

A Topology Visualisation Tool for National Research and Education Networks in Africa

Chantal YANG¹, Hussein SULEMAN², Josiah CHAVULA³

Department of Computer Science, University of Cape Town,

Private Bag X3, Cape Town, 7700, South Africa

E-mail: 1yngcha005@myuct.ac.za, 2hussein@cs.uct.ac.za 3jchavula@cs.uct.ac.za

Abstract: Previous research on National Research and Education Networks (NRENs) in Africa has shown high latency in traffic exchanged between networks, with 75% of this traffic taking circuitous routes through Europe. This paper presents a user-centered creation of a geospatial visualisation tool that can be used to show the network structure of African NRENs, using Traceroute as a method for network topology discovery. The visualisation tool allows multiple Traceroutes for a single IP address to be viewed on a map from various vantage points. The tool's effectiveness and accuracy at communicating the logical network topology of African NRENs and routes of traffic traversal was assessed. Results from usability tests showed that the tool was able to effectively and accurately communicate the logical topology on a country and continental level with over 68% of visual queries having been answered successfully.

Keywords: National Research and Education Networks; Network Topology Visualisation; Geospatial Visualisation; Information Visualisation; Network Topology Discovery; Traceroute

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 [1]. In order to facilitate better research and communication, NRENs aim to reduce latencies between educational institutions, promote bandwidth sharing and improve traffic engineering [1].

Despite an increase in the number of national Internet Exchange Points (IXPs) in Africa, high latencies still occur, indicating a need for further research into the network topology of NRENs.

Previous research on NRENs in Africa has shown that 75% of traffic originating from and destined for African institutions traverse circuitous routes through Europe [7]. This contributes to a high latency in traffic exchanged between African NRENs [7].

In order to understand the underlying topology of African NRENs, the collection of traceroute data for logical topology discovery is necessary. Such topology data needs to be presented in a format that can easily be evaluated and understood by stakeholders involved with NRENs in Africa. By presenting this information in a visual and interactive manner through an information visualisation (InfoVis); gaps, anomalies, clusters or patterns in the data may be identified [2][3][4][5]. An information visualisation (InfoVis) is the representation of abstract data on an interactive visual interface [4] [6].

A visualisation would be useful in helping identify network and routing issues (such as incorrect or suboptimal routing); identifying possible routes of network traffic traversal between NRENs (continental vs intercontinental) and aiding in the peering decision-making process [8].

This paper presents the design of a geospatial visualisation, which visualises the network topology of institutions and NRENs in Southern and Eastern Africa, using Traceroute data collected using the Ripe Atlas platform [9]. This was done using a User Centred Design methodology over two iterations. We evaluated the effectiveness and accuracy of the visualisation at communicating the network topology (physical and logical) of African NRENs and routes of potential traffic traversal having conducted usability tests.

2. Related Work

Periakaruppan and Nemeth [1999] discussed GTrace, a “Graphical Traceroute tool” used to discover routing loops and help decide route implementations. Using several heuristics, the system determines the location of a node as the traceroute is executed before displaying the traceroute on a world map as series of nodes (hops) and links. A table displays more detailed information about the Traceroute, including the hop number, hop IP Address and hop hostname [10].

Gilmore et al. [2007] generated router and Autonomous System (AS) level maps of the African Internet using data collected from traceroutes sent to selected IP addresses. At the router level, a Java-based tool, Terrapix, was created, where 2D and 3D visualisations mapped nodes and links to geographic locations. For the AS level, CAIDA’s Walrus tool was used to generate logical node-link, graph visualisations in a 3D space. Using these visualisations, a “picture” of the African Internet was pieced together - though the accuracy of this is questionable as traceroutes in the paper were only conducted from a single vantage point [11] [12].

3. Design and Implementation

A User-Centred Design (UCD) approach was used to develop the visualisation. UCD describes a process where the user is involved throughout the design cycle where user needs and goals are determined [6][13]. In the context of a visualisation, this implies gaining an understanding of the visual queries (an information need addressed by a visualisation) that potential users may have. Users of the visualisation are envisaged to be NREN managers and network engineers of tertiary education institutions that are part of the UbuntuNet Alliance.

The UCD design process consisted of three phases: the early envisioning phase, the global specification phase and the detailed specification phase [13]. In each phase, analysis, design and evaluation were conducted.

Iteration of the design cycle is concluded once a specific criteria has been reached, such as if users are able to adequately answer visual queries [6]. Two iterations were conducted before the final visualisation was designed.

Making use of a User-Centered Design approach means that access to potential or intended users is imperative. These users need to not only have the appropriate domain knowledge and be easily accessible but also be available throughout the design cycle to ensure consistent feedback. Sourcing such users was challenging as not only are they highly specialized in their role (and therefore few) but also geographically dispersed.

3.1 Initial Visualisation Design

In order to obtain real user tasks and goals, specialised individuals working within UbuntuNet Alliance-affiliated NRENs and organisations were invited to an online survey. The survey enquired about currently used network management operational tools and metrics of interest, such as latency and throughput. Of the responses received, a “lack of comprehensive routing information” was stated as a network management limitation

encountered while network down-time and congestion were cited as common network problems experienced.

The design of the first prototype of the visualisation was largely based on existing traceroute visualisation tools such as the OpenIPMap prototype [15]. A prototype (Figure 1) with limited functionality was produced using the Google Maps Javascript API v3 and a small sample of traceroute data collected from the Ripe Atlas platform [9][16]. The MaxMind GeoLite City Database was used to map IP addresses to city-level coordinates and used to position markers on the map for both probes and destination IP addresses [17]. NREN probes were represented with red marker symbols, university probes were purple markers and destination IPs were blue circles. A route was illustrated on the map connecting the probe, hops and destination IP markers as a line with an animated arrow indicating the direction of the traceroute. An image overlay was placed on the map to show the positioning of undersea fibre cables around Africa.

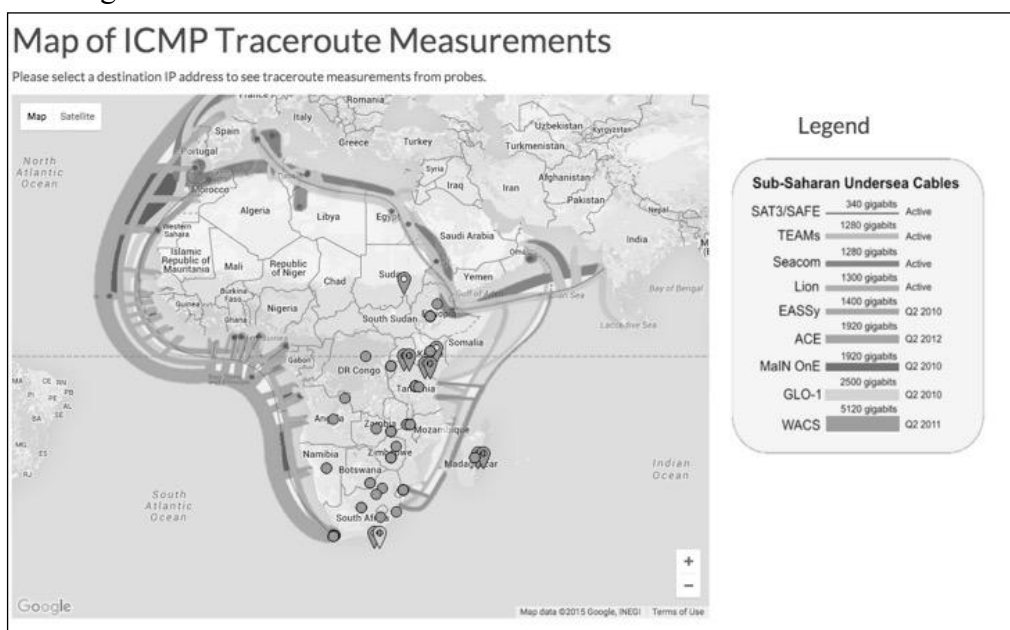


Figure 1. Screenshot of Interactive Visualisation Prototype

An informal evaluation was conducted, where the prototype was presented to an expert user group. The user group consisted of domain experts in networking from the department of Computer Science: two postgraduate students, two technical staff members and a networks lecturer.

Users had differing opinions on the undersea fibre overlay surrounding Africa: one user thought the multiple lines and colours in the overlay added too much noise to the visualisation and needed to be removed while another thought it was useful for deducing which cables were used by traceroutes to destinations.

For the final design, it was decided that the undersea cable overlay be removed since it would not be possible to identify exactly whether a certain cable would be used for intercontinental routes of network traffic.

3.2 Final Visualisation Design

A focus group was conducted with the same expert user group in section 3.1. It was discovered that users were interested in understanding whether high latencies occurred due to a lack of physical infrastructure (fibre cable) or the result of routing protocols that needed to be changed. Focus group participants were also interested in identifying problems on a per link basis between traceroute hops. Therefore, additional data needed to be

displayed in the visualisation. Other suggestions made were that rather than showing each hop, hops should be aggregated at country level and that the animation of the arrow of Traceroute links could vary by speed based on link quality.

Figures 2-5 shows the final visualisation design based on information received in the first and second iteration. In addition to the Google Maps Javascript API, the DataTable plug-in for jQuery was used [18]. The visualisation displays several dimensions of data, including the location of active Internet Exchange Points in Africa, physical fibre cables and collected traceroute information. IXP data was obtained from Packet Clearing House’s Internet Exchange Point Directory while fibre data was obtained from the AfTerFibre Project (the African Terrestrial Fibre Optic Cable Mapping Project) [19][20]. The visualisation allows for viewing multiple traceroutes sent from various vantage points to be seen on the map all at once.

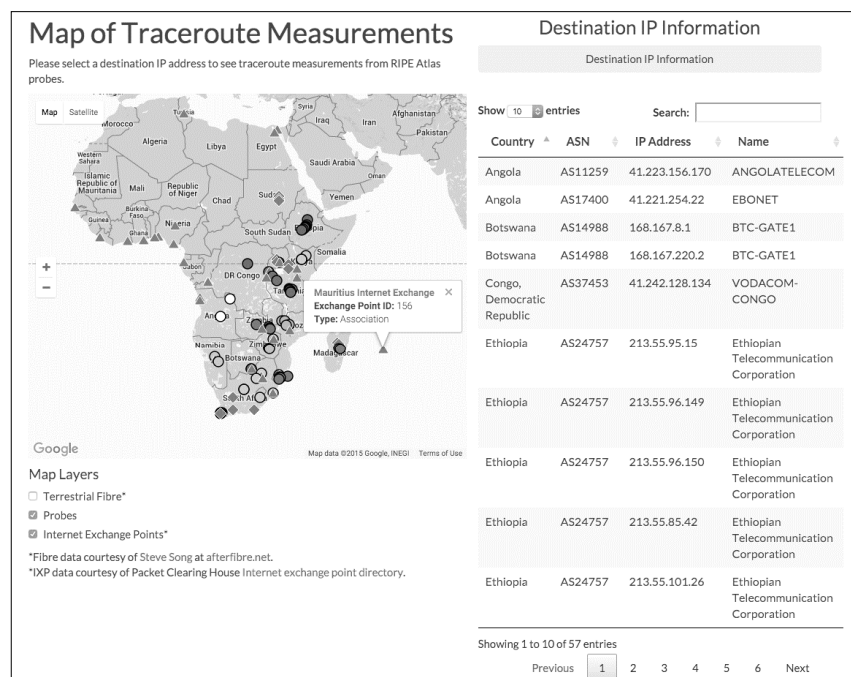


Figure 2. Initial Overview Screen with Internet Exchange Point and Probe Layers Displayed

Schneiderman’s Visual Information Seeking Mantra (overview, zoom, filter and details on demand) was applied when designing the interface and visualisation flow [5]. Users are first presented with an overview of all the available IP addresses that traceroutes are sent to on a map (Figure 2). When the user hovers the mouse pointer over an icon on the map, an info window with information related to that point of interest is displayed. After clicking a chosen destination IP address icon on the map, multiple Traceroute measurements are shown on the map from all available probes (various vantage points) to that particular IP address. (Figure 4.)

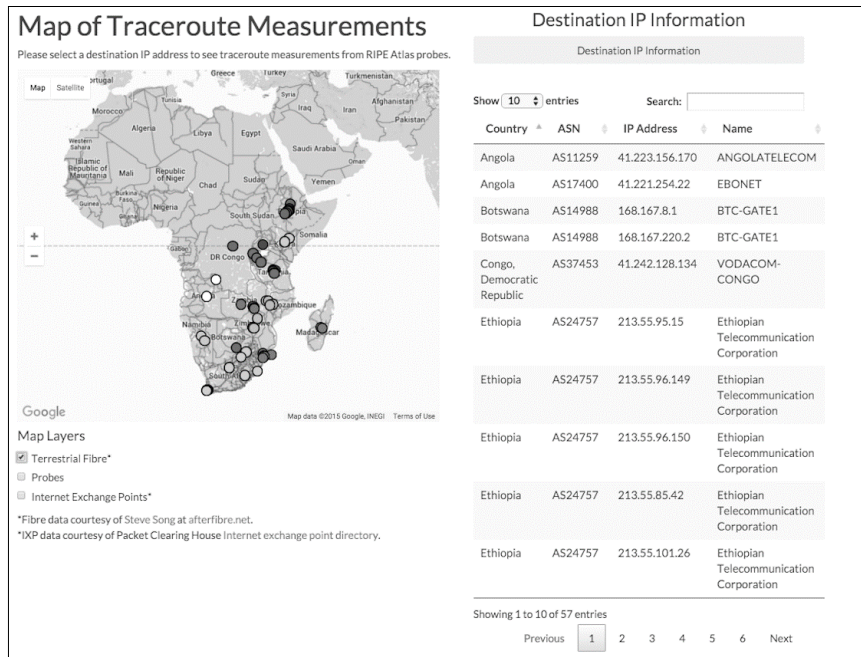


Figure 3. Initial Overview Screen with Terrestrial Fibre layer Displayed

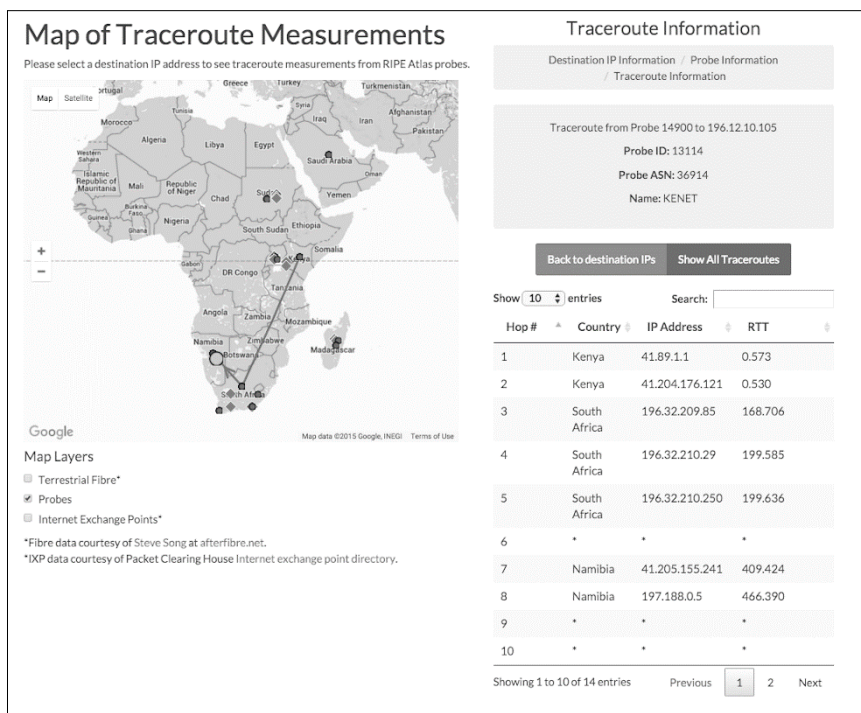


Figure 4. View Multiple Traceroutes to Selected Destination IP

Hops are illustrated on the map as small green circles which are connected by lines from source to destination indicating a Traceroute measurement. An animated arrow moves along these lines to show the direction of the packets. Clicking on a specific probe icon allows the user to view a single Traceroute as a red line with animated red arrow. Detailed hop information is then displayed in the adjacent table (Figure 5).

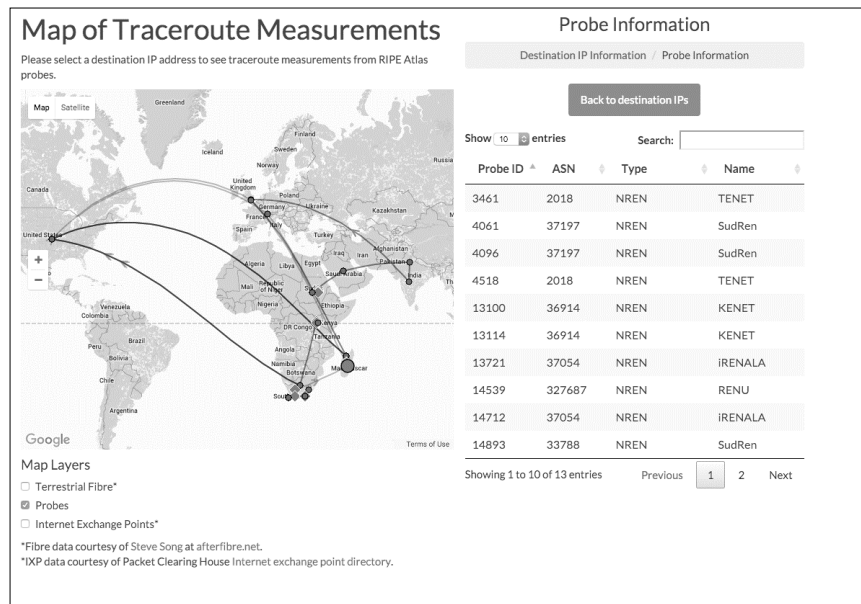


Figure 5. View a Single Traceroute to a Selected Destination IP Address

Clicking a row in the table zooms in to the related icon on the map. Selecting the checkboxes allows different layers (terrestrial fibre, probes, Internet exchange points) to be added to or removed from the map: Figure 3 shows the terrestrial fibre layer selected. The search box can be used to filter results in the table and more easily locate points of interest. Different symbols were used to represent different types of data on the map: destination IPs are represented by different coloured circles on the map and colour-coded by country so that users (NREN managers) can more easily locate a point of interest (Figure 2). Probes were shown on the map as blue diamonds, intermediate hops of the traceroute as small green circles and IXPs as pink triangles, as seen in Figure 2. These features differed on two channels of colour and shape to make items distinguishable from each other.

4. Evaluation and Results

4.1 Evaluation Metrics

The evaluation methods conducted were short-term and assessed potential use of the visualisation [21]. Users' subjective feedback and opinions of the visualisation tool were taken into account and used to determine the effectiveness, accuracy and usability of the visualisation [21]. In this case, effectiveness refers to a tool's functionality and examines a user's performance when performing tasks [22]. Accuracy refers to the way in which the information is conveyed to the user – that is, can users accomplish tasks in order to obtain the correct answers to a set of questions related to visual queries? Usability describes the quality of use of an application by a user (ease of use, satisfaction, efficiency) and is, therefore, an important aspect of the visualisation to consider [22][23].

Effectiveness and accuracy were measured using the metric of successful task completion while usability was measured using the System Usability Scale (SUS) questionnaire [24][25][26].

4.2 Usability Tests

Usability tests were conducted with a total of 23 users to assess the effectiveness, accuracy and usability of the visualisation. The 23 participants consisted of a mixture of students in

2nd, 3rd and 4th year. Tests were conducted in an uncontrolled environment of the Computer Science honours laboratory, where participants accessed a webpage that contained the visualisation. Tasks for the usability test were designed with the objective of establishing whether users could use the interactive functionality of the visualisation to answer visual queries. The tests lasted approximately 30-40 minutes and were conducted in one hour sessions.

Users performed tasks that served as a self-guided walkthrough of the available functionality of the visualisation. This allowed users to familiarise themselves with the interface before they attempted to answer a set of 10 questions. The visual queries formulated were questions related to identifying networks (physical and logical), where they connect (location of source, destination, intermediate hops) and potential routes of traffic traversal (traceroute paths at a country and continental/intercontinental level) (Table 1). Thus, the 10 questions/tasks in the usability test were based on answering the visual queries and completing relevant tasks.

Table 1. Visual Queries and Related Research Theme

| Visual Query | Visual Query Type | Research Theme |
|---|---|---|
| Which country on the African continent has the most fibre? | Physical Network: Most Fibre Cable | Network topology; identification of physical network |
| What is the route for the traceroute between institution A and institution B? | Route of Traceroute: Country level | Identification of potential routes of traffic traversal; where networks connect (source, destination and intermediate hops) |
| Does the route for a traceroute from institution A to institution B travel intracontinentally (within the continent)? | Route of Traceroute: Intracontinental Level | Identification of potential routes of traffic traversal; where networks connect (source, destination and intermediate hops) |
| Does the route for a traceroute from institution A to institution B travel intercontinentally (to a different continent)? | Route of Traceroute: Intercontinental Level | Identification of potential routes of traffic traversal; where networks connect (source, destination and intermediate hops) |

Free-response questions were appended to the SUS questionnaire to evaluate user experience and usability. These consisted of a field for “General Comments” (where users’ opinions and problems encountered could be described), “What features were useful?” (to determine commonly used features) and “What features are missing?” (to elicit overlooked requirements) [21].

4.3 Analysis of Results

Successful task completion is characterised by the ability of a participant to obtain specific data when carrying out a task [27][28]. If a question was answered correctly by a participant in the question set, then the task was deemed to have been successfully completed. This same definition applied to the accuracy of answering visual queries.

Results were grouped by Visual Query Type shown in Table and demonstrated that participants were able to correctly determine the country with the most fibre cables with 100% accuracy (Figure 6). Identification of a Traceroute’s path on a country level had 72.73% accuracy; 68.18% accuracy on an intracontinental (within the continent) level and intercontinental (between continents) level (Figure 6).

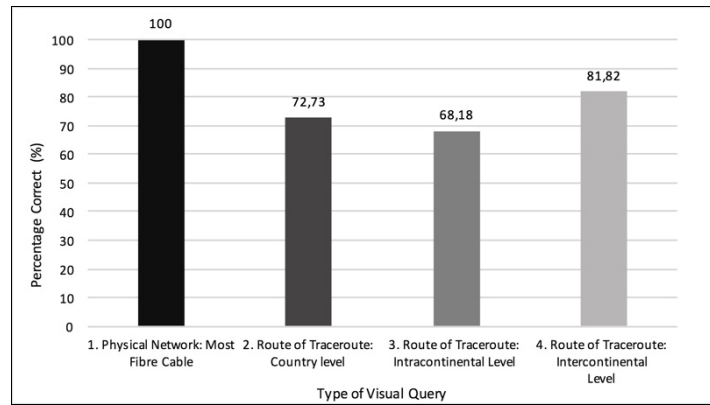


Figure 6. Correct Visual Queries Answered

Figure 7 shows the percentage of the six tasks successfully or partially completed during the usability test.

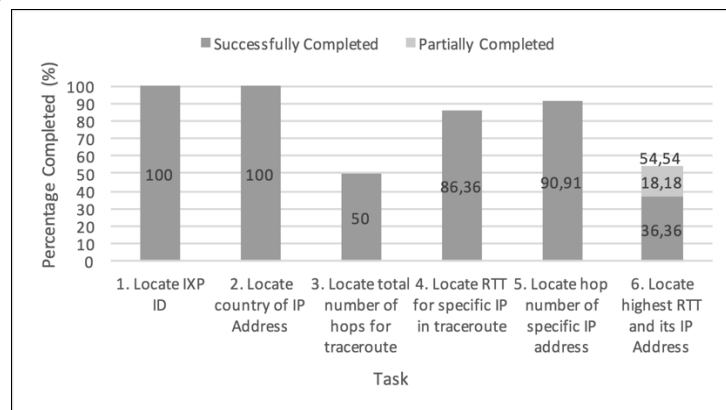


Figure 7. Successful Tasks Completed

Both Task 1 (locating the ID of an IXP) and Task 2 (locating the country of a particular IP address) had 100% successful task completion. In comparison to these tasks, Tasks 3,4,5 and 6 had lower successful task completion rates. This is understandable as Tasks 3,4,5 and 6 were tasks that required more complex subtasks to be performed, thus increasing the likelihood of an error. For Tasks 4 and 5, only 3 participants and 2 participants respectively answered these questions incorrectly (with one participant having not answered the questions at all). Without having observed these users on an individual basis, it cannot be reasoned why they were not able to complete the tasks successfully. It is possible that these problems may have resulted from a lack of familiarity with Traceroute by users doing usability tests.

The success rate of all 10 questions (4 visual queries and 6 tasks) across all participants was calculated to be 79.55%, where successfully completed tasks (correctly answered questions) were allocated 1 point and partially answered questions were allocated 0.5 points [27][28]. According to Sauro [2011], who conducted an analysis of 1200 usability tasks, the average task-completion rate is 78%, which means the task completion rate for the visualisation is above average.

Figure 8 presents the frequency of System Usability Scores by category for 23 users who participated in the usability test and a pilot usability test by an expert user. The scores shown are for a total of 24 participants, including the pilot usability test.

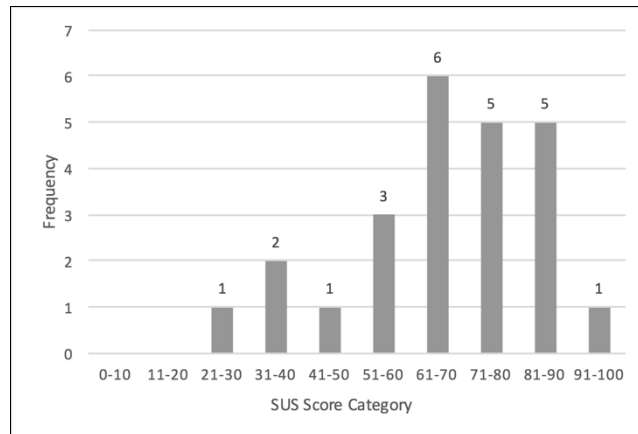


Figure 8. Frequency of Individual SUS Scores by Category

The average SUS score was then calculated to be 67.8, which is about the same average score reported from an analysis of 500 studies making use of the SUS, where the average score was 68 [28]. Furthermore, 50% of the respondents in this study scored the visualisation with a SUS score greater than the average of 68.

4.4 Limitations of Results

In terms of the visualisation, a mismatch between the model and reality was uncovered during the usability tests, where the visualisation (eg. number of hops, traceroute links joining hops) was not an accurate representation of the Traceroute data collected. One reason for this is due to constraints of the Maxmind GeoLite Cities database: as the database maps IP addresses only on a city level, it is possible that different hop IP addresses may resolve to the same coordinates, thus leading users to believe there were fewer hops for a traceroute depicted in the visualisation than was reflected in the actual traceroute data.

5. Conclusions and Future Work

In this paper, we presented the design of a geospatial visualisation of the network topology of National Research and Education Networks in Africa using Traceroute data. A user centered design approach was used and applied over two iterations. The final visualisation shows various dimensions of data on the map, including the location of Internet Exchange Points, the placement of terrestrial fibre and Traceroute information. In comparison to other Traceroute visualisation tools (such as GTrace and Terrapix mentioned in Section 2.1.1), which either only visualise a single Traceroute measurement or display routes from a single vantage point, the tool designed allows multiple traceroutes sent from various vantage points to be viewed on the map.

The final evaluation investigated whether the geospatial visualisation designed can effectively and accurately communicate the network topology of African NRENs and allow users to identify networks where they connect and the routes traffic possibly traverses. Effectiveness and accuracy were evaluated by checking the correctness of visual queries answered. These visual queries related to the identification of networks (physical and logical), the location of the source, destination and intermediate hops of these networks and potential routes of traffic traversal at a country and continental level. It was demonstrated that the visualisation was able to effectively and accurately communicate the logical topology and that potential routes of traffic traversal could be identified at a country and continental level with an accuracy of 68% and above.

Future work would be to determine the extent to which topology visualisation supports decision making, such as the placement of NREN interconnection points. This however, would require expanding the current topology data set used in the visualisation by using

more network probes in more network locations. A current challenge with regards to visualizing the UbuntuNet's topology is that Atlas probes are available only in a few NRENs. To have a complete view of the topology, there is need to deploy network probes in as many locations as possible. This would require a deliberate effort by the Alliance to increase the number of Atlas probes available across different campuses. The ideal situation would be to have at least one Atlas probe in each campus network, although having a probe in each and every NREN would be a good starting point. This would allow for the collection of a complete and accurate data set for the network topology, which can then be aggregated to a visualization server managed by the Alliance. Added features of the visualisation could be to allow users to view and compare traffic routing for different protocols (e.g ICMP, UDP, TCP) and adjusting the speed of the Traceroute animation to allow for the identification of slow links between hops.

References

- [1] Fryer, T., Research and Education Networks around the World and their Use. in *Segunda Conferencia de Directores de Tecnología de Información*, (Lima, 2012), Red Clara, 1-13.
- [2] Becker, R. A., Eick, S. G., and Wilks, A. R. 1995. Visualizing network data. *Visualization and Computer Graphics, IEEE Transactions on*, 1(1), 16-28.
- [3] Carr, D. A. 1999. Guidelines for designing information visualization applications. *Proceedings of ECUE, 99*, 1-3.
- [4] Keim, D. A., 2002. Information visualization and visual data mining. *IEEE Transactions on Visualization and Computer Graphics*, 8(1), pp.1-8.
- [5] Shneiderman, B. 1996. The eyes have it: A task by data type taxonomy for information visualizations. in *Visual Languages, 1996. Proceedings., IEEE Symposium on*, IEEE, 336-343.
- [6] Wassink, I., Kulyk, O. and Van Dijk, B., 2009. Applying a user-centered approach to interactive visualisation design. *Trends Interaction. Visualisation.* (2009), 175-199. DOI:<http://dx.doi.org/10.1007/978-1-84800-269-2>
- [7] 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.
- [8] Steenberg, R. A Practical Guide to (Correctly) Troubleshooting with Traceroute. 2009
- [9] Atlas.ripe.net., What is RIPE Atlas? - RIPE Atlas — RIPE Network Coordination Centre. 2015. <https://atlas.ripe.net/about/>.
- [10] Periakaruppan, R. and Nemeth, E. 1999. "race - Graphical Traceroute Tool". *LISA '99: Proceedings of the 13th USENIX conference on System administration*: 69-78.
- [11] Gilmore, J. S., Huysamen, N. F., & Krzesinski, A. E. 2007. Mapping the african internet. in *Proceedings Southern African Telecommunication Networks and Applications Conference (SATNAC), (Sept 2007), Mauritius*, South African National Research Foundation.
- [12] Shavitt, Y., and Weinsberg, U. 2009. Quantifying the importance of vantage points distribution in internet topology measurements. In *INFOCOM 2009, IEEE*. 792-800.
- [13] Abras, C., Krichmar D. M., and Preece. J 2004. User-centered design. *Bainbridge, W. Encyclopedia of Human-Computer Interaction. Thousand Oaks: Sage Publications* 37, 4: 445-56. <http://doi.org/10.1.1.94.381>
- [14] UbuntuNet Alliance, (2014), *What is UbuntuNet Alliance?*, Retrieved April 25, 2015 from UbuntuNet: http://www.ubuntu.net/sites/default/files/What_Is_UbuntuNet_v1_20141030.pdf
- [15]. Aben, E. emileaben/django-openipmap. *GitHub*, 2014. <https://github.com/emileaben/django-openipmap>.
- [16] Google Developers., Google Maps Javascript API | Google Developers. 2015. <https://developers.google.com/maps/documentation/javascript/>.
- [17] Dev.maxmind.com., GeoLite2 Free Downloadable Databases « Maxmind Developer Site. 2015. <http://dev.maxmind.com/geoip/geoip2/geolite2/>.
- [18] Datatables.net., DataTables | Table plug-in for jQuery. 2015. <https://www.datatables.net/>.
- [19] Prefix.pch.net., Packet Clearing House (PCH) - Internet Exchange Directory. <https://prefix.pch.net/applications/ixpdir/>.
- [20] Song, S. AfTerFibre (Archived) Mapping Terrestrial Fibre Optic Cable Projects in Africa - Archived. *Many Possibilities*, 2011. <https://manypossibilities.net/afterfibre-old/>.

- [21] Lam, H., Bertini, E., Isenberg, P., Plaisant, C., Carpendale, S. 2012. Evaluation To cite this version : Seven Guiding Scenarios for Information Visualization Evaluation.
- [22] Dal Sasso Freitas, C.M, Luzzardi, P.R.G., Cava R.A, Winckler M.A.A., Pimenta, M.S, Nedel, L.P, 2002. Evaluating Usability of Information Visualization Techniques. *Proc. 5th Symp. Hum. Factors Comput. Syst.* (2002), 40–51.
- [23] Valiati, E.R, Pimenta, M.S, and C M D S Freitas. 2006. A taxonomy of tasks for guiding the evaluation of multidimensional visualizations. *Proceedings of the 2006 AVI workshop on Beyond time and errors: novel evaluation methods for information visualization*, October:1–6. <http://doi.org/http://doi.acm.org/10.1145/1168149.1168169>
- [24] Nielson, J. Success Rate: The Simplest Usability Metric. *Nngroup.com*, 2001. <http://www.nngroup.com/articles/success-rate-the-simplest-usability-metric/>.
- [25] Brooke, J. 1996. SUS - A quick and dirty usability scale. *Usability Evaluation in Industry.*, 189 (194). 4–7.
- [26] Koua, E.L., Kraak, M.J. 2004. A usability framework for the design and evaluation of an exploratory geovisualization environment. *Proceedings. Eighth International Conference on Information Visualisation, 2004. IV 2004.*: 0–5. <http://doi.org/10.1109/IV.2004.1320138>
- [27] Usability.gov,. Planning a Usability Test. <http://www.usability.gov/how-to-and-tools/methods/planning-usability-testing.html>.
- [28] Sauro, J. Measuring Usability with the System Usability Scale (SUS): MeasuringU. *Measuringu.com*, 2011. <http://www.measuringu.com/sus.php>.