [19] NEWMAN, M. E. Analysis of weighted networks. *Physical review E* 70, 5 (2004).
- [20] PHOKEER, A., DENSMORE, M., JOHNSON, D., AND FEAMSTER, N. A first look at mobile internet use in township communities in south africa. In *Proceedings of the 7th Annual Symposium on Computing for Development* (2016), p. 15.
- [21] QUE, X., CHECCONI, F., PETRINI, F., AND GUNNELS, J. A. Scalable community detection with the louvain algorithm. In *International Parallel and Distributed Processing Symposium (IPDPS)* (2015), pp. 28–37.
- [22] RAMASWAMI, R., SIVARAJAN, K., AND SASAKI, G. *Optical networks: a practical perspective*. 2009.
- [23] SATHIASEELAN, A., TYSON, G., AND SEN, S. Exploring the role of smart data pricing in enabling affordable internet access. In *IEEE INFOCOM Smart Data Pricing Workshop (SDP)* (2015).
- [24] SHAVITT, Y., AND ZILBERMAN, N. A geolocation databases study. *IEEE Journal on Selected Areas in Communications* 29, 10 (2011), 2044–2056.
- [25] SIGANOS, G., FALOUTSOS, M., FALOUTSOS, P., AND FALOUTSOS, C. Power laws and the as-level internet topology. *IEEE/ACM Transactions on Networking (TON)* 11, 4 (2003), 514–524.
- [26] WANG, X. S., BALASUBRAMANIAN, A., KRISHNAMURTHY, A., AND WETHERALL, D. Demystifying page load performance with WProf. In *NSDI* (2013), pp. 473–485.
- [27] ZAKI, Y., CHEN, J., PÖTSCH, T., AHMAD, T., AND SUBRAMANIAN, L. Dissecting web latency in ghana. In *Internet Measurement Conference (IMC)* (2014), pp. 241–248.