Analysing and Visualising National Research and Education Networks in Africa: A Literature Review

ROSLYN SANBY, UNIVERSITY OF CAPE TOWN

ABSTRACT

National Research and Education Networks (NRENS) aim to provide IP connectivity between educational institutions to improve education and research. In Africa, about 75% of the network traffic between African NRENs follow circuitous routes through Europe and other areas outside the continent. In order to find out what the effect of these circuitous routes are, research needs to be done into discovering the topology of the Internet, especially in Africa and what traffic is being exchanged between and within the NRENs in Africa. Multiple vantage points in a variety of locations help with limiting the bias in the topology and gives more accurate data. A visualisation tool should also be developed to aid in analysis of the information collected.

Categories and Subject Descriptors

C.2.1 [Computer-Communication Networks]: Network Topology

General Terms

Design, Measurement

Keywords

National Research and Education Networks, Autonomous Systems, Distributed Network Probing, Visualisations

1. INTRODUCTION

A National Research and Education Network (NREN) is a mesh of interconnected networks that support the needs of education and research communities in a country [Fryer. 2014]. NRENs aim to reduce latencies between educational institutions in order to facilitate better research and communication [Fryer. 2014]. Some well-known NRENs include GÉANT and RedCLARA [Fryer. 2014]. In Sub-Saharan Africa, most of the NRENs are inter-connected through the UbuntuNet Alliance which, as of September 2014, consisted of 15 NRENs [Fryer. 2014].

However, the objective of reducing latencies between African NRENs is facing challenges. Currently, according to Chavula et al. [2014], about 75% of the traffic between NRENs in Africa gets routed outside the continent, often being boomeranged through Europe, resulting in high latencies [Chavula, et al. 2014, Gupta, et al. 2014]. This could be due to lack of direct physical links or lack of local peering. Furthermore, more research needs to be done, though, on the level of traffic that is being exchanged between these networks [Chavula, et al. 2014] so as to improve communication and research efforts in Sub-Saharan Africa.

Surveys have been conducted by Motamedi et al. [2013] and Donnet and Friedman [2007], discussing different mechanisms for discovering the Internet topology at different levels - namely the IP interface, the router, the Point of Presence (PoP) and the Autonomous System (AS) level. Other methods of probing the Internet have also been researched, all based on Van Jacobson's original implementation. Bush et al. [2009] propose dual probing, however Augustin et al. [2006] suggest the use of Paris traceroute to avoid anomalies seen with traceroute. Scamper, which is fully described in the paper by Luckie et al. [2008], is used by Chavula et al. [2014] to collect their data, and this software makes use of the Paris traceroute. A paper by Bush et al. [2009] looks at reachability of the Internet, questioning the methods that are used to understand reachability and looking at a method of traceroute called dual probing.

This project will help researchers see what information is being sent between and within African NRENs and how the circuitous routes through Europe affect the NRENs in Sub-Saharan Africa. Information will be collected from log data for traffic profiling, as well as from probing of the Internet to get a better understanding of the topology of the Internet in Africa.

As the data collected will be used for future research, a visualisation of the data will be very advantageous to researchers in this field and to help NRENs to see what information is being passed within and between them in order to aid with future research. There have been many papers looking into visualisations of the Internet as this topic has been researched for many years. Bhardwaj and Singh [2015] list twenty different visualisations that have been designed to map the topology of the Internet, but this project will be focusing on visualisations that map the PoP and AS levels and that display dynamic data.

The rest of this paper is organised as follows. Section 2 discusses papers pertaining to National Research and Education Networks. Section 3 analyses papers that discuss and evaluate techniques for discovering the topology of the Internet. As most research in this field has looked into the topology of the Internet at an AS or PoP level, these levels will be the focus of discussion. However, the other two levels (namely IP Interface and Router) will also be looked at. Section 4 looks at distributed probing platforms and advantages of using more than one vantage point when collecting data and Section 5 looks at various methods of visualising the topology of the Internet and networks.

2. NATIONAL RESEARCH AND EDUCATION NETWORKS (NRENS

A National Research and Education Network (NREN) is a mesh of interconnected networks that support the needs of education and research communities in a country [Fryer. 2014]. The UbuntuNet Alliance is an association of NRENs in Southern and Western Africa that consisted of 15 NRENs in September 2014 [Fryer. 2014]. TENET in South Africa and KENET in Kenya are NRENs that are part of this alliance. The goals of an NREN, apart from providing IP connectivity, include providing point-to-point connectivity or bandwidth-on-demand. This is done well with other continental NRENs, such as GÉANT in Europe and RedCLARA in Latin America. However, these goals are not being met as well in Africa [Fryer. 2014].

Internet Exchange Points (IXPs) are infrastructures where Internet traffic is exchanged, allowing networks to connect directly and reduce latency, bandwidth and cost. The purpose of establishing NRENs in an area is to take advantage of these IXPs and peerings to improve the communication between educational institutions. But this objective has not yet been reached in African NRENs as the traffic is being exchanged in Europe to reach other ASes in Africa [Chavula, et al. 2014, Gupta, et al. 2014]. It has been found that about 75% of traffic originating in Africa and destined for other African NRENs travel out of the continent to reach their destination [Chavula, et al. 2014].

Both Chavula et al. [2014] and Gupta et al. [2014] found that these circuitous paths result in higher latencies. Gupta et al. [2014] focused their research on the connectivity at JINX in Johannesburg and KIXP in Nairobi, which are major local Internet exchange points (IXP) in Africa, while Chavula et al. [2014] researched NRENs in Africa, gathering information from five vantage points targeting 95 universities. Even though the research was done from different perspectives, the results seem to be the same. NRENs in Africa are currently not meeting their objectives of reduced latencies, due to these circuitous Internet paths.

Gupta et al. [2014] give two suggestions to avoid these circuitous paths and reduce latencies. One is to add more local cache servers in Africa but this is not beneficial for local content hosted in Africa. If one looks at content hosted overseas, such as Google or BBC, more local cache servers would help reduce latencies but not so for local content being sent within Africa. An example of this content would be inter-university virtual classrooms where students learn online [Harper, et al. 2004]. The second suggestion is to add more peering relationships in Africa. This suggestion could be helpful as it would create more local links and help NRENs avoid circuitous routes [Shavitt and Shir. 2005].

3. INTERNET TOPOLOGY DISCOVERY

Measurements of the Internet can be performed in the control or data planes. As explained in Motamedi et al. [2013], measurements in the control plane consist of information about Internet routing often found in Border Gateway Protocol (BGP) tables. On the other hand, measurements in the data plane look at which paths packets travel along to discover the reachability of the internet as well as Round Trip Time (RTT) of packets. These measurements can be either active or passive. The difference between active and passive measurements are that active measurements tools, like traceroute, send packets into the network and collect data from the response, whereas passive measurements, like BGP monitors, collect information that is already flowing over the wire [Motamedi, et al. 2013]. Active measurements are often used for collection data to discover topologies whereas passive measurements, like logs, are used for traffic profiling.

Topology visualisation can be done on four different levels – Internet Protocol (IP) Interface, Router, Point of Presence (PoP) and Autonomous System (AS) levels [Donnet and Friedman. 2007, Motamedi, et al. 2013]. NRENs mainly operate on the AS level but could constitute several PoPs. Topology discovery at the PoP level provides information and limitations about latencies between PoPs, which helps with understanding the geographical properties of Internet paths, such as where ASes can connect and the coverage of ASes [Motamedi, et al. 2013]. Therefore the project will look into methods of discovery on the AS and PoP levels.

Autonomous Systems (ASes) are privately managed networks, identified by a unique 16-bit AS number, which are all interconnected making up the Internet [Donnet and Friedman. 2007, Motamedi, et al. 2013]. To collect information for the AS level, data is collected from BGP tables, data from traceroute measurements and Internet Routing Registries (IRR). This indicates measurements in the control plane.

A Point of Presence (PoP) is a collection of routers belonging to one AS [Motamedi, et al. 2013]. There are three main methods to collect data at the PoP level. Firstly, one can aggregate data from traceroute measurements to identify PoPs. Secondly, one can obtain delay estimates from ping measurements. Finally, one can retrieve information from websites where ISPs have published their data. However, although this will provide more accurate data than measurements, this technique is not always reliable as the information could be outdated [Motamedi, et al. 2013].

Chavula et al. [2014] used both the AS and PoP levels to map their results and used Scamper as a measurement tool to collect these results. Scamper is an implementation of four probing techniques [Luckie, et al. 2008]. These include: Van Jacobson's original traceroute tool; using Internet Control Message Protocol (ICMP) echo request probes instead of User Datagram Protocol (UDP) probes; using Transmission Control Protocol (TCP) SYN probes to well-known ports as used in tcptraceroute; and Paris traceroute, which is made up of ICMP-Paris and UDP-Paris [Luckie, et al. 2008]. Paris traceroute helps avoid abnormalities in traffic paths that crop up because of load balancing [Augustin, et al. 2006]. Gupta et al. [2014] also make use of traceroutes as well as BGP routing tables for collecting their data.

According to papers by Cheswick et al. [2000], Mao et al. [2003], Luckie et al. [2008], Pansiot et al. [2010], and Augustin et al. [2006], traceroutes appear to be the best method for probing the Internet to gather information about its topology and what information is passed between and within networks. However there are many limitations with Van Jacobson's original implementation of traceroute, as described by Pansiot et al. [2010], Bush et al. [2009] and Augustin et al. [2006]. As a result, different applications of traceroute have been implemented and extensions added with arguments for each as to why they are better than Van Jacobson's original method. Examples include a variation called dual probing [Bush, et al. 2009] and mrinfo probing [Pansiot, et al. 2010].

Measuring the topology of the Internet from only one or a few vantage points gives a biased and incomplete view of the Internet [Shavitt and Weinsberg. 2009, Shavitt and Shir. 2005], which is why the use of many vantage points is advantageous. Shavitt and Weinsberg [2009] describe two methods that help one know which elements the vantage points should probe – namely an omniscient and an oblivious method.

The *omniscient method* assigns elements to vantage points optimally, but because this method requires that we know how well vantage points probe the inspected elements, this is computationally hard. This method is good for infrastructures that deploy measuring devices, like iPlane and Archipelago, statically [Shavitt and Weinsberg. 2009].

There is no need to know about vantage points' ability to probe elements for the *oblivious method* because this method randomly assigns each vantage point a set of random targets. The probability that the same element would be probed twice is very small though and virtually zero [Shavitt and Weinsberg. 2009]. This method is appropriate for community-based projects like DIMES [Shavitt and Shir. 2005] where we cannot assume that vantage points will be active all the time. DIMES only sees a vantage point as active when it connects to a server [Shavitt and Shir. 2005].

4. DISTRIBUTED NETWORK PROBING

Distributed network probing involves making use of many vantage points to get a more accurate view of the Internet [Shavitt and Weinsberg. 2009, Shavitt and Shir. 2005]. However, if one only uses a few vantage points, there is not a big advantage in adding only a few more. There is a big advantage, though, in adding thousands more points as they will add a significant percentage of new links [Shavitt and Shir. 2005]. Chavula et al. [2014] ran traceroutes from five different vantage points targeting 95 different universities in Africa in order to get a more complete view of the topology.

When discovering the topology of the Internet, especially on the AS level, Shavitt and Weinsberg [2009] propose that only considering having many active vantage points is not enough. More vantage points reveal new links between ASes [Shavitt and Shir. 2005], but if the vantage point is situated outside the two local ISPs between which we are trying to find a link, we might not be able to know of the existence of the connection because of BGP policies [Shavitt and Shir. 2005]. Therefore, one must also consider having a high diversity of locations to get an unbiased topology. According to Lakhina et al. [2003], if you infer AS degrees using traceroute-like sorting techniques, highest-degree ASes also tend to be nearer to the measuring sources than lower-degree ASes. Low-degree ASes are harder to detect so fewer vantage points observe these [Shavitt and Weinsberg. 2009].

It was also found that vantage points located in different types of ASes resulted in a bias towards these types of ASes in the observed topology. Because of this, one should also consider a broad diversity of types of ASes in which vantage points are located rather than just how many vantage points and where these vantage points are. Therefore, when considering vantage points, one should use many vantage points and also consider a large variety of locations with a broad variety of types of ASes in which they are located [Shavitt and Weinsberg. 2009].

There are many platforms available to aid with distributed probing. Some examples are Archipelago, DIMES, iPlane and RIPE Atlas.

Archipelago (Ark) is an active measurement infrastructure based on skitter and developed by the Centre for Applied Internet Data Analysis (CAIDA) [Hyun. 2006]. Skitter is a tool for actively probing the Internet and was used until about 2008 [McRobb, et al. 1999, Hyun. 2006]. Ark uses Scamper as the general-purpose measurement engine and provides a library to control aspects of the tool. In mid-December of 2008, there were 31 active Ark monitors situated all over the world with at least one monitor on each continent [Claffy, et al. 2009]. Now there are 107 [Hyun. 2015]. The majority of these monitors are deployed in academic or research organisations [Claffy, et al. 2009]. Unfortunately, there are still only five monitors in Africa, [Hyun. 2015] which does not help with analysing the network traffic in Africa.

DIMES aims to build a map of the Internet that is accurate, timely and comprehensive in terms of latency and topology and possibly bandwidth in the future [Shavitt and Shir. 2005]. As opposed to mapping the Internet at a very course AS level or very fine router level, DIMES aims to generate a mid-level map on top of these two where nodes are a group of routers working together [Shavitt and Shir. 2005]. DIMES makes use of thousands of measurement agents that interact with each other. Each host node runs light weight, low signature measurement agents as background processes [Shavitt and Shir. 2005]. These agents perform traceroute periodically to a set of IP addresses. This gives information pertaining to the router level in Internet connectivity [Cohen and Raz. 2006]. This is still useful as connectivity in the AS level can be induced by connectivity in the router level [Cohen and Raz. 2006]. One disadvantage of DIMES is that there are not many DIMES agents in Africa and as the purpose of this project is to collect and analyse data sent between African NRENs, many vantage points are needed in Africa. One also cannot run custom experiments with the DIMES infrastructure. This is disadvantageous because one cannot select probe destinations and get the data specifically needed for a particular research project [Shavitt and Shir. 2005].

iPlane is another measuring platform that measures Internet topology at an AS level. It also has many vantage points spread through the world and has a significantly higher average degree of ASes in academic ASes and IXPs than DIMES. However, neither DIMES nor iPlane have any vantage points in Central Africa, or the Arab countries in the Middle East [Shavitt and Weinsberg. 2009]. As this project focuses on interconnectivity within African NRENs, vantage points in Africa are very important.

RouteViews is another tool that collects BGP messages passively, providing data for discovery of the AS level topology. It has been observed that RouteViews observe a larger average degree for ASes than DIMES and iPlane [Shavitt and Weinsberg. 2009].

RIPE has a platform called RIPE Atlas that make use of thousands of active probes around the world to measure Internet connectivity and reachability in real-time [RIPE. 2010]. These probes are small USB-powered hardware devices attached to an Ethernet port that conduct measurements, such as ping, traceroute, DNS and SSLcert, and relay the data to the RIPE NCC. This data is then aggregated with data collected from other RIPE Atlas probes. In the space of two years (end of 2011 to end of 2013), the number of people hosting RIPE Atlas probes grew from over 1000 to over 5000. Cohen and Raz [2006] describe a method of collecting information for Internet topology discovery by gathering information from the Internet Routing Registry (IRR), which is a set of databases containing connectivity information for registered ASes. However, not all ASes want to publish who they are peering with and sometimes the information in the IRR is outdated. The most up to date and complete IRR database is maintained by RIPE [Cohen and Raz. 2006]. The RIPE Atlas platform is very advantageous to this project as many of the probes are situated in Africa [RIPE. 2010]. One can also send custom probes from a probe to an NREN so as to collect data that is required to research traffic between African NRENs.

5. VISUALISATIONS

The Internet can be mapped at the four different levels mentioned above - namely the IP Interface, router, PoP and AS levels. NRENs are on the AS level but many universities or campuses can be operating on the same AS so we will also consider visualisations on the PoP level. When mapping the Internet, one considers nodes and links to be mapped.

At the *IP Interface level*, a node represents a network interface with a designated IP address while the links show connectivity between these interfaces. At the *router level*, a node represents an IP-compliant network device or router while the links show that the devices' interfaces are on the same IP broadcast domain [Donnet and Friedman. 2007, Motamedi, et al. 2013].

At the *PoP level*, a map consists of nodes, which represent a PoP belonging to an AS, and links, which represent a physical link between the two PoPs' routers. A map at the PoP level provides information and limitations about latencies between PoPs, which helps with understanding the geographical properties of Internet paths - such as where ASes can connect and the coverage of ASes [Donnet and Friedman. 2007].

At the *AS level*, a map consists of nodes, which represent the ASes, and links, which show a business relationship between the two ASes. A business relationship implies that Internet traffic can be exchanged between the ASes. There are three main types of relationships – customer-provider, peer-peer and sibling relationships [Donnet and Friedman. 2007, Motamedi, et al. 2013].

There are many examples of visualisations of the mapping of the Internet as this problem has been thoroughly researched for many years now [Bhardwaj and Singh. 2015]. Examples are a visualisation by Cheswick et al. [2000], FlowScan [Plonka. 2000], a jellyfish structure [Siganos, et al. 2006], a visualisation metaphor [Da Lozzo, et al. 2014], Grid and Platter [Bhardwaj and Singh. 2015].

Fowler et al. [2014], Siganos et al. [2006] and Da Lozzo et al. [2014] have documented visualisations at the AS level of topology. Fowler et al. [2014] use heat maps superimposed onto a map of ASes in a tool called IMap. Siganos et al. [2006] proposed what they call a jellyfish structure to visualise the AS level topology of the Internet, intending for this design to be easily remembered and manually recreated. Da Lozzo et al. [2014] use a visualisation metaphor that shows a combination of a geographical representation of ASes as well as other data pertaining to ASes in a non-geographical manner. Although each of these authors have come up with quite different visualisations, they all steer away from purely geographical visualisations.

Although ASes have geographical information, such as their locations attached to them, a purely geographical visualisation of their relationships is not logical. Geographically correct maps contain a lot of information that is not of use to the understanding of the topology of the Internet, such as oceans and countries with limited Internet presence [Fowler, et al. 2014]. The properties of maps also don't relate to the Internet.

The surface area of a country has no correlation with the presence and contribution of global traffic and the physical distance between countries does not relate to the level of connectivity between those countries [Fowler, et al. 2014]. Some ASes are so large that they can span countries or even continents which would be illogical and cluttered to display on a geographical map [Da Lozzo, et al. 2014]. Also, displaying data on a geographical map would result in dense collections of information in places such as North America and Europe as they have a high Internet presence [Cheswick, et al. 2000]. As our research will be focused on traffic in Africa, these concerns are not as great but must be considered because one of the main research areas is analysing the traffic that travels via Europe to reach other African areas.

Another area of interest is visualising dynamic data. Research has been done into visualising dynamic data mainly to help prevent security attacks. Gilmour et al. [2007] provides a good visualisation of the Internet. However, this visualisation is static and doesn’t provide a changing view of the topology and how latencies are being affected by different routes. IMap visualises a large volume of dynamic data [Fowler, et al. 2014]. The visualisation metaphor by Da Lozzo et al. [2014] also makes use of animation to display the data in a dynamic sense. Plonka [2000] proposes a system called FlowScan that displays a continuous stream of near real-time data of network traffic.

6. CONCLUSIONS

There are two main components for this project – collecting data on the AS and PoP levels and visualising this data to analyse what traffic is being exchanged between and within NRENs in Africa. Many platforms have been developed to measure the topology of the Internet, but not many of them do not have enough vantage points in Africa. Since African NRENs is the main point of research of this project, vantage points in Africa are necessary for an unbiased dataset. It would appear that measurement tools such as ATLAS would be ideal to measure the connectivity between and within African NRENs as it appears to have more vantage points inside African NRENs, and also because it allows the use of many probes from many vantage points in different locations. User input would be required to determine the best method of visualising the data collected. A visualisation with some geographic properties would be valuable as one aspect of this project is to see which paths are going through Europe and whether inter-NREN traffic is being affected by their circuitous routes.

REFERENCES

FRYER, T. 2014. Research and education networks around the world and their use.

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.

GUPTA, A., CALDER, M., FEAMSTER, N., CHETTY, M., CALANDRO, E. AND KATZ-BASSETT, E. 2014. Peering at the Internet’s frontier: A first look at ISP interconnectivity in Africa. In *Passive and Active Measurement,* Springer, 204-213.

MOTAMEDI, R., REJAIE, R. AND WILLINGER, W. 2013. A Survey of Techniques for Internet Topology Discovery.

DONNET, B. AND FRIEDMAN, T. 2007. Internet topology discovery: a survey. *Communications Surveys & Tutorials, IEEE* 9, 56-69.

BUSH, R., MAENNEL, O., ROUGHAN, M. AND UHLIG, S. 2009. Internet optometry: assessing the broken glasses in internet reachability. In *Proceedings of the 9th ACM SIGCOMM conference on Internet measurement conference,* ACM Press, 242-253.

AUGUSTIN, B., CUVELLIER, X., ORGOGOZO, B., VIGER, F., FRIEDMAN, T., LATAPY, M., MAGNIEN, C. AND TEIXEIRA, R. 2006. Avoiding traceroute anomalies with Paris traceroute. In *Proceedings of the 6th ACM SIGCOMM conference on Internet measurement,* ACM Press, 153-158.

LUCKIE, M., HYUN, Y. AND HUFFAKER, B. 2008. Traceroute probe method and forward IP path inference. In *Proceedings of the 8th ACM SIGCOMM conference on Internet measurement,* ACM Press, 311-324.

BHARDWAJ, A.K. AND SINGH, M. 2015. Data mining-based integrated network traffic visualization framework for threat detection. *Neural Computing and Applications* 26, 117-130.

HARPER, K.C., CHEN, K. AND YEN, D.C. 2004. Distance learning, virtual classrooms, and teaching pedagogy in the Internet environment. *Technology in Society* 26, 585-598.

SHAVITT, Y. AND SHIR, E. 2005. DIMES: Let the Internet measure itself. *ACM SIGCOMM Computer Communication Review* 35, 71-74.

CHESWICK, B., BURCH, H. AND BRANIGAN, S. 2000. Mapping and Visualising the Internet. In *Proceedings of USENIX Annual Technical Conference (ATEC),* The USENIX Association.

MAO, Z.M., REXFORD, J., WANG, J. AND KATZ, R.H. 2003. Towards an accurate AS-level traceroute tool. In *Proceedings of the 2003 conference on Applications, technologies, architectures, and protocols for computer communications,* ACM Press, 365-378.

PANSIOT, J., MÉRINDOL, P., DONNET, B. AND BONAVENTURE, O. 2010. Extracting intra-domain topology from mrinfo probing. In *Passive and Active Measurement,* Springer, 81-90.

SHAVITT, Y. AND WEINSBERG, U. 2009. Quantifying the importance of vantage points distribution in internet topology measurements. In *INFOCOM 2009, IEEE,* IEEE Infocom, 792-800.

LAKHINA, A., BYERS, J.W., CROVELLA, M. AND XIE, P. 2003. Sampling biases in IP topology measurements. In *INFOCOM 2003. Twenty-Second Annual Joint Conference of the IEEE Computer and Communications. IEEE Societies,* IEEE, 332-341.

HYUN, Y. 2006. Archipelago measurement infrastructure. (November 2014) Retrieved April 24, 2015 from http://www.caida.org/projects/ark/.

MCROBB, D., CLAFFY, K. AND MONK, T. 1999. Skitter. *Skitter: CAIDA's macroscopic Internet topology discovery and tracking tool* 2015.

CLAFFY, K., HYUN, Y., KEYS, K., FOMENKOV, M. AND KRIOUKOV, D. 2009. Internet mapping: from art to science. In *Conference For Homeland Security, 2009. CATCH'09. Cybersecurity Applications & Technology,* IEEE, 205-211.

HYUN, Y. 2015. CAIDA Monitors: The Archipelago Measurement Infrastructure. (April 2015). Retrieved April 20, 2015 from http://www.caida.org/data/monitors/monitor-map-ark.xml.

COHEN, R. AND RAZ, D. 2006. The Internet Dark Matter-on the Missing Links in the AS Connectivity Map. In Proceeding of IEEE *INFOCOM 2006,* IEEE.

RIPE, N. 2010. Ripe atlas. (2010). Retrieved April 24, 2015 from https://atlas.ripe.net.

PLONKA, D. 2000. FlowScan: A Network Traffic Flow Reporting and Visualization Tool. In *LISA,* USENIX, 305-317.

SIGANOS, G., TAURO, S.L. AND FALOUTSOS, M. 2006. Jellyfish: A conceptual model for the as internet topology. *Communications and Networks, Journal of* 8, 339-350.

DA LOZZO, G., DI BATTISTA, G. AND SQUARCELLA, C. 2014. Visual discovery of the correlation between BGP routing and round-trip delay active measurements. *Computing* 96, 67-77.

FOWLER, J.J., JOHNSON, T., SIMONETTO, P., SCHNEIDER, M., ACEDO, C., KOBOUROV, S. AND LAZOS, L. 2014. IMap: visualizing network activity over internet maps. In *Proceedings of the Eleventh Workshop on Visualization for Cyber Security,* ACM Press, 80-87.

GILMORE, J., HUYSAMEN, N. AND KRZESINSKI, A. 2007. Mapping the African Internet. In *Proceedings 2nd IFIP International Symposium on Wireless Communications and Information Technology in Developing Countries, (WCITD),* (Sept 2007), Pretoria, South Africa.