Quantifying the Effects of Circuitous Routes on the Latency of Intra-Africa Internet Traffic: A Study of Research and Education Networks

Josiah Chavula[1], Nick Feamster[3], Antoine Bagula[2], and Hussein Suleman[1]

[1] University of Cape Town, Department of Computer Science,
Private Bag X3, Rondebosch, 7701, South Africa
jchavula@cs.uct.ac.za
[2] University of Western Cape, Department of Computer Science,
Private Bag X17, Bellville 7535, South Africa
[3] Georgia Tech, School of Computer Science,
266 Ferst Drive, Room 3348, Atlanta, Georgia 30332, USA

Abstract. Despite an increase in the number of Internet eXchange Points (IXP) in Africa, as well as proliferation of submarine and terrestrial fibre optic cable systems, the level of peering among Africa's Internet service providers remains low. Using active network measurements, this work characterizes the level of interconnectivity and peering among Africa's National Research and Education Networks (NRENs), and examines the performance of traffic exchange in terms of latencies. This paper shows that over 75% of Africa's inter-university traffic follows circuitous inter-continental routes, and is characterised by latencies that are more than double those of traffic exchanged within the continent.

Key words: Round-trip time, latency, peering, active topology measurements, Internet exchange points, National Research and Education Networks

1 Introduction

Latency, measured as round trip time (RTT) for traffic to move from source to the destination, and for the acknowledgement packet to be received by the sender [1], is an important characteristic of Internet connectivity that affects the performance and responsiveness of Internet applications, especially real-time and interactive ones. This is particularly important for National Research and Education Networks (NRENs), where many education and collaboration oriented applications are used. High latency makes it difficult for research communities to make use of Internet-based collaborative tools such as real-time remote lectures or sharing of virtualized resources such as computer processors [2]. Furthermore, many scientific research facilities, such as the Large Hadron Collider (LHC) and other astronomical observatories, generate lots of data at high speed that needs to be exchanged among research centres around the world. These applications necessitate low latency networks among the research community.

In Africa, interconnectivity and peering among Internet Service Providers (ISPs) remains low [3]. With many of universities obtaining Internet connectivity from such ISPs [2, 4], traffic exchanged among African NRENs traverse higher tier transit providers and global Internet Exchange Points (IXP) through long intercontinental links, resulting in high latency and data transmission costs.

The contribution of this work is two fold; first through active topology measurements, it takes a fresh look at the logical topology of the African Internet, with a specific focus on education and research networks. As detailed in section 3, topology measurements are carried out from 5 vantage points targeting a selected set of 95 universities. The results are presented in the form of AS-level and PoP-level maps, presented in section 4. The second contribution is the performance analysis for traffic that uses inter-continental links, and this paper shows that links connecting Africa to other continents contribute over 50% of the latency for the traffic that uses those links. As presented in Section 4.2, more than 75% of intra-Africa traffic utilizes inter-continental links.

2 Background and Related Work

The Internet is an interconnection of many privately managed networks known as Autonomous Systems (ASes) [5]. Traffic exchange among ASes is facilitated through the Border Gateway Protocol (BGP), a single path routing system that conveys AS-level paths between domains, enabling them to interconnect and exchange traffic. Any two ASes can exchange traffic if they have some direct logical connection between them, or if they both have access to other higher level providers that can transit traffic between them [6]. Due to this hierarchical structure of the Internet, traffic whose source and destination networks are geographically close may sometimes have to traverse circuitous remote links in search of interconnecting paths - a phenomenon know as boomerang routing.

To obtain a better understanding of an internetwork, topology discovery techniques are used to obtain data for network visualization. Internet topology discovery techniques can largely be grouped into two: passive techniques and active techniques [7]. Passive techniques involve analysing network management data such as BGP routing tables in order to infer the network topology in terms of logical relationships among ASes. Active measurements, on the other hand, rely on sending specially crafted packets into the network with the aim of soliciting topology information [8]. Active measurement techniques attempt to exploit network management protocols such as SNMP and ICMP to solicit responses from a set of network destinations, and then analyse such responses to infer topological characteristics such as route paths, round-trip-times (RTTs) and packet loss.

Recent work on the African Internet topology [3] has shown that about 66% of traffic between South African Internet vantage points and Africa-based Google cache servers is routed outside the continent. The same work also characterized the IXP peering situation in Africa and showed that most African ISPs do not peer among themselves at national or regional IXPs, but rather prefer to peer at larger European IXPs such as London and Amsterdam, presumably to achieve better economies of scale with access to global networks. Gilmore et al. [9] also carried out topology measurements on the African Internet from a South African vantage point, and showed that routes originating from the South African Tertiary Education Network were mostly routed via the UK, Scandinavia and the USA.

3 Dataset

This work used active network topology discovery techniques to characterise the performance of the traffic whose source and destination is within Africa, and also to obtain a logical connectivity map of the African NRENs. In particular, this work performs active Internet measurements from 5 vantage points that are part of the CAIDA's Internet measurement platform - Archipelago, and are located in North, West, East and Southern Africa (Morocco, Gambia, Senegal, South Africa and Rwanda). The Archipelago measurement platform is based a network measurement tool called Scamper [10], which implements Paris-traceroute [11], a variant of traceroute based on Multi-path Discovery Algorithm (MDA) [12, 13].

3.1 Traceroute Measurements

In this work, Internet path traces were performed to a set of 95 university campus IP addresses across 29 African countries. Traceroute probes from 5 vantage points to each of the IP addresses were repeated daily for 14 days from 6 April to 20 April, 2014.

Of interest from each of the traces is the round-trip time for each source-destination pair, as well as the geo-location of the IP path hops used. Another interest is to compare the latency of the traffic that is routed through inter-continental links (boomerang traffic) in comparison to locally routed traffic. For this purpose, traces from each vantage point are grouped into two; inter-continental traffic that gets routed at Points of Presence (PoPs) outside Africa, and intra-Africa traffic that gets routed within the continent. For the inter-continental traffic, a further interest is to quantify the effect of the inter-continental links (i.e., latency from the vantage point to remote inter-continental gateway) on the overall RTT.

To observe the extent of inter-continental link utilization from each vantage point, traceroute traces are grouped by source and destination pairs. From each vantage point, the traces for each source-destination are analysed as follows: starting from the source (vantage point), the next hop and its corresponding RTT is extracted; using MaxMind's GeoIPLite database, each hop's geographical location (country and longitude/latitude coordinates) are obtained; if the extracted hop is located outside Africa, then the traffic trace is categorised as inter-continental. For each inter-continental route, the first hop outside Africa, together with its corresponding RTT, is recorded as the inter-continental gateway RTT for the route.

The RTT for each source-destination pair is taken as the RTT of the last responding hop. However, since not all probes reach the end point (among others due to blocking by routers [7]), only probes that reach the destination network are considered valid RTT values for each source-destination pair. To determine if the destination network is reached, the last responding hop's autonomous system number (ASN) and country is compared with that of the probe destination. Finally, the RTT values for each source-destination pair are computed as the mean of the multiple RTTs.

3.2 Dataset limitations

Although the topology measurements were carried out from multiple vantage points, the drawback is that the paths discovered are only forward paths from the vantage points to the destinations. This is the case because Internet traffic is largely asymmetric, ie, forward paths are not necessarily the same as reverse paths [5]. Therefore, the maps obtained from outgoing traceroute measurements are still incomplete. A more complete picture can be obtained by increasing the number and distribution of vantage points [14].

Another major challenge on the analysis of the dataset is the inaccuracy of the geolocation information for the IP addresses. Maxmind's free geolocation database - GeoLite2 - reported to have about 80% accuracy for IP to city resolution (within 40km) for most countries [15]. For example, the accuracy level reported for South Africa's IP to city database is 71%, whereas for Kenya it is reported to be as low as 55%. However, the accuracy for IP to country resolution, on which route categorisation is based, is higher at 99.8%.

4 Results

4.1 Circuitous Logical Links

A key observation from the traceroute data is that a larger percentage of traffic originating in Africa and destined for African universities gets routed through PoPs that are outside the continent. On average, 75% of the traces from African vantage points to African NRENs traversed inter-continental links through PoPs in Europe, such as Amsterdam, London, Lisbon, and Marseille. However, depending on the geographical location of the vantage points, different levels of inter-continental traffic is observed. For example, the vantage points along the north-west coast of Africa used inter-continental links for as much as 95% of the traces, whereas vantage points in central and southern Africa had a relatively lower usage of inter-continental links. From the Rwandan vantage point, 70% of the traceroute traffic used inter-continental links, while the South African vantage point had only 60% of the traffic using inter-continental links. The lower usage of inter-continental links by the South African vantage point can be attributed to the direct logical links observed between South Africa and some of its neighbouring countries, such as Mozambique, Zambia and Zimbabwe, as well as links to East Africa, such as the EASSY submarine fibre-optic cable.

Figure 1 shows the PoP-level connectivity map for traffic originating at the five vantage points to the addresses in the sample.

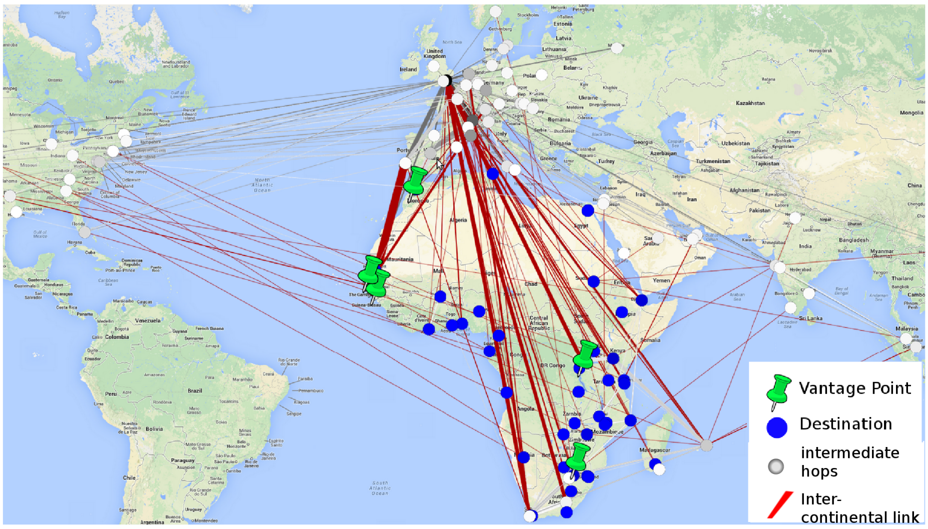

Fig. 1. Logical paths for African traffic, showing logical links interconnecting in Europe and North America

4.2 Round-Trip Times (RTT)

Results from the traceroute measurements show that inter-continental traffic from Africa to African universities experiences an average RTT of 300 ms, in contrast to an RTT of 139 ms for traffic that did not leave the continent (intra-Africa). Inter-continental RTT has a standard deviation of 120ms while the intra-Africa RTT has a standard deviation of 82 ms. The standard deviations in both cases show an overlap in latencies obtained by the two sets of traffic, indicating that, for certain source-destination pairs, better performance is obtained by using exchange points outside Africa, while for other traffic, better performance is obtained when traffic is routed within the continent. Figure 2 shows a scatter plot of the RTTs for both cases and shows an overlap of performance between the two sets of traffic.

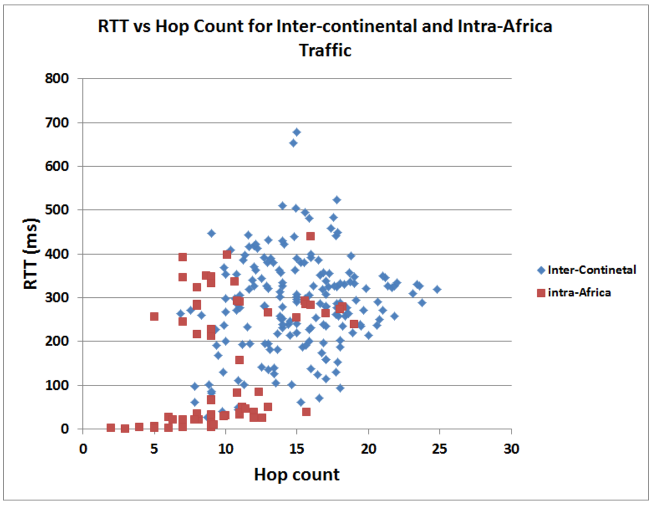

Fig. 2. Scatter plot showing the distribution of RTT for intra-Africa traffic and inter-continental traffic.

4.3 AS-Level Peering

Using Maximind's database and the WHOIS database, IP hops in the traced paths are mapped to their corresponding Autonomous System Number (ASN). Each ASN is represented as a node in a graph, and edges are created where there is direct link between ASes. The observed interconnection among the African NRENs is largely through ASes that peer at global IXPs in Europe, such as Cogent Communications(ASN 174), TATA (ASN 6453), Level3 (ASN 3356) and Seacom (ASN 37100). These ASes peer mainly at the London Internet Exchange (LINX), Amsterdam Internet Exchange (AMS-IX), and Frankfurt Internet Exchange. There is also high connectivity to the UbuntuNet Alliance (ASN 36944) peering at LINX and AMS-IX.

The AS-level graph shows that there is minimal peering within Africa, as most of the probed networks have direct AS-level paths to the ASes in Europe, resulting in high node degree (ranging from 18 to 32) at the European ASes, and low node degree (between one and three) for Africa-based ASes. The AS-level graph indicates an overall network diameter of 8, an average path length of 3.37, and an average clustering coefficient of 0.180. On the other hand, a path length of 3.37 would suggest a densely interconnected AS-level topology, a low clustering coefficient of 0.180 suggests that, overall, the AS-level topology is sparse. This finding is consistent with the previously reported low peering among Africa based ASes and a high connectivity to the global IXPs [3]. Figure 3 shows the degree distribution of the AS-level topology.

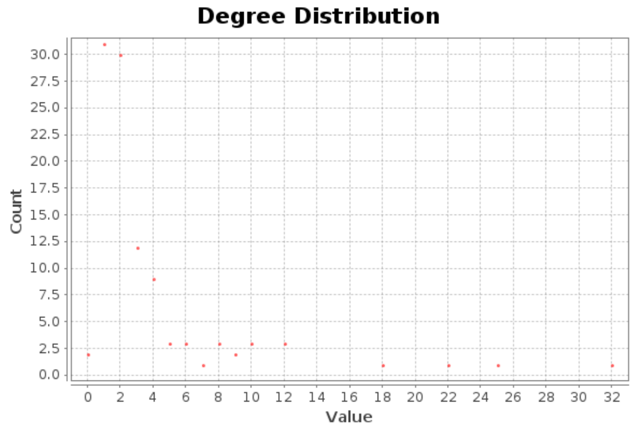

Fig. 3. AS-level node degree distribution.

5 Discussion

5.1 Effect of inter-continental latency

Round-trip times as well as geographical distribution of routes for inter-NREN traffic in Africa reveals a lack of peering among the African service providers. Despite a growing number of national IXPs in Africa, there is still limited interconnection at regional IXPs to facilitate cross-border peering. As a result, inter-NREN traffic in sub-Saharan Africa follows circuitous inter-continental routes, resulting in much higher latencies compared to traffic exchanged within the continent. The physical length of the transmission medium travelled by the packets (linearised path) has a significant contribution to the internet path delay. This delay is relative to the signal propagation time, which is bounded by the speed of light, i.e. the speed cannot be greater than the speed of light (299.792 km/ms) [16]. In practice, however, the speed of light in a fibre optic cable is slower at about 200 km/ms.

We observed that traffic to countries with no direct fibre connection from the vantage points experienced much higher latencies. For example, traffic from Johannesburg to Malawi is routed first through London, then Maputo Mozambique, before being forwarded to Lilongwe, and experiences RTT of around 380 ms. Traffic from southern Africa to southern Africa routed via London, covers a distance of roughly 30,000km, hence a round trip of about 60,000km (the West Africa Cable System fibre-optic cable from Cape Town to London is about 14,530 km long). At optic speed of 200km/ms in fibre cable, this translates to a minimum RTT of about 300 ms. The observed RTT for this round-trip is around 370 ms, which suggests that about 80% of the RTT in this case is due to the

[0] http://afterfibre.net/

linearlized distance factor alone. The effect of inter-continental routing is demonstrated further in the case of universities within the same country that achieve remarkably different RTTs due to where their traffic is routed. For example, in Kenya, one university has its traffic from Johannesburg South Africa routed via A msterdam (AIMS-IX), then back to South Africa (CINX, Cape Town) before being forwarded to Kenya, and achieves an RTT average of about 400 ms. In comparison, another university in the same country (Kenya) has a direct logical link from Johannesburg to Kenya and achieves an RTT of only 80 ms.

Lower latencies are observed between countries where direct fibre cable connection and peering exists. For example, the Zambia NREN (ZamREN) recently established a direct connection to the Johannesburg IXP where it is peering with the UbuntuNet alliance, and using this link, traffic from South Africa to ZamREN experiences a low latency of 55 ms. Other locations achieving low RTT from the South African vantage point are in the neighbouring countries such as Mozambique (45 msec) and Namibia (80 msec). The countries have direct fibre links to South Africa. Futhermore, where functional NRENs are present and traffic is routed locally within national IXPs, much lower latency is observed. For example, within South Africa, members of South African NRENs (SanRen) are linked though a fibre-optic backbone and exchange traffic locally at national IXPs - the Johannesburg Internet Exchange (JINX) and the Cape Town IXP (CINX). In the experiments, traffic from the vantage point located within the SanRen, to other SanRen members achieved an average RTT of 20 msec.

Further more, vantage points experience varying degrees of delay on the inter-continental link depending on their proximity to their remote gateways (i.e., the first hop outside Africa). For example, Morocco, which is closest to Europe among all the vantage points used in this study, had an average RTT of 39 ms to its remote gateways, whereas the South African vantage point had an average RTT of 170 ms to its remote gateways. The highest latency for reaching the remote gateway was from the Rwandan vantage point, which was at 199 ms. In general, vantage points with a higher inter-continental link latency obtained higher overall RTT, which shows that the high latency for inter-continental traffic is largely due to the delay on the egress and ingress links. On average, the RTT from the vantage points to the remote gateways is 150 ms, almost half the average RTT obtained for inter-continental traffic. This is also more than the average RTT of 139 ms obtained for traffic exchanged within Africa.

5.2 Improving the Routing and Traffic Engineering Environment

One solution that is used to reduce latencies and the circuitous route problem is the use of Internet Exchange Points (IXPs) [17], which can enable ASes to establish mutual peering agreements for direct exchange of local traffic. For end-to-end communication, the challenge is that selection of Internet route paths is mostly influenced by routing policies that are optimized for interests of individual autonomous systems. Such policies do not provide guarantees for optimal end-to-end connections and sometimes result in packets not traversing the shortest paths to their destinations. To deal with this challenge, inter-domain traffic engineering (TE) techniques aim to optimize resource utilization and internetwork performance through mechanisms that identify and dynamically use optimal low-latency paths. This requires routing systems that are able to learn and make use of inter-domain topology information in choosing routing paths.

Solving the problem of circuitous routes also requires effective TE techniques. With multiple end to end paths, it is possible to optimise usage of inter-continental links as well as cross-border terrestrial links. One challenge however, is that the inter-domain routing protocol, BGP, being a single-path system, is inflexible with regard to multi-path routing and inter-domain TE. On the other hand, novel hierarchical routing protocols such as the Locator/Identifier Separation Protocol (LISP) [18], provide new opportunities for inter-domain traffic engineering by allowing networks to announce multiple gateways (route locators) for reachability, making it possible to have multi-path routing. Furthermore, Software Defined Networking (SDN) protocols allow for dynamic and remote configuration of traffic forwarding paths. These protocols can make possible collaborative routing and dynamic peering among edge networks, by allowing ASes to exchange multiple routing paths and to dynamically respond to varying network path conditions and QoS requirements. Furthermore, networks could employ remote peering strategies, dynamically selecting and peering at optimal open exchange points. For example, SEACOM [1], a major fibre cable operator in Africa, provides open peering points that allow networks to peering remotely.

6 Conclusion

This paper has looked at the effect of circuitous inter-continental IP routes on the performance of traffic exchanged between education and research institutions in Africa. Traceroute probes on a sample of 95 African IP addresses has shown that on average 75% of the traffic originating in Africa and destined for African universities traverse links outside the continent, thereby performing with a latency that is more than double that of the intra-Africa traffic. Future work will evaluate how this actually affects the performance of the African Internet, and will undertake further monitoring of the actual Internet traffic to determine how much traffic is actually exchanged between the African NRENs.

Acknowledgement: We are thankful to the Hasso-Plattner-Institute (HPI) for their financial support towards this work.

References

1. M. Ahmad and R. Guha, "Evaluating end-user network benefits of peering with path latencies," in *Computer Communications and Networks (ICCCN), 2012 21st International Conference on*, pp. 1–7, July 2012.

[1] http://seacom.mu/services-solutions/

2. B. Barry, C. Barton, V. Chukwuma, L. Cottrell, U. Kalim, M. Petitdidier, and B. Rabiu, "egy-africa: better internet connectivity to reduce the digital divide," in *IST-Africa, 2010*, pp. 1–15, IEEE, 2010.
3. 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 *Passive and Active Measurement Conference 2014*.
4. R. Steiner, A. Tirivayi, N. Tirivayi, M. Jensen, P. Hamilton, J. Buechler, A. Jeffries, U. H. Emdon, I. A. Ibrahim, *et al.*, "Promoting african research and education networking," *International Development Research Center, January*, 2005.
5. Y. Shavitt and U. Weinsberg, "Quantifying the importance of vantage point distribution in internet topology mapping (extended version)," *Selected Areas in Communications, IEEE Journal on*, vol. 29, pp. 1837–1847, October 2011.
6. M. Ahmad and R. Guha, "Internet exchange points and internet routing," in *Network Protocols (ICNP), 2011 19th IEEE International Conference on*, pp. 292–294, Oct 2011.
7. B. Donnet and T. Friedman, "Internet topology discovery: a survey," *Communications Surveys Tutorials, IEEE*, vol. 9, pp. 56–69, Fourth 2007.
8. R. Govindan and H. Tangmunarunkit, "Heuristics for internet map discovery," in *INFOCOM 2000. Nineteenth Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings. IEEE*, vol. 3, pp. 1371–1380 vol.3, Mar 2000.
9. J. Gilmore, N. Huysamen, and A. Krzesinski, "Mapping the african internet," in *Proceedings Southern African Telecommunication Networks and Applications Conference (SATNAC),(Sept 2007), Mauritius*, 2007.
10. M. Luckie, "Scamper: a scalable and extensible packet prober for active measurement of the internet," in *Proceedings of the 10th ACM SIGCOMM conference on Internet measurement*, pp. 239–245, ACM, 2010.
11. B. Augustin, T. Friedman, and R. Teixeira, "Multipath tracing with paris traceroute," in *End-to-End Monitoring Techniques and Services, 2007. E2EMON '07. Workshop on*, pp. 1–8, Yearly 2007.
12. B. Augustin, T. Friedman, and R. Teixeira, "Measuring load-balanced paths in the internet," in *Proceedings of the 7th ACM SIGCOMM conference on Internet measurement*, pp. 149–160, ACM, 2007.
13. B. Augustin, T. Friedman, and R. Teixeira, "Measuring multipath routing in the internet," *IEEE/ACM Transactions on Networking (TON)*, vol. 19, no. 3, pp. 830–840, 2011.
14. Y. Shavitt and E. Shir, "Dimes: Let the internet measure itself," *ACM SIGCOMM Computer Communication Review*, vol. 35, no. 5, pp. 71–74, 2005.
15. Y. Shavitt and N. Zilberman, "A geolocation databases study," *Selected Areas in Communications, IEEE Journal on*, vol. 29, pp. 2044–2056, December 2011.
16. R. Landa, J. Araujo, R. Clegg, E. Mykoniati, D. Griffin, and M. Rio, "The large-scale geography of internet round trip times," in *IFIP Networking Conference, 2013*, pp. 1–9, May 2013.
17. N. Chatzis, G. Smaragdakis, A. Feldmann, and W. Willinger, "There is more to ixps than meets the eye," *ACM SIGCOMM Computer Communication Review*, vol. 43, no. 5, pp. 19–28, 2013.
18. K. Li, S. Wang, and X. Wang, "Edge router selection and traffic engineering in lisp-capable networks," *Communications and Networks, Journal of*, vol. 13, pp. 612–620, Dec 2011.