IST-Africa 2016 Conference Proceedings
Paul Cunningham and Miriam Cunningham (Eds)
IIMC International Information Management Corporation, 2016
ISBN: 978-1-905824-55-7

Efficient Topology Discovery for African NRENs

Roslyn SANBY[1], Hussein SULEMAN[2], Josiah CHAVULA[3]
Department of Computer Science, University of Cape Town,
Private Bag X3, Cape Town, 7700, South Africa
E-mail: [1]sanbyra@gmail.com; [2]hussein@cs.uct.ac.za; [3]jchavula@cs.uct.ac.za

Abstract: The proliferation of under-sea and terrestrial fibre optic cables in Africa is leading to an evolution of the interconnection among Africa's National Research and Education Networks (NRENs). To fully monitor and track this evolution, NREN stakeholders and researchers need to continuously run active network measurements to map the inter-NREN topology. However, active topology measurement tools, such as Traceroute, introduce additional traffic in the network, which maybe undesirable for networks that have limited bandwidth. Furthermore, many network routers are configured to block/drop network probing packets, thereby making the task of topology discovery a lot more difficult. The paper aims to show that by carefully analysing Internet path overlaps, significant reductions can be made on the number of packets required to probe the topology. Furthermore, by using multiple probing protocols and vantage points, it is possible to alleviate the impact of blocked packets. The paper reports on a Traceroute study carried out on African NRENs, which shows that by jointly using different probing protocols, and taking into consideration path overlaps, it is possible to map all probe destinations, while achieving a 65% reduction in the number of probe packets sent.

Keywords: Network measurements; NRENs; Traceroute; RIPE Atlas; latency; UbuntuNet Alliance; Africa; path diversity; topology discovery.

1. Introduction

A National Research and Education Network (NREN) is a mesh of interconnected networks that supports the needs of education and research communities in a country [4]. Among other goals, NRENs aim to reduce latencies between educational institutions [4]. At a regional level, the UbuntuNet Alliance connects NRENs in southern and eastern Africa, and as of September 2014, consisted of 15 NRENs [4].

Through projects like the AfricaConnect, the UbuntuNet Alliance is aiming towards keeping African NRENs' traffic within Africa. In order to help monitor traffic as the African NREN logical topology is evolving, there is a need for a platform that can continually collect and display accurate data about NREN topologies in Africa. By generating and displaying accurate topological data, as well as latencies experienced by African NRENs, such a platform would help researchers and NREN decision makers to monitor and evaluate how and where their interconnectivity can be improved.

Active measurements are often used for collecting data relating to network topologies, with Traceroute being the most commonly used method for deducing the paths through which traffic is routed between networks [1] [2] [7] [8] [9]. However, active topology measurements are costly in the sense that they introduce additional traffic into the topology, which may not be desirable in networks that are resource constrained. Prominent Internet topology measurement infrastructures, such as Ripe Atlas, assign a cost and limit to the number and rate of probe packets that one can introduce into the platform. Furthermore, network probing packets are generally blocked by routers in the network, which entails that it may not be possible evaluate all the possible network paths. Even where probe packets are allowed, the presence of multiple alternate paths and load balancing renders it difficult to evaluate all the alternate paths.

This paper aims to show that Traceroute data can be collected reliably and efficiently for the purpose of discovering the logical topology of the African NRENs. In this regard, reliability implies the ability to discover an accurate and complete representation of the topology. Efficiency on the other hand, refers to minimizing the number of probe packets required to obtain a complete topology. This study intends to demonstrate this by using the Ripe Atlas Internet measurement infrastructure to conduct Traceroute probes from 12 vantage points located within UbuntuNet NRENs, targeted targeting 50 IP addresses located in various research and education institutions, also within the UbuntuNet Alliance zone.

2. Methodology

With the key objective being to achieve a reliable and efficient mechanism for collecting Traceroute data for the discovery of the UbuntuNet Alliance's topology, a distributed network measurement method was used. Distributed network probing uses many vantage points to get a more accurate view of a network topology. Furthermore, varying location of vantage points further increases the accuracy and completeness of the discovered topology.

Many distributed topology probing platforms are in existence, but unfortunately, many of them have very few vantage points on the African continent, let alone inside the NRENs. Some of the prominent platforms globally are Archipelago [3] [5] [6], DIMES [11], iPlane [12] and RIPE Atlas [10]. On the African continent, RIPE Atlas [10], has over 100 active probes within the UbuntuNet countries, although only about 19 of these probes were seen to be hosted inside NREN networks. Atlas is a platform that makes use of thousands of active probes around the world to measure Internet connectivity and reachability in real-time. These probes are small USB-powered hardware devices (attached to an Ethernet port) that conduct measurements, such as Ping, Traceroute, DNS and SSLcert, and relays the data to the RIPE Network Coordination Centre (NCC). Custom measurements can be sent from any probe to any IP address, and therefore to any NREN, allowing the collection of data that is required to study routes between African NRENs. The RIPE Atlas platform is thus very advantageous for researching the topology of African NRENs.

2.1 Topology Measurements

To run topology measurements on the African NRENs topology, 14 Atlas probes were selected from various NRENs within the UbuntuNet Alliance. Although some institutions were seen to host several Atlas probes, a maximum of two were used from each ASN. The NREN institutions and the AS numbers in which RIPE Atlas probes were selected are shown in Table 1.

Table 1. List of AS Numbers and Their Names

Name	AS number	Country
RENU	327687	Uganda
KENET	36914	Kenya
iRENALA	37054	Madagascar
SudRen	37197, 33788	Sudan
TENET	2018	South Africa
University of Cape Town	36982	South Africa
Rhodes University	37520	South Africa

In order to increase the reliability and accuracy of the topology data, multiple traceroute measurements were performed using 3 different protocols; TCP, UDP, and ICMP. The measurements were conducted from the selected vantage points to a set of 50 IP addresses, each one representing a university or a research institution within each NREN in the UbuntuNet Alliance. Given the 14 Atlas probes as vantage points, and the 50 IP addresses as probe destinations, the experiment had in total 700 source-destination pairs. The probes and the destinations used for the Traceroute measurements are depicted in Figure 1. Probes are represented by diamonds while destinations are represented by circles. The colour of the circle indicates the AS in which the IP address is located.

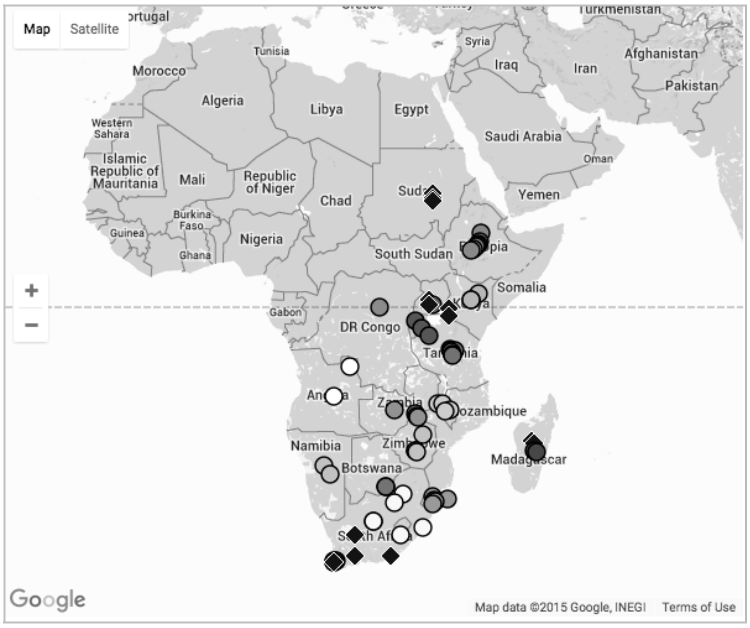

Figure 1. Visualisation of Probes and Destinations for Traceroute Measurements

2.2 Reducing Redundancy in Measurements

In order to reduce probing redundancy and enhance efficiency in performing measurements, overlapping paths had to be identified. For example, if paths from a vantage point to a set of destinations share the first N hops, then the traceroute for all but one subsequent measurement were configured to skip those first N hops. Similarly, if a set of measurements to a particular destination converged at a hop N, then for subsequent measurements, all but one traceroute measurement were configured to trace only up to hop N towards the destination.

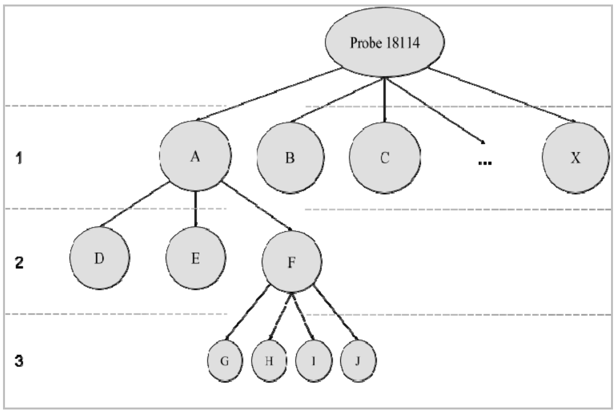

Figure 2. Example n-ary tree Created to Find Overlapping Paths at Beginning of Traceroute Measurements

In the illustration given in Figure 2, measurements from probe 18114 going to G, H, I and J all go through A and F. So, the first hop for measurements from 18114 to H, I and J, for example, was incremented to 2 (hop F). A record was made of the measurement which had the same beginning path to allow for reconstruction of the path information.

Paths which overlapped near the destination were found by comparing paths from different vantage points (Atlas probes) to the same destination. In the example in Figure 3, measurements from X and Y going to destination A both go through C and B. The Traceroute could therefore be sent to C instead of A. A record was also made of the measurements that had the same end path so that full path information could be reconstructed.

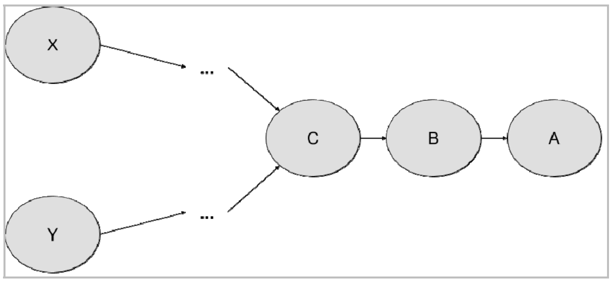

Figure 3. Example of Overlapping Paths at the End of Two Traceroute Measurements

In total, eighteen measurements were conducted for each source-destination pair. These measurements consisted of six ICMP-based measurements, six TCP-based measurements and six UDP-based measurements. The measurements for each protocol consisted of three full measurements and the three partial measurements, where either the first hop or last hop was adjusted taking into consideration the overlapping paths.

3. Results and Discussion

3.1 Destinations Reached

Of the three protocols, measurements configured with TCP had the highest of measurements that reached the probe destination. Sixty-three percent (63%) of the TCP-based measurements reached their destinations, whereas only 50% of the measurements configured with ICMP, and 26% of the UDP-based measurements reached probe destinations. All the destinations were reachable from at least one TCP-based measurement, whereas only 70% and 56% of all the destinations were reachable via ICMP and UDP measurements respectively. This indicates a significant amount of blocking for network probing traffic within the African NRENs, especially for ICMP and UDP based probes. This is not surprising as network administrators employ probe packet filtering as a mechanism for avoiding network attacks and abuse (e.g denial of service attacks). However, by combining the different probing protocols, all of the destinations were reached.

This demonstrates the importance of not only using multiple vantage points, but also employing different probing protocols to ensure that more probe destinations are reached. Using only a single protocol would leave significant gaps in the topology discovered, as many different networks block different probing packets based on the protocol used. On the other hand, by using multiple protocols and multiple vantage points, a more complete topology can be discovered.

3.2 Path Diversity

When there is more than one unique path from a source to a destination, there is path diversity. For an internetwork topology, path diversity can be expressed as the average number of unique paths between a source and a destination.

Results from probing the African NRENs topology has shown, as is seen in Figure 4, that about 79% of the source-destination pairs have more than one IP path, with an average path diversity of 3.

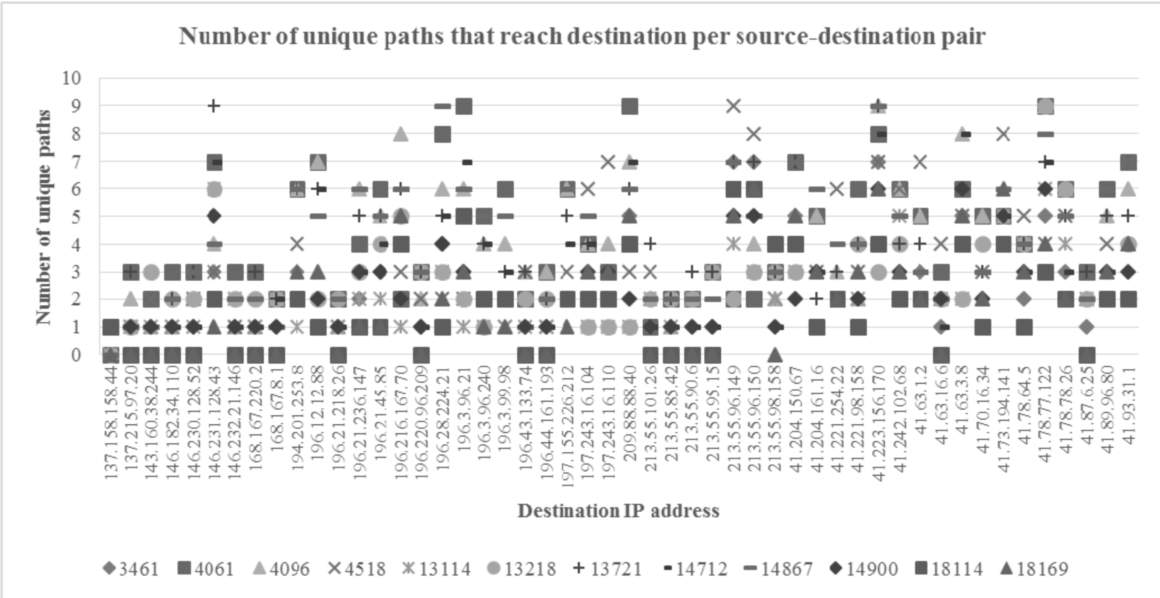

Figure 4. Graph of Number of Unique Paths from a Source to a Destination that Reached the Destination

In the example Figure 5 below, there are 9 unique IP paths from Atlas probe 14867 to destination 196.28.224.21.

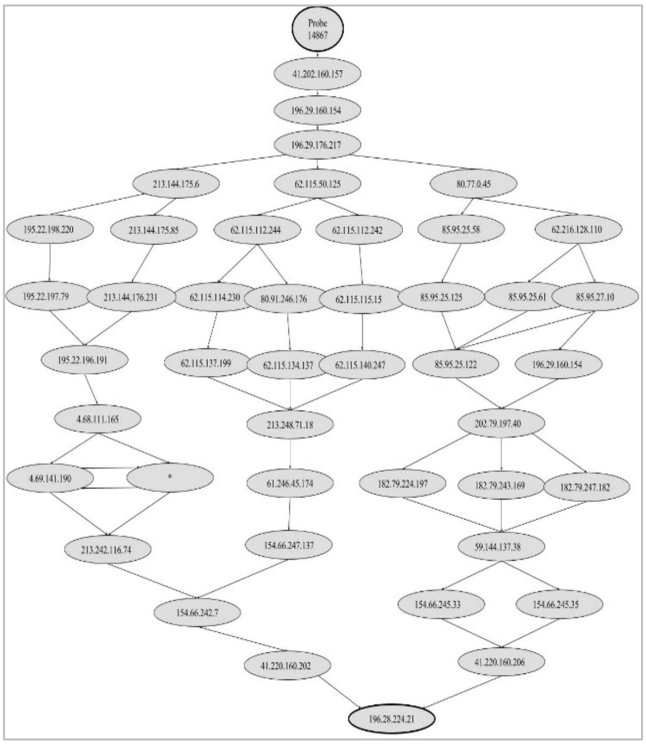

Figure 5. Example Showing the Path Diversity from Probe 14867 to Destination 196.28.224.21

It is worth noting that there is an increase in paths discovered when different Traceroute protocols are used. After running 18 traceroute measurements for each of the 600 source-destination pairs, a total of 3352 IP paths are observed. ICMP-based measurements alone have 1168 paths; TCP-based measurements have 1029 paths, and UDP-based measurements have 1443 paths. Many of the paths discovered were common to the different protocols, as depicted in the Venn diagram in Figure 6.

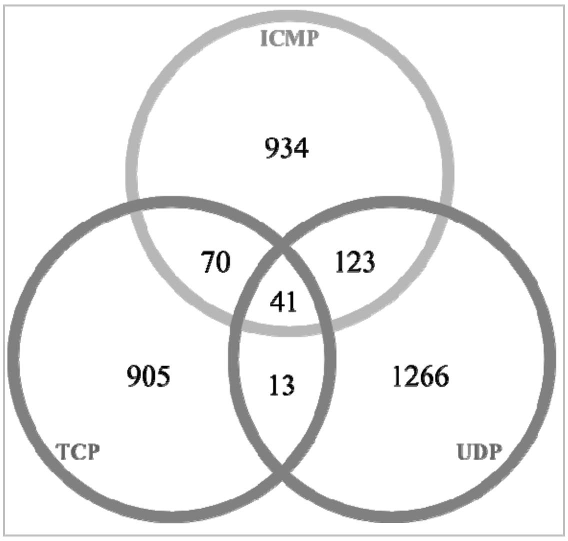

Figure 6. Venn Diagram Showing the Number of Unique Paths Discovered by Each Protocol

3.3 Round Trip Times (RTTs)

As shown in the Cumulative Distribution Frequency graphs in Figure 7, the average latency for each protocol is roughly the same at about 293ms and the cumulative frequency graphs generally follow the same curve for all three protocols. As seen in Figure 7, there are more low latencies for TCP-based measurements than for UDP and ICMP-based measurements. About 40% of the latencies for TCP-based measurements are less than 100ms, whereas only about 25% of the ICMP and UDP-based measurements are less than 100ms. This observation is analogous to the level of blocking associated with the three protocols, as reported in section 3.1, where TCP probes are less likely to be blocked that ICMP and UDP. The lower TCP latencies would also suggest that ICMP and UDP Traceroute traffic is generally given lower priority compared to TCP.

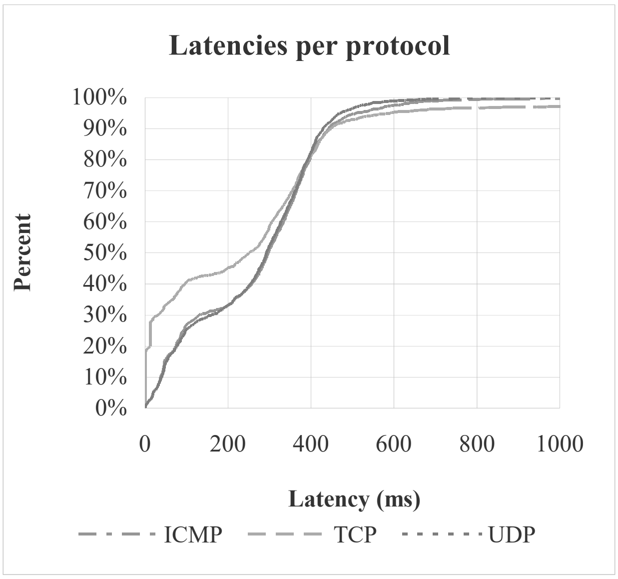

Figure 7. Cumulative Distribution Frequency graph showing distribution of latencies per protocol for full measurements

3.4 – Efficiency

By skipping parts of the overlapping paths during probing, it was noted that the partial paths measurements had on average 6 hops per Traceroute measurement, compared to an average of 17 hops per Traceroute in full path measurements. This represents a 65% reduction in the number of packets that needed to be sent by Traceroute for each destination. The significant reduction in the number of hops indicates that there are a lot of path overlaps. This is not surprising considering that the IP addresses and the vantage points used in the experiment are all located in institutions that somehow share Internet paths through their respective NRENs. For example, traffic from SanREN institutions in Cape Town, to universities in KENET, would generally follow the same path between the TENET gateway in Cape Town and KENETs gateway in Mombasa.

4. Related Work

Chavula et al. [13] used the CAIDA Archipelago platform to conduct logical topology mapping for African NRENs, making use of CAIDA's five vantage points in Africa. Gupta et al. [14] used Traceroute and BGP routing tables to map Africa's logical topology. Aben et al. [15] used the RIPE Atlas platform to measure and analyse Internet traffic paths in Sweden. Aben et al. [15] had access to hundreds of probes around Europe to investigate traffic traversing IXPs.

5. Conclusion

This paper presented work aimed at efficiently collecting topology mapping data for the UbuntuNet Alliance network. The focus was on how to efficiently and reliably collect Traceroute data by making use of three protocols; TCP, ICMP, and UDP. In the experiments, although not all destinations were reached for each source-destination pair, the use of multiple vantage points ensured that all destinations were. Furthermore, using multiple protocols from each vantage point ensures that more paths can be discovered between vantage points and destinations, which increases the completeness of the topology discovered.

This paper further presented a mechanism for analysing overlapping paths and taking into consideration the overlaps to reduce redundancy in the path probing. By skipping overlapping parts of the paths in Traceroute probes, the number of probe packets needed to probe the topology was reduced by 65%. This result highlights an important mechanism for reducing probe packets in persistent topology measurement experiments, especially in networks where bandwidth is limited.

Future work aims to streamline the topology discovery system to allow for new data to be collected, stored, aggregated and displayed in a graphical visualisation. This would help researchers, network managers and other interested parties in planning new routing policies based on an accurate and up-to-date set of data.

References

[1] Augustin, B., Cuvellier, X., Orgogozo, B., Viger, F., Friedman, T., Latapy, M., Magnien, C., and Teixeira, R. Avoiding Traceroute anomalies with Paris Traceroute. in *6th ACM SIGCOMM conference on Internet measurement,* (New York, USA, 2006), ACM Press, 153-158.

[2] Cheswick, B., Burch, H., and Branigan, S. Mapping and Visualising the Internet. in *USENIX Annual Technical Conference,* (Berkeley, USA, 2000), USENIX Association, 1-1.

[3] Claffy, K., Hyun, Y., Keys, K., Fomenkov, M., and Krioukov, D. Internet mapping: from art to science. in *Conference for Homeland Security, 2009. CATCH'09. Cybersecurity Applications & Technology,* (Washington, DC, USA, 2009), IEEE, 205-211.

[4] Fryer, T. Research and education networks around the world and their use. in Segunda Conferencia de Directores de Tecnologia de Informacion, (Lima, Peru, 2012), 1-13, http://hdl.handle.net/10786/847

[5] Hyun, Y. Archipelago measurement infrastructure, 2015. Retrieved April 24, 2015 from Centre for Applied Internet Data Analysis: http://www.caida.org/projects/ark/.

[6] Hyun, Y. CAIDA Monitors: The Archipelago Measurement Infrastructure, 2015. Retrieved April 20, 2015 from Centre for Applied Internet Data Analysis: http://www.caida.org/data/monitors/monitor-map-ark.xml.

[7] Luckie, M., Hyun, Y., and Huffaker, B. Traceroute probe method and forward IP path inference. in *8th ACM SIGCOMM conference on Internet measurement,* (New York, USA, 2008), ACM Press, 311-324.

[8] Mao, Z.M., Rexford, J., Wang, J., and Katz, R.H. Towards an accurate AS-level Traceroute tool. in *2003 conference on Applications, technologies, architectures, and protocols for computer communications,* (New York, USA, 2003), ACM Press, 365-378.

[9] Pansiot, J., Mérindol, P., Donnet, B., and Bonaventure, O. Extracting intra-domain topology from mrinfo probing. in *11th International Conference, PAM,* (Zurich, Switzerland, 2010), Springer Berlin Heidelberg, 81-90.

[10] Ripe Atlas, 2010. Retrieved April 24, 2015 from Ripe Network Coordination Centre: https://atlas.ripe.net.

[11] Shavitt, Y., and Shir, E. DIMES: Let the Internet measure itself. *ACM SIGCOMM Computer Communication Review 35* (5), 71-74.

[12] Shavitt, Y., and Weinsberg, U. Quantifying the importance of vantage points distribution in internet topology measurements. in *INFOCOM 2009, IEEE,* (Rio de Janeiro, 2009), IEEE, 792-800.

[13] Chavula, J., Feamster, N., Bagula, A. and Suleman, H., 2014. Quantifying the Effects of Circuitous Routes on the Latency of Intra-Africa Internet Traffic: A Study of Research and Education Networks. In e-Infrastructure and e-Services for Developing Countries (pp. 64-73). Springer International Publishing.

[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. in 15th International Conference, PAM, (Los Angeles, CA, USA, 2014), Springer International Publishing, 204-213.

[15] Aben, E. Measuring countries and IXPs with RIPE Atlas, 2015. Retrieved 3 August 2015 from Ripe Network Coordination Centre: https://labs.ripe.net/Members/emileaben/measuring-ixpswith-ripe-atlas.