

REDUCING LATENCY IN AFRICAN NRENS USING PERFORMANCE-BASED LISP/SDN TRAFFIC ENGINEERING

Josiah Chavula, Melissa Densmore, and Hussein Suleman
Department of Computer Science, University of Cape Town, South Africa
{jchavula,mdensmore,Hussein}@cs.uct.ac.za

ABSTRACT

Active topology measurements on the African Internet have showed that over 75% of the intra-Africa traffic destined for Africa's National Research and Education Networks (NRENs) uses intercontinental links, resulting in high latencies and data transmission costs. The goal of this work is to investigate how latency-based path selection using Locator/Identifier Separation Protocol (LISP) and Software Defined Networking (SDN) in NRENs can be used to reduce inter-NREN latencies. We present aspects of an experimental prototype implementation for real-time topology probes to discover lower-latency remote gateways and dynamic configuration of end-to-end Internet paths. Simulation results indicate that ranking remote ingress gateways, and dynamic configuration of end-to-end paths between gateways can lower the average latency for inter-NREN traffic exchange.

KEYWORDS

Software Defined Networking, Traffic Engineering, Latency, Internet exchange points, African National Research and Education Networks (NRENs)

1. INTRODUCTION

Research collaboration and e-resource sharing in sub-Saharan Africa continues to be hampered by the limited interconnectivity that is not only expensive, but fails to meet the quality of service (QoS) required for collaborative applications among the National Research and Education Networks (NRENs). Despite the ongoing interconnection efforts by African NRENs, over 75% of the inter-NREN traffic is being exchanged through circuitous routes traversing inter-continental links and Internet exchange points in Europe and North America, thereby experiencing high latencies [1].

Efforts have been made to improve the Internet traffic exchange among the NRENs in Africa. A key player in this effort is the UbuntuNet Alliance, an association of the NRENs in Eastern and Southern Africa. Between 2011 and 2014, the UbuntuNet Alliance has been implementing the Africa-Connect Project to create a regional research and education Internet network interconnecting the NRENs in the region. The project has aimed to interconnect Southern and Eastern African NRENs into a regional network through the use of terrestrial network facilities. The project involves establishment of Points of Presence (PoPs) in major cities in the region -

notably in Mtunzini, Maputo, Dar es Salaam, Nairobi, Kampala and Kigali, and interconnecting them with broadband cross-border links to create a regional research network. So far, the intra-Africa interconnection serves six NRENs: TENET (South Africa), MoRENet (Mozambique), TERNET (Tanzania), KENET (Kenya), RENU (Uganda) and RwEdNet (Rwanda). Furthermore, transcontinental links between Nairobi and the UbuntuNet Alliance PoP in Amsterdam, as well as from Cape Town to London, have been established, thereby linking UbuntuNet Alliance with GEANT, the European research network (Figure 1).



Fig. 1. The UbuntuNet Alliance regional network

The establishment of multiple PoPs, multiple intra-Africa Internet links, as well as multiple transcontinental links, provides new opportunities for implementing multipath routing and traffic engineering mechanisms with the aim of improving performance of traffic exchange among Africa's NRENs. NRENs could implement mechanisms that would allow them to announce to each other, and make use of multiple Internet attachment points to exchange traffic. One protocol that allows edge networks to exchange traffic over multiple Internet gateways is the Locator/Identifier Separation Protocol (LISP) [2, 3]. LISP allows edge networks to announce

multiple Internet gateways, known as Route Locators (RLOCs), and to influence the selection of incoming paths. Through a mapping system, LISP allows networks to announce preferences for multiple RLOCs. The availability of multiple locators for the same destination increases path diversity[4] and enable multipath routing, as source networks can select among multiple gateways to reach a destination network.

This paper demonstrates the potential for performance improvement in the Pan-African NRENs by employing, at the network edge, traffic engineering techniques that are based on end-to-end multipath ranking. Using a Software Defined Network (SDN) experimental topology, and a LISP mapping system, the paper examines the potential for dynamically ranking egress and ingress links between multihomed NRENs based on end-to-end path metrics. The objective is to minimize latency for two-way delay sensitive traffic flows (e.g. real-time classroom streaming and video conferences between universities), and minimize intercontinental bandwidth utilization.

2. BACKGROUND AND RELATED WORK

Given the opportunities for traffic engineering provided by the multiple intra-continental and transcontinental links provided by the Africa Connect network, one way of improving the performance and optimization of traffic exchange across African NRENs is to enable dynamic selection of optimal traffic exchange routes based on application QoS needs. For example, path selection for delay sensitive applications can be made based on prevailing end-to-end latencies through either the intercontinental transit links or through the intra-Africa links.

2.1 LISP-based Multipath Traffic Engineering

The Locator/Identifier Separation Protocol (LISP) [3] simplifies multihoming. LISP divides the Internet's address space into two - locally routable Endpoint IDentifiers (EIDs) and global Route LOCators (RLOCs). By separating the host address space from the locator address space, LISP introduces a level of indirection that allows networks to specify preferences for multiple ingress gateways (locators). The availability of multiple locators for the same destination increases path diversity since the source networks are able to forward traffic for a particular destination through multiple remote locators (gateways).

Thus, additional end-to-end performance gains can be achieved with the ability to dynamically select the ingress link at the destination network. In Figure 2, for example, edge networks A and B are multi-homed to networks x,q and y,z respectively. Depending on how the routing is done in the Internet core, the choice of the egress link by network A, i.e. (A,x or A,q), has potential to influence selection of ingress link towards B i.e. 3,y or 4,z. Since each end-to-end path has its own unique path metrics in terms of bandwidth, delay, and loss, selection of particular egress and ingress links at A and B impacts the overall quality of the end-to-end path.

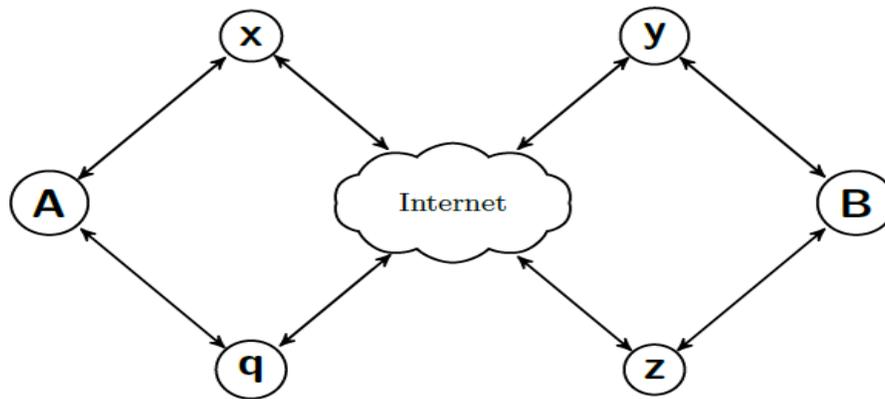


Fig. 2. Multihomed Networks A and B, multihomed through providers (x,q) and (y,z) respectively

One prominent work that uses LISP for traffic engineering is the ISP-Driven Informed Path Selection (IDIPS) [5]. IDIPS is a request/response service where centralized server nodes perform network measurements towards popular destinations, and clients request path rankings for a set of sources and destinations. The IDIPS server ranks the available paths based on a client's ranking preference and measured path metrics. In IDIPS implementation, path ranking is further influenced by the destination's preferences in the locator mapping. The selected paths therefore reflect not only the source network's ranking criteria, but also the destination's preferences for incoming traffic.

2.2 SDN-based Path Enforcement

One challenge with inter-domain multi-path routing and end-to-end traffic engineering is with regard to enforcement of paths across different domains. Software Defined Networking (SDN) provides new opportunities for flexible management of Internet routing and packet forwarding [6]. An SDN-based IXP [7] allows IXP participants to have access to an SDN controller and to write policies that override the default policies of the IXP's BGP route server. SDN has three important characteristics that are useful for interdomain traffic engineering [7]. Firstly, in contrast to traditional switches that forward traffic based only on the destination MAC address, SDN enables packet forwarding based on multiple header fields. Secondly, an SDN controller consolidates control messages from multiple remote networks, such that source and destination networks can remotely configure forwarding paths through a controller. Thirdly, the controller's direct control of the data plane enables dynamic/programmatic configuration of the forwarding tables. With these SDN opportunities, it is possible to allow edge networks some control over selection of inter-domain forwarding paths at Internet exchange points, thereby having more control on the end-to-end paths.

3. A MODEL FOR PERFORMANCE-BASED PATH SELECTION

The traffic engineering framework depicted in Figure 3 is based on the ability of traffic source gateways to select the destination's ingress link based on metrics of the edge-to-edge path. This requires that the source gateway should have a mechanism for learning the destination's multiple ingress links.

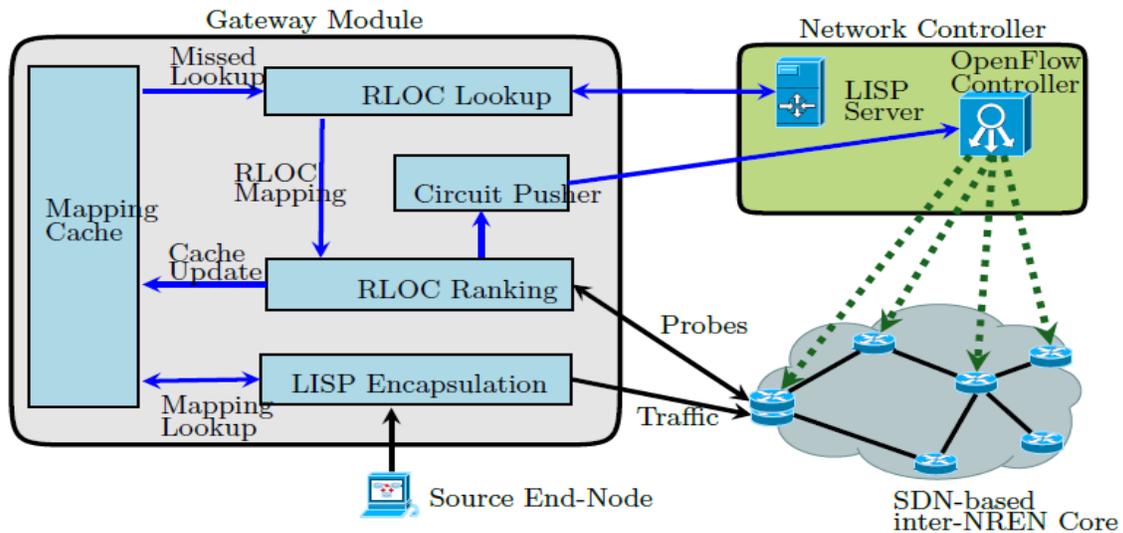


Fig. 3. LISP/SDN Multi-Path Traffic Engineering Model

In summary, the model works in the following manner: For each new traffic flow from a local source node to some remote network, the source network queries the LISP mapping system (through an RLOC lookup API) to obtain the destination network's RLOCs. After obtaining the remote RLOCs, a Locator Ranking module performs active measurements towards remote RLOCs. The Locator Ranking module uses the network metrics obtained from the active measurements to rank the local and remote gateways, after which it updates the local mapping cache. For each new traffic flow, the source network selects the egress and ingress gateways based on the rankings in the mapping cache. Once the egress and ingress links are selected, a Circuit Pusher module is invoked to configure, through an Open Flow SDN control, an end-to-end switching path between the source and destination RLOCs via the selected egress/ingress links. If no mapping exists in the mapping cache for a destination network, the RLOC Mapping module is invoked to perform the lookup, after which the RLOC Ranking module performs the ranking.

3.1 Path Performance Measurement

The key network metrics in this model are end-to-end latency, jitter and packet loss. Latency, measured as round trip time (RTT) for traffic to move from source to destination, and for the acknowledgement packet to be received by the sender, is an important characteristic that affects the performance and responsiveness of Internet applications. Jitter, on the other hand, is the variation in latency over time between a traffic source and destination node. Packet loss is a measure of the percentage of packets lost along the data path for each traffic flow. To obtain traffic characteristics of the network in terms of the key path metrics, active measurement techniques are used. In particular, a ping-based tool is used for sending probe packets, from each edge-network gateway, to the destination networks' RLOCs. Upon retrieving the destination network's RLOCs from the mapping server, the source gateway sends a ping probe, through each of its egress links, to each of the destination's RLOCs. By analysing the solicited responses, topological characteristics such as round-trip-times, jitter and packet loss are obtained. The values obtained from the responses are used to rank the different paths.

3.2 LISP-based Path Ranking

With LISP, multihomed edge networks are able to achieve some degree of path diversity, as multiple alternate gateways become visible between source and destination networks. Achieving optimal end-to-end performance in such environments requires that the source and destination networks should be able to evaluate the alternate paths, and to dynamically select both the source network's egress link and the destination network's ingress link. In particular, the source network needs a way of discovering and evaluating end-to-end links through alternate egress and ingress links.

The routing cost for an edge-to-edge path can be modelled as a vector comprising the measured performance metrics and the network RLOC preferences [8]. Let $P(A_{xy})$; $P(B_{yx})$ be the performance cost vectors for the two edge networks A and B, with respect to forwarding traffic through their access links x and y respectively. The performance cost P from each edge network comprises a set of end-to-end path metrics K, (eg. latency, packet loss, jitter) weighted by variable A_i , such that:

$$P(A_{xy}) = \sum_{i=1}^n A_i \cdot K_{xyi}$$

$$P(B_{yx}) = \sum_{i=1}^n A_i \cdot K_{yx_i}$$

where $\sum_{i=1}^n A_i = 1$.

To calculate the total cost, $T(A_{xy})$; $T(B_{yx})$, the source preference cost ($\phi(x)$) is combined with the performance cost P, using a variable scaling factor λ :

$$T(A_{xy}) = \lambda P(A_{xy}) + (1 - \lambda)\phi(x);$$

$$T(B_{yx}) = \lambda P(B_{yx}) + (1 - \lambda)\psi(y);$$

where $0 \leq \lambda \leq 1$.

In LISP implementation, selection of the local egress RLOC as well as the remote ingress RLOC is determined by priority and weight values recorded in RLOC records that are retrieved from the mapping system and stored in a local cache. If multiple locators for the same destination exist, the priority values, ranging from 0 to 255, are used to select the locator that is most preferred. In this work, the calculated path costs are translated into RLOC priority values using a log function that scales the costs into values between 0 and 255. The egress-ingress RLOC pair that has the lowest resultant cost is the one that is used for the outgoing traffic flow between the locator pair. While the mapping remains valid, all subsequent matching flows between the locator pair uses the cached locator ranking. The cached locator mapping remains valid for a configurable TTL period of 60 seconds, after which the mapping and locator ranking process is repeated.

3.3 End-to-end Path Configuration

Each edge network gateway has a Circuit Pusher, an SDN module for setting up end-to-end circuits between locators. After the source gateway has selected both the local and remote locators, it invokes the SDN module to configure, through an SDN network controller, a unidirectional end-to-end circuit between locators. This is achieved by installing flow entries on all switches that are part of the shortest path between two selected locators. The installed path is unidirectional because each source gateway independently performs RLOC lookups and path ranking, and sets up a circuit toward the remote RLOC.

4. TESTBED IMPLEMENTATION

The overarching purpose of this study is to evaluate the extent to which LISP and SDN can support dynamic selection of end-to-end paths between multi-homed edge networks. The objective was therefore to assess network performance in a topology that uses a dynamic performance-based path selection, versus static default path routing. Figure 4 depicts the testbed implementation.

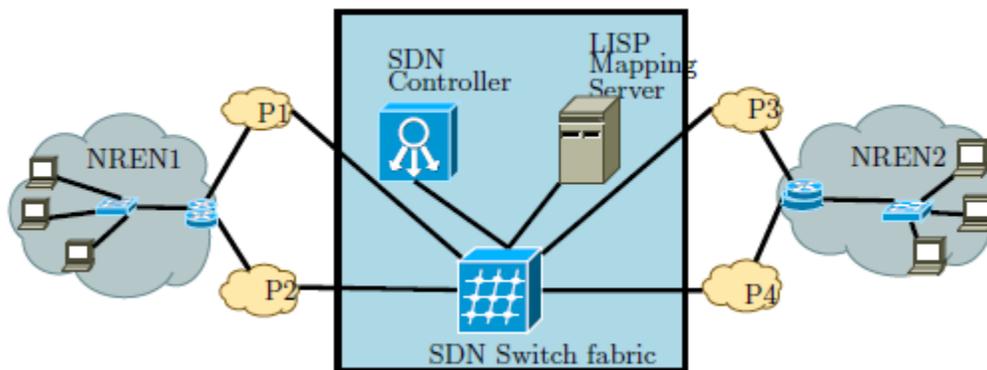


Fig. 4. SDN/LISP Traffic Engineering for dual homed NRENs

4.1 Topology

The testbed was built using a LISP-based SDN simulated network consisting of an SDN (Openflow) controller, an Open flow switch, and a LISP mapping server. The topology was built in a virtual environment, using the Mininet network emulator [9]. Mininet allows the creation of virtual hosts, switches, controllers and links. Furthermore, Mininet nodes run a standard Linux kernel and network stack, and can therefore run real network applications. The switches and network controllers are based on the standard SDN protocol Open Flow. An Internet-like topology was built in Mininet and was used together with a LISP mapping system and LISP transit routers (xTRs), which are based on an open source implementation of LISP called Open LISP[10]. The interconnection among the edge networks is through ISPs that interconnect at a common Internet exchange point. More specifically, the network was designed with the following features:

- 48 routers connected to the IXP switch; the routers represent participant networks at the IXP. The JINX has, as of December 2014, 48 participant networks.
- Each IXP participant has 10 edge networks connected to it, and each edge network is dual-homed and serves a total of 5 hosts, representing campus networks. In total, there were 1,200 end hosts as potential sources and destinations for traffic flowing through the IXP fabric.
- To simulate multi-homing, each edge network is connected to two provider networks present at the IXP. The primary link represents the intra-Africa link and is configured with lower latency, while the secondary link represents an transcontinental link and is configured with higher latency.
- The access links between the edge networks and the provider networks are configured with band-width evenly distributed between 2 Mbps and 4 Mbps. The links between the IXP participants (provider networks) and the IXP are equally provisioned with bandwidth of 100 Mbps.
- The end-to-end latencies are modelled on latencies measured for traffic exchanged between African NRENs. Round trip latencies in this experiment are distributed between 60 ms and 700 ms.

4.2 Network Traffic

A realistic evaluation of a network model requires emulating the network with traffic that characteristically resembles the traffic pattern of the emulated networks. This is important, as a major scalability issue with centralized network architectures, such as the Open flow controllers and LISP mapping systems, hinges on their ability to cope with the traffic flow characteristics in the network [11].

Researchers have characterized Internet traffic based on flow metrics such as byte volume, packet volume, flow duration, and flow inter-arrival time. The percentage of UDP traffic has increased with the advent of many UDP based P2P applications and streaming multimedia, which transport large volumes of data[12]. In 2009, a CAIDA survey showed that the ratio of UDP to TCP traffic was almost 0.21 in terms of packet numbers, 0.11 in terms of byte count, and 3.09 in terms of flows.

The length of data flows impacts the relative latency introduced at the controller and, furthermore, the number of active flows has implications for the size of the forwarding tables maintained at each Open Flow forwarding device. For example, a characterization of university campus network traffic [13] established that 21.4% of the traffic was carried by flows longer than 10 minutes, 12.6% by flows longer than 20 minutes, and nearly 2% was carried by flows longer than 100 minutes. Short flows are bursty and have flow speeds ranging from 1 Bps to over 10 kBps, while longer flows are slower, around 50 Bps for 40 min flows, [13].

The test traffic for the experiment is therefore based on the following Internet traffic characteristics:

1. Protocol flow: UDP to TCP ratio: 3:1
2. Flow Duration:
 - 0 - 2 sec : 45% of all the traffic
 - 2 sec - 5 mins : 55% of all the traffic

3. Flow inter-arrival time: 4 ms - 40 ms (or 25 to 250 new flows per sec)
4. Flow rate:
 - Short flows (0 - 60 sec) : 1 Bps - 10 kBps
 - Medium flows (1 min - 5 mins) : 100 Bps - 50 Bps

4.3 Test Traffic Generation

Some of the most widely used traffic generators include Iperf, PackETH, D-ITG, and Ostinato [14,15]. PackETH [16] is a stateless packet generator designed for Ethernet networks, and supports a number of protocols including UDP, TCP and ICMP. Iperf[17] is mostly used for evaluating topology parameters such as bandwidth, delay, window size and packet loss, for both TCP and UDP traffic. Iperf provides an estimation of received and transmitted data rates. Ostinato[14] is a user level traffic generator tool that supports UDP and TCP protocols at multiple rates. DITG (Distributed Internet Traffic Generator) [18] can generate Internet traffic with a user defined packet inter-departure times. This work made use of D-ITG for traffic generation and for obtaining performance metrics.

5. RESULTS

A key objective of the RLOC ranking was to discover and direct traffic flows through lower latency paths towards multi-homed remote networks. TCP traffic is particularly impacted by network round-trip-times, and this paper has considered performance of TCP traffic when RLOC ranking and dynamic path configuration is employed. Furthermore, the evaluation also considers how jitter is affected due to path ranking and circuit configuration.

5.1 Round Trip Times

The key results from a series of experiments suggest that in cases where the paths towards the different RLOCs of an edge network have significantly different RTTs, latency based ranking and selection of RLOCs does help to lower the overall latency in the network. Figure 5 shows the RTT dispersion and mean for TCP traffic in a LISP/SDN topology, with each flow lasting between 1sec and 300 sec. The vertical lines represent the dispersion of flow RTTs over the time interval, with each traffic flow RTT averaged over 2 sec intervals. The instantaneous average RTT for all the flows is indicated by the blue and red lines. The results show that RLOC ranking results in a 20 % lower overall latency compared to the default gateway forwarding.

The performance gains from RLOC ranking appear to diminish significantly with increased network load. As the network gets more congested, the observable gain from RLOC ranking is significantly reduced. This can be explained from the fact that as the network links get more congested, the otherwise shorter links begin to exhibit equally higher RTTs. With higher RTTs, many of the probe packets time-out, prompting the egress RLOC to forward packets towards the destination's default RLOC. Even when the probe packets generate responses, the RTT values of the otherwise shorter paths tend to be just as high, thereby being ranked lower and resulting in selection of the other paths for traffic forwarding. Figure 6 shows the dispersion and mean of the RTT at the point when the network is congested. Figure 7 and Figure 8 further illustrates the RLOC ranking effects in normal network operation and when there is congestion.

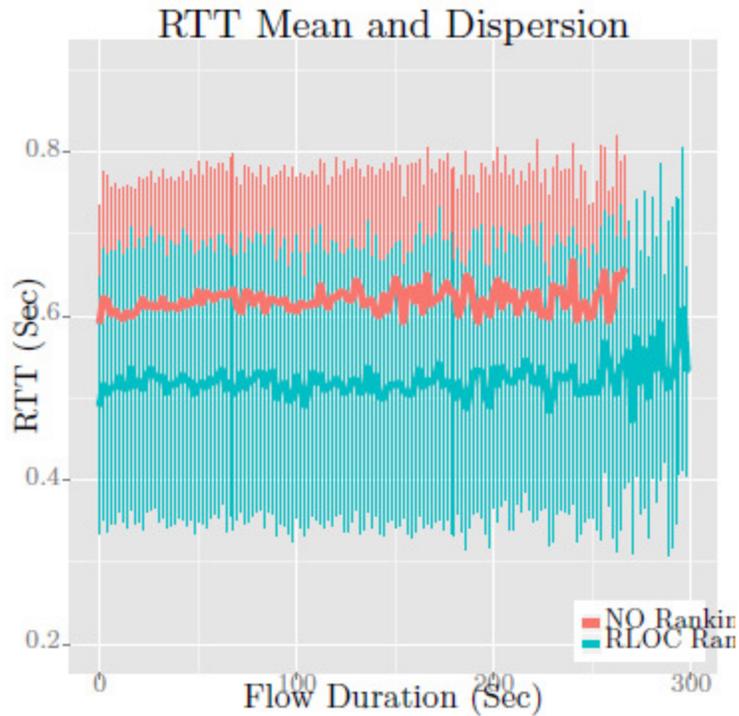


Fig. 5. Round-trip times for network operating with LISP gateway ranking vs non-path ranking LISP operation

5.2 Jitter

Apart from latency, jitter is another key metric that affects performance of interactive Internet applications. Some causes of Internet jitter include congestion in the core network as well as in the access links. Results from the simulation network suggest that RLOC ranking and dynamic path configuration does increase the overall jitter in the network. Given a network topology and traffic profile, RLOC ranking in a LISP network results in higher overall jitter than when no RLOC ranking is employed. Figure 9 shows the jitter for both RLOC ranking and the normal LISP operation. On average, RLOC probing is seen to increase jitter by 20 %. The jitter can be attributed to the delay experienced by some packets during the time the local gateway performs path measurement and ranking.

However, as the network reaches congestion point, both the RLOC ranking and non-ranking scenarios experience similarly higher jitter. This is illustrated in Figure 10, where jitter for both ranking and non-ranking LISP operations have their average jitter increase significantly. As congestion occurs, the probe engine fails to discover any lower latency RLOCs and resorts to using default paths.

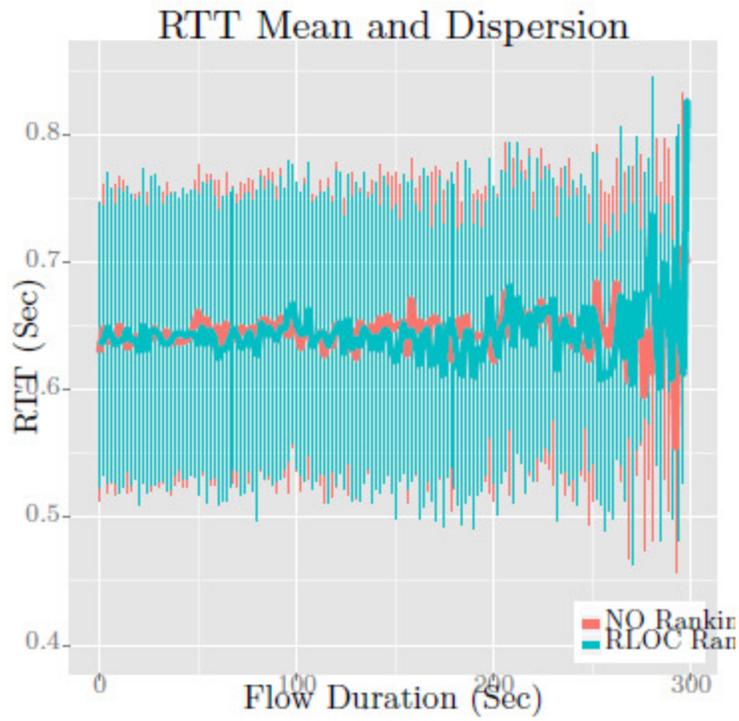


Fig. 6. Round trip times for path ranked and non-path ranked LISP operation at network congestion point

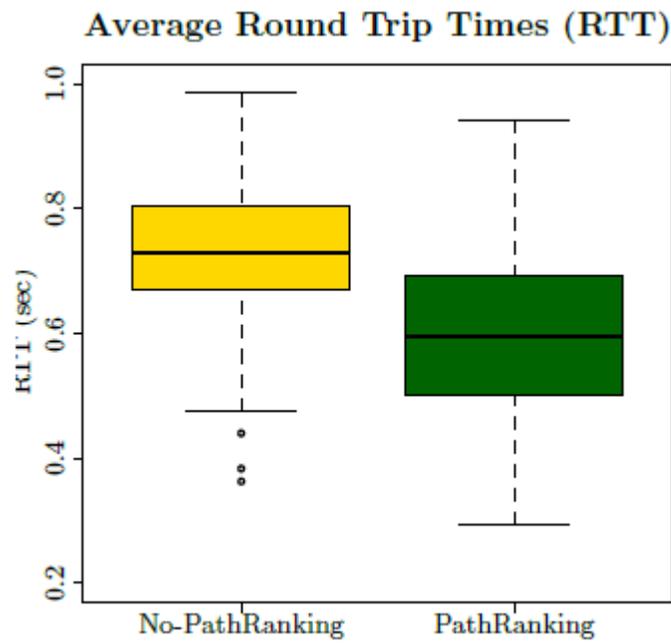


Fig. 7. RTT without congestion

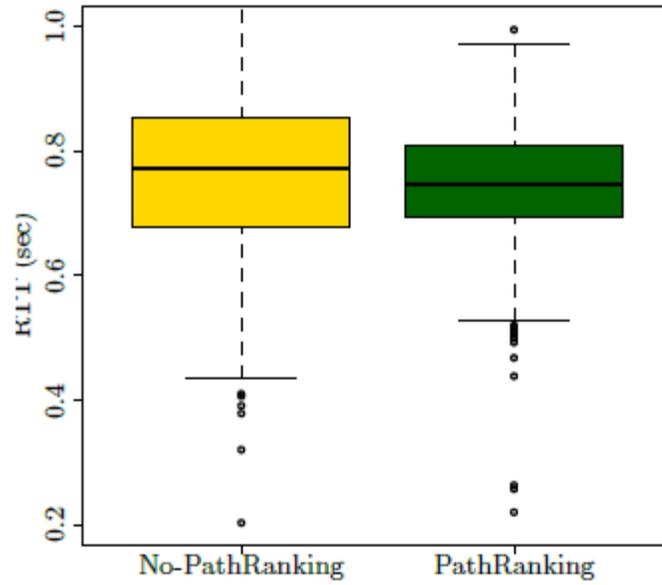
Average Round Trip Times (RTT)

Fig. 8. RTT with congestion

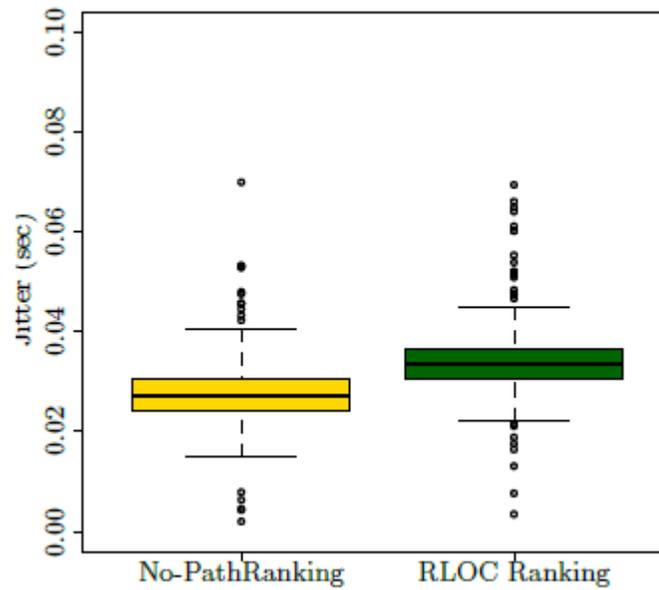


Fig. 9. Jitter with no congestion

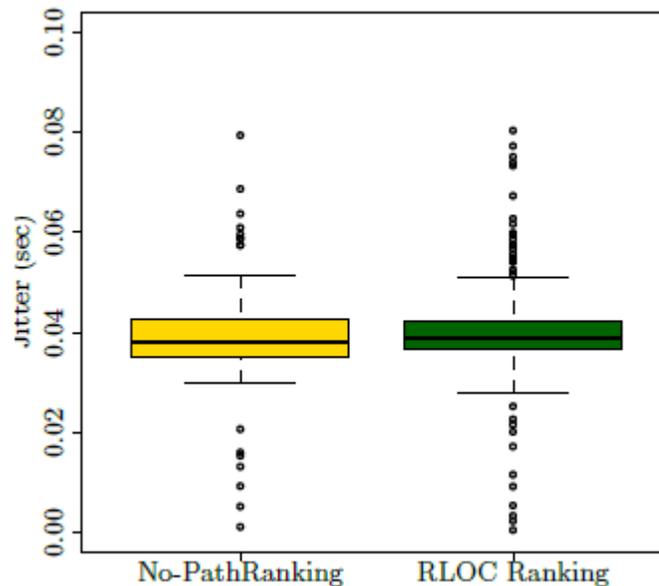


Fig. 10. Jitter with congestion

6. DISCUSSION

6.1 RLOC Ranking During Congestion

The RTT results show that the RLOC ranking mechanism fails to produce positive results under network congestion. In general, centralized systems are vulnerable to performance bottlenecks under system overload. For instance, inter-arrival times of traffic flows have implications on the performance of centralized network controllers [11]. In an Open Flow network architecture, for example, a scalability challenge stems from the fact that the first packet of each flow is forwarded to a central controller, which is responsible for determining and configuring the forwarding path for the flow. Similarly, for LISP, the egress gateway performs a lookup from a mapping server to determine each new flow's destination network's RLOC. In either case, the flow inter-arrival time has a scalability impact on the network, as higher rates for new flows result in bottlenecks at the SDN and LISP controllers, thereby introducing latency and jitter. Although the scalability and performance bottleneck would affect both the ranking and normal LISP operations, the RLOC ranking mechanism would experience more severe impact as it is dependent on receiving replies from probe packets, which take longer when there is congestion. One way of dealing with RLOC ranking during congestion is to reduce the amount of probing required by using historical performance information to select the RLOCs. Also, the amount of probing needs to be reduced by performing ranking only for critical flows (eg. delay intolerant applications).

6.2 Effects of RLOC Ranking on Jitter

The observed jitter in the experiments reveals that the process of RLOC ranking and path configuration does increase the overall jitter in the network. One way of minimising the jitter is to reduce the path setup time. In the presented model, a new end-to-end path is not setup until path

probing/ranking and SDN path configuration is complete. Reducing jitter would require pre-ranking the RLOCs, so that new flows don't wait too long before circuits are set up. Pre-ranking of paths could be based on RLOC performance history as well as known traffic characteristics in the network.

6.3 Model Limitations

One challenge with the model presented in this paper is the assumption that NRENs are multihomed. This is true to a large extent as, in general, campus networks that are part of NRENs have more than one Internet attachment points; a specific network attachment point for traffic destined within the NRENs' network and another for traffic destined to non-NREN networks. Where an NREN has only one Internet attachment point, it can still appear multi-homed by performing prefix de-aggregation and announcing separate prefixes with different RLOCs. This would enable them to still benefit from multipath traffic engineering.

Another challenge is to do with independent selection of the remote RLOCs by the source network, which could result in violation of the destination network's preferences and policies. This could negatively impact on the destination's policies. For edge networks that have some form of cooperation, such as the case with NRENs within the UbuntuNet Alliance, a mutually beneficial approach would be to employ some level of coordination in selection of gateways. While each NREN would aim to optimize traffic cost and QoS performance (latency), collectively, the NRENs can optimize the performance of some common preferred applications. Balancing the individual NREN's optimization objectives with those of the peering domains requires coordination and routing cooperation among the peers. For UbuntuNet Alliance, benefits from this level of cooperation could include better performance of the network applications, reduction in usage of intercontinental links, as well as reduction in the cost of inter-NREN traffic exchange.

7. CONCLUSION AND FUTURE WORK

Results from this study show that through dynamic ranking of the local and remote LISP locators, a source network can perform traffic engineering towards a destination without requiring any form of cooperation in the network. It is evident that by leveraging LISP capabilities through integration with SDN, there is potential for improving traffic exchange performance. This however implies granting NRENs more exibility and control of routing and traffic engineering across Internet exchange points, such as to be able to dynamically select routing paths among multiple ingress and egress links. This could have important applications for NRENs that experience high bandwidth costs [19]. The ability to perform multipath routing has potential to over significant performance enhancements and cost savings. For example, delay-sensitive applications such as VoIP, can be mapped onto lower latency intra-continental links, while traffic for bandwidth intensive _le-sharing applications can be routed through higher capacity inter-continental links. Furthermore, by using Openow's ability to customize the packet forwarding rules, and by appropriately matching packets into ows using header tags, it is possible to set up application specific traffic engineering mechanisms. With this capability, a group of NRENs can jointly form traffic engineering strategies specifically for certain applications of common interest, e.g. inter-university video streaming, or access to e-library sites within the domain.

This paper has motivated for and described a performance based traffic engineering mechanism that involves path measurement, gateway ranking, and SDN-based edge-to-edge path configuration. While Internet infrastructure being implemented by the UbuntuNet Alliance through the Africa Connect project has great potential to improve traffic exchange among African NRENs, the inability of traditional protocols to fully take advantage of the available path diversity remains a challenge. For this reason, African NRENs should, apart from implementing the physical interconnectivity, also consider appropriate traffic engineering mechanisms to allow individual NRENs to discover and use optimal inter-NREN paths. However, a globally optimal solution requires coordination and collaboration among several domains that form part of the end-to-end path. Future work will investigate mechanisms that would enable African NRENs to perform collaborative performance based and application specific traffic engineering.

REFERENCES

- [1] J. Chavula, N. Feamster, A. Bagula, and H. Suleman, "Quantifying the effects of circuitous routes on the latency of intra-africa internet traffic: A study of research and education networks," vol. 147, pp. 64{73, 2015.
- [2] 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.
- [3] D. Saucez, L. Iannone, O. Bonaventure, and D. Farinacci, "Designing a deployable internet: The locator/identifier separation protocol," *Internet Computing, IEEE*, vol. 16, pp. 14{21, Nov 2012.
- [4] S. Secci, K. Liu, G. Rao, and B. Jabbari, "Resilient traffic engineering in a transit-edge separated internet routing," in *Communications (ICC), 2011 IEEE International Conference on*, pp. 1{6, June 2011.
- [5] D. Saucez, B. Donnet, and O. Bonaventure, "Idips: Isp-driven informed path selection," 2008.
- [6] C. E. Rothenberg, M. R. Nascimento, M. R. Salvador, C. N. A. Correa, S. Cunha de Lucena, and R. Raszuk, "Revisiting routing control platforms with the eyes and muscles of software-defined networking," pp. 13{18, ACM, 2012.
- [7] A. Gupta, M. Shahbaz, L. Vanbever, H. Kim, R. Clark, N. Feamster, J. Rexford, and S. Shenker, "Sdx: A software defined internet exchange," 2013.
- [8] S. Secci, K. Liu, and B. Jabbari, "Efficient inter-domain traffic engineering with transit-edge hierarchical routing," *Computer Networks*, vol. 57, no. 4, pp. 976{989, 2013.
- [9] B. Heller, N. Handigol, V. Jeyakumar, B. Lantz, and N. McKeown, "Reproducible network experiments using container based emulation," in *Proc. ACM CoNEXT*, Dec. 2012.
- [10] L. Iannone, D. Saucez, and O. Bonaventure, "Openlisp: an open source implementation of the locator/id separation protocol," *IEEE INFOCOM, Demo paper*, 2009.
- [11] T. Benson, A. Akella, and D. A. Maltz, "Network traffic characteristics of data centers in the wild," in *Proceedings of the 10th ACM SIGCOMM conference on Internet measurement*, pp. 267{280, ACM, 2010.
- [12] M. Zhang, M. Dusi, W. John, and C. Chen, "Analysis of udp traffic usage on internet backbone links," in *Applications and the Internet, 2009. SAINT'09. Ninth Annual International Symposium on*, pp. 280{281, IEEE, 2009.
- [13] L. Quan and J. Heidemann, "On the characteristics and reasons of long-lived internet flows," in *Proceedings of the 10th ACM SIGCOMM conference on Internet measurement*, pp. 444{450, ACM, 2010.
- [14] A. Botta, A. Dainotti, and A. Pescap, "A tool for the generation of realistic network workload for emerging networking scenarios," *Computer Networks*, vol. 56, no. 15, pp. 3531 { 3547, 2012.
- [15] S. S. Kolahi, S. Narayan, D. D. T. Nguyen, and Y. Sunarto, "Performance Monitoring of Various Network Traffic Generators," in *International Conference on Computer Modeling and Simulation*, 2011.
- [16] M. Jemec, "packeteth{ethernet packet generator," 2012.

- [17] A. Tirumala, F. Qin, J. Dugan, J. Ferguson, and K. Gibbs, "Iperf: The tcp/udp bandwidth measurement tool," <http://dast.nlanr.net/Projects>, 2005.
- [18] S. Avallone, S. Guadagno, D. Emma, A. Pescape, and G. Ventre, "D-itg distributed internet traffic generator," in Quantitative Evaluation of Systems, 2004. QEST 2004. Proceedings. First International Conference on the, pp. 316-317, IEEE, 2004.
- [19] B. Barry, C. Barton, V. Chukwuma, L. Cottrell, U. Kalim, M. Petitdidier, and B. Rabiou, "egy-africa: better internet connectivity to reduce the digital divide," in IST-Africa, 2010, pp. 1-15, IEEE, 2010.

AUTHORS

Josiah Chavula is a PhD candidate at the Hasso-Plattner-Institute in Computer Science, at the University of Cape Town. His research focuses on Internet performance and traffic engineering in Africa's National Research and Education Networks (NRENs). He is interested in approaches for achieving flexible collaborative Internet traffic engineering using Software Defined Networking (SDN). He obtained a Master of Science degree (Networking and Internet Systems) from Lancaster University (UK), and a BSc (Computer Science) from University of Malawi.



Melissa Densmore is a Senior Lecturer in the Department of Computer Science at University of Cape Town. Melissa completed her PhD at University of California, Berkeley in Information Management and Systems, a 3 year ethnographic study of the use of Internet and mobile technologies by health practitioners and NGO staff in a health financing program in Uganda, has an MSc in Data Communications, Networks and Distributed Systems from University College London, and holds a BA in Computer Science from Cornell University.



Hussein Suleman is an Associate Professor in Computer Science at the University of Cape Town. His research is situated within the Centre for ICT for Development (ICT4D) and the Digital Libraries Laboratory. Hussein's main research interests are in digital libraries, ICT4D, information retrieval, cultural heritage preservation, Internet technology, high performance computing and computer science education. He completed his undergraduate degrees and MSc at the then University of Durban-Westville (now University of Kwazulu-Natal) and finished a PhD at Virginia Tech in 2002, in the area of component-based digital libraries.

