

Investigating the Causes of Congestion on the African IXP substrate

Rod  rick Fanou
IMDEA Networks Institute and
Universidad Carlos III de Madrid
roderick.fanou@imdea.org

Francisco Valera
Universidad Carlos III de Madrid
fvalera@it.uc3m.es

Amogh Dhamdhare
CAIDA/UC San Diego
amogh@caida.org

ABSTRACT

The goal of this work is to investigate the prevalence, causes, and impact of congestion on the African IXP substrate. Towards this end, we deployed Ark probes (within networks peering) at six African IXPs and ran the time-sequence latency probes (TSLP) algorithm, thereby collecting latency measurements to both ends of each mapped AS link for a whole year. We were able to detect congestion events and quantify their periods and magnitudes at four IXPs. We then verified the events and investigated the causes by interviewing the IXP operators. Our results show that only 2.2% of the discovered IP links experienced (sustained or transient) congestion during our measurement period. Our findings suggest the need for ISPs to carefully monitor the provision of their peering links, so as to avoid or quickly mitigate the occurrence of congestion. Regulators may also define the maximum level of packet loss in those links to provide some protection to communications routed through local IXPs.

CCS CONCEPTS

• **General and reference** → **Measurement**; • **Networks** → **Network measurement**; **Network monitoring**;

KEYWORDS

Congestion, IXP, Performance.

ACM Reference format:

Rod  rick Fanou, Francisco Valera, and Amogh Dhamdhare. 2017. Investigating the Causes of Congestion on the African IXP substrate. In *Proceedings of IMC '17, London, United Kingdom, November 1–3, 2017*, 7 pages. <https://doi.org/10.1145/3131365.3131394>

1 INTRODUCTION

The growing popularity of bandwidth-hungry applications such as streaming video has generated renewed interest in understanding the nature, location, and causes of performance degradations in the Internet infrastructure. In the US or in Europe, previous studies have found that congestion often occurs at the boundaries between networks, due to disputes about which party should pay to upgrade the infrastructure necessary for carrying traffic [16, 28]. However,

much less is known about the nature of congestion and its causes at Internet Exchange Points (IXPs), particularly those located in developing regions such as Africa. While the possibility of performance problems due to congestion is not unique to IXPs and could also occur within ISPs operating in these regions, IXPs are of particular interest due to their position as hubs that facilitate traffic exchange between hundreds of connected networks. As there is a great push to promote peering at IXPs in Africa [2, 13, 14, 25], it is of interest to quantify the performance at those infrastructures. The absence of congestion or performance problems on the IXP may help to dispel the doubts of ISPs that are still reluctant to join local IXPs. In cases where there is evidence of poor performance, it is also important to be aware of the causes (peering disputes or other reasons).

To fill the lack of congestion related measurements at IXPs in Africa, we selected six IXPs located in three of the five African sub-regions [1, 3] for reasons discussed in §3. Notably, the only sub-regions involved are West, East, and Southern Africa, as we were not able to find hosts to deploy our probes at IXPs in the other African sub-regions. We used techniques allowing continuous, fine-grained, and longitudinal measurements. Using time-sequence latency probes (TSLP) [28] measurements run from Ark probes deployed at those selected IXPs over a year (from February 2016 to April 2017), we inferred whether any of the discovered links were congested. We then evaluated the extent to which this phenomenon influenced RTTs to the near and far ends of those links and the characteristics of the observed patterns. We also evaluated the impact on the links in terms of packet loss. Finally, we investigated the causes by interviewing the IXP operators.

We detected cases of congestion at four IXPs. We show how RTTs and loss rates to the far end increase drastically during the congestion events, and delve into the root cause of the observed congestion. Although we did not find any evidence of widespread congestion (only 2.2% of the measured links showed evidence of congestion during our measurement period), our findings suggest the need for ISPs to monitor the provisioning of their peering links to avoid or quickly mitigate the occurrence of congestion. Regulators may also define the maximum permissible level of packet loss in those links with the goal of improving performance at local IXPs and thereby making those facilities attractive hubs for local interconnection.

The rest of this paper is organized as follows. We discuss related work in §2. We then give a description of our measurement infrastructure in §3. In §4 we present the data collection process as well as the AS link inference and validation. We detail the analysis performed on the dataset in §5. In §6 we present case studies of congestion discovered on links probed from our vantage points (VPs) located at two of the studied IXPs, discuss their causes, and

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

IMC '17, November 1–3, 2017, London, United Kingdom

   2017 Association for Computing Machinery.

ACM ISBN 978-1-4503-5118-8/17/11...\$15.00

<https://doi.org/10.1145/3131365.3131394>

examine their consequences. Finally, we conclude and outline our perspectives for future work in §8.

2 RELATED WORK

Chetty *et al.* [11] measured broadband performance in South Africa using measurement software implemented on mobile phones and home routers. They found that users in South Africa do not get advertised speeds, and the interconnection (or lack of it) between local ISPs mainly determines reliability and user performance. Gupta *et al.* [18] found that 66.8 % of paths from their vantage points toward Google caches, both located in the region, detoured through Europe. Using RIPE Atlas probes scattered throughout Africa, Fanou *et al.* [13] underlined the reliance on ISPs based outside the continent for serving intra-continental traffic, detected the launch of new IXPs, and showed the positive impact of peering at local IXPs on AS path lengths and delays. Chavula *et al.* [22] examined, using Ark monitors, communications among African research networks. They found that 75 % of the paths are routed via Europe and the US, increasing RTTs by 150 ms on average. More recently, Fanou *et al.* performed a longitudinal study [14] of the evolution of connectivity among local African networks from 2013 to 2016, highlighting the prevailing dominance of intercontinental ASes, mapping both existing and newly launched IXPs, and exhibiting with diverse case studies the significant improvements in QoS resulting from more peering in the region.

3 MEASUREMENT INFRASTRUCTURE

The time-sequence latency probes (TSLP) method [28] consists of frequently performing round trip time (RTT) measurements from a vantage point (VP) within a network to the near and far routers of an interdomain link of the hosting network. If the interdomain link is congested, then the buffer occupancy at the link increases and RTTs measured across the link also increase. We can thus infer from a pattern showing an increase of RTTs to the far end of an AS link (but no increase to the near end of the link), that a queue between the routers on both sides of the link caused the observed delay. Luckie *et al.* [28] used traffic data from a research network to validate the TSLP technique. An advantage of adopting the said technique is that it makes it possible for a vantage point placed within or at the edge of a network to monitor congestion in that network without explicit cooperation from the network operator.

We used the CAIDA Archipelago (Ark) measurement platform [9] because of its capacity to perform fine-grained measurements. Ark allows us to run scamper [27] on its monitors (*e.g.*, for limiting the TTL value of the ICMP packets, sending a burst of packets through the congested link, etc.) and to gain more visibility on the events occurring on the IP layer, while performing a longitudinal TSLP-based study. We used probes deployed at six African IXPs [15, 42] located in three of the five African sub-regions: Ghana Internet eXchange Association (GIXA, launched in Ghana in 2005) [17], Johannesburg Internet eXchange (JINX, South Africa, 1996) [21], Kenya IXP (KIXP, Kenya, 2002) [41], Serekunda IXP (SIXP, Gambia, 2014) [38], and Tanzania IXP (TIX, Tanzania, 2004) [39].

These are interesting IXPs, as (i) they are mature and large Internet markets that will allow us to analyze a large number of peering

links from each VP, (ii) they may become regional IXP hubs in the near future, as they have the potential to attract more members. By *regional IXP hub* we refer to a local IXP at which most networks operating in a sub-region peer, which thus helps localize traffic among countries located in that sub-region. We deployed the VPs in two different settings: some (VP1–3) are plugged into the content network of the IXP. The term *content network* refers to the network, usually connected to the IXP switches, which hosts all resources designed to offer common services to the members *i.e.*, cache instances, Internet portals, search engines, NTP servers, routing registry, looking glass, etc. Commonly, the content network is not separated from the peering network. As we show later (§6.2.1), this is not the case for all IXPs. From VPs deployed on the content network of the IXP, we expect to discover all networks accessing the content available at the IXP. Others (VP4–6) are hosted by ASes that peer at the IXPs. From these VPs we expect to discover, among others, the peers of the host network at the IXP.

4 DATA COLLECTION

We detail how we collected the studied dataset and how we inferred and validated the AS relationships.

We automatically inferred the host networks' boundaries and discovered their respective border links using the CAIDA's border mapping tool *bdrmap* [10, 29]. The border mapping process first gathers routing and addressing data used for data collection and analysis. The input datasets are prefix-AS mappings constructed from public BGP data (RouteViews [30] and RIPE RIS [36]), CAIDA's AS-rank algorithm [8] used to infer AS relationships, delegation files published by the 5 Regional Internet Registries (RIRs) [4, 5, 7, 26, 35], a list of IXP prefixes from PeeringDB [33] and Packet Clearing House (PCH) [31], and a list of sibling ASes of the VP's AS. The creation of the sibling list is a semi-manual process seeded with CAIDA's AS-to-organization mapping to which we then manually add missing siblings and remove spurious ones. *bdrmap* uses an efficient variant of traceroute to trace the path from each VP to every routed prefix observed in BGP. It then applies alias resolution techniques to infer routers and point-to-point links used for interdomain interconnection. This collected data is used to assemble constraints that guide the execution of heuristics to infer router ownership. The border mapping process aims to obtain sufficient information about the links observed from the VP's AS toward every other AS to constrain the border router inferences. To validate the *bdrmap* output we first checked the inferred links against public datasets [20, 32, 33, 37]. We emailed the probe hosts for cross-checking when our results were in contradiction with those public datasets. Four of the six involved VP hosts responded to our queries. They also gave us more insights into the setup of links that the border mapping process had correctly discovered (§6.2). This cooperation allowed us to better analyze the collected data. On average the border mapping process correctly discovered 96.2 % of the neighbors of the VP networks.

Next, we periodically probed both ends of each discovered IP link every 5 minutes using TTL-limited probes set to expire at the near and far ends of the link. Regarding ethical considerations, we ensured that our measurements would not adversely affect the VP network by using a low probing rate (small packets sent at the rate of 100 packets per second). Moreover, the targets of the probing traffic

(both ends of each mapped IP link) do not put the Ark probe hosts at risk. Our probes do not collect traffic data or any information that may be considered sensitive due to privacy reasons. We ran our measurements for a year from 22/02/2016 to 27/03/2017. In §5.2 we detail how we detect congestion events from the analysis of these TSLP measurements. In case we detected repeated occurrences of congestion on a link (§5.2), we set up measurements attempting to measure packet loss on those links by probing both ends of those links at a higher rate, *i.e.*, one packet per second, and then computed the loss rate over every batch of 100 probes. We ran the loss rate measurements from 19/07/2016 to 01/04/2017; these began roughly 5 months after latency measurements, because we made sure the targeted links were all suffering from repetitive congestion events before launching the loss rate measurements.

5 DATA ANALYSIS

5.1 Evolution of number of discovered links

For each VP we identified the links discovered from that VP that were at the IXP, since some VPs are hosted by an IXP member, whereas others are in the content network of an IXP (§3). To achieve this, we categorized the links having any of their IPs belonging to the (peering or management) prefix of any studied IXP as links established at those IXPs. After that, we validated the *bdrrmap output* (§4) with the corresponding IXP operator and inspected the evolution of the number of neighbors of the VP’s AS over time (§6.1). We also geolocated both IPs of each link using the Netacuity Edge Database [12] and hints in Reverse DNS outputs [19, 34] as added checks that those links were indeed established at the IXPs.

5.2 Analysis of congestion cases

We began by gathering the RTT time series collected in §4 per VP and discovered neighbor. Then, we applied an algorithm to detect *level shifts* in the measured time series, which indicate that the router queue at the interdomain link was filling up, possibly due to the link being congested. The level-shift algorithm identifies changes in the direction of the rank-based non-parametric statistical cumulative sum (CUSUM) test [40] as evidence of a level-shift. We tuned the algorithm to use 5-minute latency samples and to detect level shifts that last at least 30 minutes. The magnitude of a level shift that results from congestion corresponds to the size of the router buffer. We impose a threshold on the minimum magnitude of the level shifts that we label as potentially caused by congestion. The objective of this threshold is to eliminate false detections that result from noise in the RTT times series or slow ICMP response generation from the routers. Next, we show that this objective is achieved reasonably well with a threshold of 10 ms.

We inspected the sensitivity of selecting 10 ms as opposed to 5 ms, 15 ms, or 20 ms by analyzing the variation in the number of inferred congested links. For each value of the threshold, we obtained the links flagged as potentially congested (Table 1), and manually checked whether those links had a persistent diurnal pattern indicating peak-hour congestion. We flagged 11.2 % more links as potentially congested when using 5 ms; however, the number of links for which we identified a recurring diurnal pattern was the same as that with a 10 ms threshold. In contrast, we flagged 50 % fewer links with recurring diurnal patterns when using

a 15 ms or 20 ms threshold. Finally, we contacted the IXP operators to confirm whether 10 ms was a reasonable threshold. We received 2 responses; both stated that they considered 10 ms a reasonable threshold.

Table 1: Sensitivity analysis of the threshold value used for labeling potentially congested links in our datasets.

VP	# Potentially congested links (with a diurnal pattern) for a threshold of			
	5 ms	10 ms	15 ms	20 ms
VP1	4 (2)	4 (2)	3 (1)	2 (1)
VP2	6 (2)	5 (2)	4 (1)	3 (1)
VP3	80 (1)	56 (1)	48 (1)	40 (1)
VP4	2 (1)	1 (1)	0 (0)	0 (0)
VP5	147 (0)	147 (0)	147 (0)	146 (0)
VP6	100 (0)	88 (0)	88 (0)	71 (0)
All VPs	339 (6)	301 (6)	290 (3)	262 (3)

To analyze the flagged links that presented recurring diurnal patterns, we ensured that we detected no level shift on the near side, which would mean that the observed congestion was not at the targeted link. In this step we also tagged for further analysis links showing unclear patterns, *i.e.*, RTTs to the far end presented a diurnal waveform, while those to the near end were inconclusive. To make robust inferences about whether any observed congestion was at the targeted links, we used the Record-routes method [24, 28] to check path symmetry, thereby ensuring that an increase in RTTs from a near to a far router was solely due to traffic on that link.

We then investigated the level shift sensitivity to decide whether to directly use its output to calculate the width of the congested period, or to sanitize it before doing so. We computed the average magnitude A_w and the average duration Δt_{UD} between consecutive upshift and downshift. For links showing recurring diurnal patterns, we investigated whether congestion had a measurable effect on packet loss. Finally, we interviewed the IXP operators to validate and corroborate the obtained results as well as the suggested causes.

6 RESULTS AND DISCUSSION

We summarize our measurements per IXP and quantify how many observed links experienced congestion during the study (Table 2). We then shed light on the evolution of the number of discovered links, AS neighbors, and peers of each VP’s AS. Then we analyze in depth the most interesting results per VP, characterizing whether the congestion was sustained or transient, the impact on packet loss rate, and the causes of the observed phenomenon.

6.1 Evolution of number of discovered links

We summarize per VP in Table 2, the total number of discovered IP links, inferred IP peering links, as well as AS neighbors, and peers obtained from the border mapping process (§4), when considering three snapshots. Discovered IP links gather all router-level links found to connect the VP’s AS to that of any of its neighbors. Inferred IP peering links are the subset of discovered IP links having any side that belongs to the IXP prefix (§5). The number of neighbors and peers of the AS host are the highest (1,215 and 197 respectively) for our VP in AS30844 (Liquid Telecom) that peers at KIXP. We noticed that the number of neighbors and peers decreased from 13 on 17/03/2016 to 7 on 15/11/2016 for AS30997 (GIXA): this

Table 2: Evolution of the number of discovered IP links, AS neighbors, and peers per vantage point.

ID	IXP Country host (African sub-region)	IXP name (IXP-AS)	Measurements Duration (Total # traceroutes)	Total # record routes	AS hosting the probe (AS name)	Total # snap-shots	Snapshots dd/mm/yyyy	# Discovered IP (peering) links	# Congested IP peering links	# Neighbors (peers)
VP1	Ghana (West Africa)	GIXA (AS30997)	27/02/2016 to 27/03/2017 (241,848,566)	34,343	AS30997 (GIXA)	397	17/03/2016 18/06/2016 15/11/2016	46 (36) 13 (13) 10 (10)	2 1 1	13 (13) 8 (8) 7 (7)
VP2	Tanzania (East Africa)	TIX (AS33791)	28/02/2016 to 27/03/2017 (597,083,978)	166,605	AS33791 (TIX)	991	19/03/2016 18/06/2016 16/11/2016	59 (59) 98 (98) 36 (36)	2 2 0	31 (26) 30 (30) 36 (29)
VP3	South Africa (Southern Africa)	JINX (AS37474)	05/03/2016 to 27/03/2017 (555,641,317)	209,250	AS37474 (JINX)	889	27/07/2016 15/11/2016 19/02/2017	193 (171) 212 (130) 212 (120)	1 0 0	32 (27) 42 (42) 44 (39)
VP4	Gambia (West Africa)	SIXP (AS327719)	22/02/2016 to 27/03/2017 (89,387,074)	0	AS37309 (QCell)	127	18/03/2016 22/07/2016 07/09/2016	14 (11) 4 (3) 6 (5)	1 1 1	7 (6) 4 (3) 6 (5)
VP5	Kenya (East Africa)	KIXP (AS4558)	25/02/2016 to 27/03/2017 (415,583,808)	103,392	AS30844 (Liquid Telecom)	668	11/03/2016 23/03/2017 07/04/2017	288 (4) 9,754 (557) 10,466 (601)	0 0 0	244 (4) 1,208 (199) 1,215 (197)
VP6	Rwanda (East Africa)	RINEX (AS37224)	08/07/2016 to 27/03/2017 (200,749,695)	0	AS37228 (RDB)	318	27/07/2016 15/11/2016 19/02/2017	79 (4) 82 (4) 72 (4)	0 0 0	9 (1) 9 (1) 9 (1)

drop is due to the commercialization of the content network of the IXP (§6.2.1), causing the disconnection of non-registered members. Meanwhile, AS33791 (TIX) and AS37228 (RINEX) have a roughly constant number of peers over our measurement period.

Table 2 also presents the number of detected congested links. *Congested links* are those for which RTTs to the far end show a recurring diurnal pattern, whereas those to the near end stay constant. A congestion case that is later mitigated is described as being *transient* in the rest of this paper; otherwise we refer to it as *sustained*. While for the first four probes, we found one or more cases of congested links, we did not detect any cases for the last two (VP5 and VP6). In fact, the fraction of observed links that experienced any congestion is at most 7.7 % for VP1, 3.3 % for VP2, 0.6 % for VP3, and 33 % for VP4. In total, 2.2 % of the discovered IP peering links experienced congestion. We thus did not find any evidence of widespread congestion. That said, we analyze in depth, in the next section, striking congestion cases observed from VP1 and VP4, while highlighting their causes and consequences.

6.2 Analysis of congestion cases

6.2.1 Cases seen from VP1 deployed at GIXA. Only two of the links mapped by VP1 hosted at GIXA [17] experienced congestion during our measurement period: the links to GHANATEL (Vodafone, AS29614) and KNET (AS33786).

GIXA – GHANATEL

The waveform registered for the first link presents different amplitudes over a total of roughly 5 months. First, RTTs to the far end sometimes peak at 20 ms and 50 ms at other times, while those to the near end remain low and constant during the first 3.5 months (03/03/2016 to 14/06/2016) termed *phase 1*. Figure 1 shows part of phase 1. Our analysis of the record-routes (RR) probes during that period gives us confidence that the route is symmetric. Since the RR probes showed symmetry, then the peak on top of the peak depicted by the shape of the red curve of Figure 1 is interesting: it likely indicates congestion in both directions on the link.

From the level shifts that occurred periodically between 15/03/2016 to 14/06/2016, we inferred the average magnitude A_w

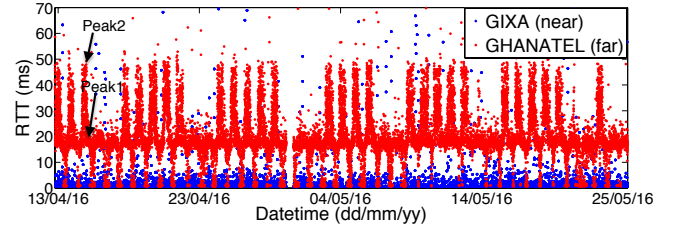


Figure 1: RTTs GIXA–GHANATEL in part of *phase 1* of the shifts to be 27.9 ms and Δt_{UD} , roughly 20 hours, implying long congestion events. While discussing with the IXP operator about the possible causes of such phenomenon, the operator first informed us that the GIXA peering and content networks are separated. The content network (hosting VP1) contains Google caches (GGCs) that need to be updated through transit links. In *phase 1*, GHANATEL was the ISP providing the required transit services through a 100 Mbps link, whereas its clients were served by its main peering link of 1 Gbps size. The 100 Mbps transit link was the one identified by our measurements as suffering from congestion. Thus, GHANATEL users [6] were likely not directly impacted during *phase 1*.

The amplitude of the waveform then dropped to 10 ms from 15/06/2016 to 06/08/2016 (the date from which our latency probes to the far end were unsuccessful): we term this *phase 2* (figure 2a). The beginning of this period coincides with the shutdown of the transit service. The IXP operator explained that GHANATEL shut off the transit service to force the IXP to pay for it. GHANATEL then used that link for peering until early October, leaving the GGCs non-functional. We observed a diurnal pattern confirmed by the loss rate increase during that phase (figure 2b). Though figure 2b depicts loss rate up to 25%, our measurements show that loss rate varied between 0% and 85% between 21/07/2016 and 06/08/2017. We conjecture that during *phase 2*, GHANATEL end-users may have been affected by the congested peering link; in addition, all end-users of networks peering at GIXA may also have been affected by the detour of their packets while accessing Google content, which was no longer cached at the IXP.

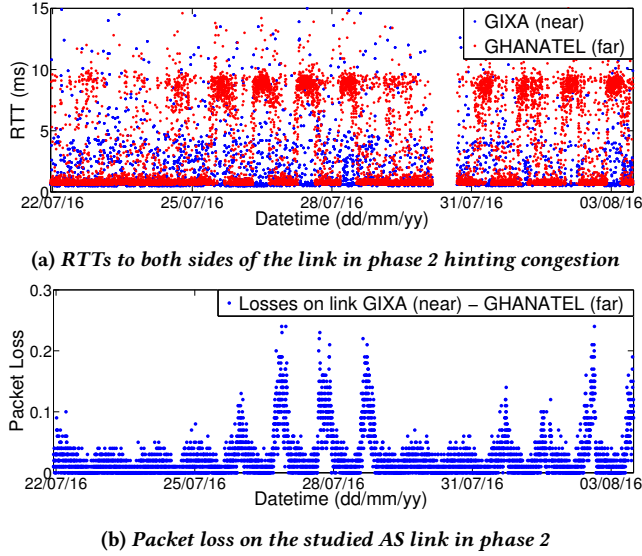


Figure 2: RTTs and losses GIXA–GHANATEL

In early October, GHANATEL stopped using the problematic link. This corresponds to a change made by the IXP wherein they started using an intercontinental ISP as a transit provider for the GGC with a higher capacity link of 620 Mbps; the IXP is now paying for the transit services and members of the IXP are now required to register in order to access content. This policy change led to the decrease in the number of peers connected to the content network as noted in §6.1 and in Table 2.

Finally, we noticed that in both phases, the elevation in far-end RTTs correlates with days of the week. For *phase 1*, the five large spikes correspond to the business days, whereas the rest, to those of the weekend (figure 1). As congestion events occurred till the shutdown of the link, the congestion was *sustained*.

GIXA – KNET

We now consider the link GIXA–KNET, for which figure 3 presents RTTs to both ends of the link, along with the loss rates. To begin with, KNET delivers high quality video, data, and voice solutions throughout West and Central Africa [23]. Its link with GIXA was discovered by bdrmap on 29/06/2016. From 06/08/2016, RTTs to the far end present a diurnal waveform, while those to the near end remain constant and stay below 1 ms (figure 3a). We observed the same pattern consistently until the end of our measurements for a total of approximately 8 months. The analysis of record route (RR) probes during that period provided evidence of route symmetry for the duration of our measurements. After that, we evaluated the characteristics of the waveform to find that A_w is 17.5 ms, while Δt_{UD} is of 2 hours 14 min after level shifts sanitization *i.e.*, a single congestion event lasts roughly 2 hours.

One might assume, since we started seeing evidence of congestion on the GIXA–KNET link on the same day (06/08/2016) as the link GIXA–GHANATEL disappeared, that there is a causal relationship between the two events. Further investigation showed that this was not the case: although KNET has a regional footprint, it does not provide transit. On 06/10/2016, during the GIXA operator interview, the operator told us that they did not believe the KNET

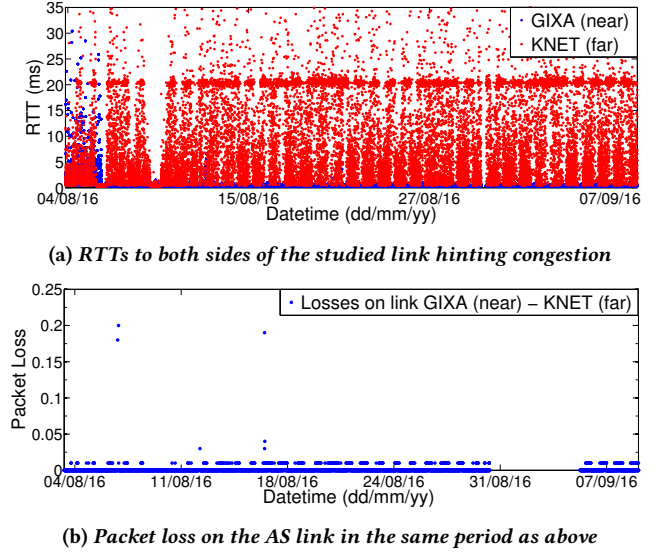


Figure 3: RTTs and losses GIXA–KNET

port at the IXP was congested. In such a context, there may be other causes for the observed phenomenon, which would need cooperation with KNET to investigate fully: (i) whether the KNET router is overloaded at peak times, resulting in slow ICMP responses or (ii) whether the link with the GIXA content network is congested.

On 05/05/2017, KNET informed the IXP that they are not experiencing congestion, and expressed that they have not received any complaints from their customers accessing content. The lack of complaints may be explained by the fact that the average loss rate measured on the link from 21/07/2016 to 29/03/2017 (figure 3b) is low (0.1 %). The observed pattern is the same regardless of the type of the day (business or not). It shows an obvious decrease everyday around midnight, an increase at different times of the day, and a constant RTT value around 20 ms in the afternoon. As this pattern is observed till the end of the campaign, the phenomenon was *sustained*.

6.2.2 Case seen from VP4 in QCELL at SIXP: QCELL–NETPAGE. VP4 is hosted within QCELL (AS37309), a SIXP member. Previously, Fanou *et al.* [13] found in August 2014 (a month after the launch of SIXP) that RTTs between QCELL and NETPAGE were constant around 1.5 ms. However, we noticed that the RTTs across that link showed repeating diurnal patterns from 29/02/2016 to 28/04/2016 (*phase 1*, shown in Figure 4a) indicating congestion on the link. From 28/04/2016 to 30/03/2017, the diurnal waveform disappeared and most RTT values were below 10 ms (*phase 2*).

While interviewing the SIXP operator, we were told that during *phase 1*, the demand to access the GGCs (for which QCELL provides transit) from NETPAGE was huge: NETPAGE’s engineers noticed that high bandwidth usage by Google traffic from their users was degrading performance and causing congestion. They thus asked for an upgrade of their link with SIXP from 10 Mbps to 1 Gbps. After the upgrade (done on 28/04/2016 according to our data), the congestion events disappeared and were not evident until the end of the measurement period (figure 4b). We believe NETPAGE’s users may have been affected by these congestion events.

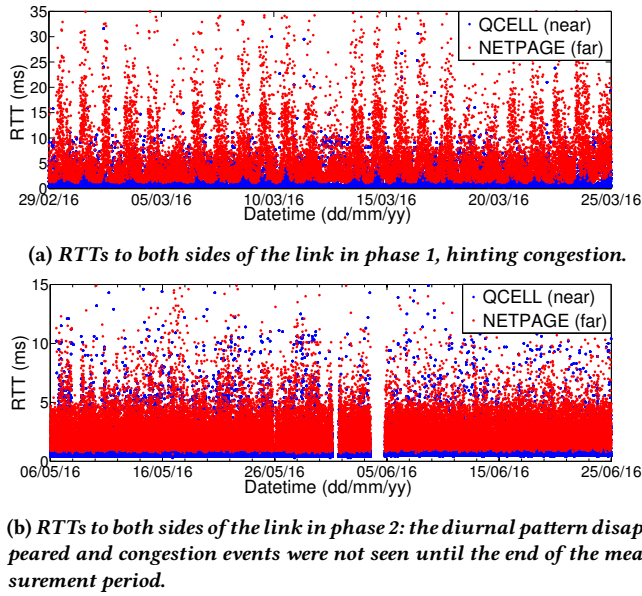


Figure 4: RTTs QCELL–NETPAGE

Regarding the characteristics of the waveform, the average magnitude A_w of the level shift during *phase 1* was 10.7 ms, with a periodicity of approximately 1 day. Moreover, congestion events lasted on average a third of the duration of those registered during *phase 1* for the link GIXA–GHANATEL (§6.2.1), since Δt_{UD} was 6 hours 22 min. Finally, we noticed that the waveform was the same over weeks (figure 4a) with a spike corresponding to each day. The height of the spike reached 35 ms in the week, whereas it stayed around 15 ms during the weekends. The reasons behind this may be increased access to Google content for daily activities combined with communications among clients of both ISPs during business days as compared to weekends.

7 IMPLICATIONS OF OUR RESULTS

In this section, we highlight the takeaways of our work and discuss the implications for research and network operations.

A key takeaway of our study is that we observed congestion on only a small fraction of the monitored links during this measurement period. However, we also noted that the IXP ecosystem is highly dynamic in Africa, as ISP presence at IXPs, policies employed by the IXPs themselves, and the presence of content providers can change over time. Further, with the push for peering in the African region, it is likely that the IXP substrate will become more mature in the future, supporting more peering between interconnected networks and hence increased traffic volumes. All these factors motivate the need for longitudinal measurement and monitoring of this evolving infrastructure.

We showed that the TSLP technique can detect congestion without requiring access to data from network operators. However, we emphasize that judicious interpretation of the *causes* of the observed congestion events requires collaboration and validation from the stakeholders, as these are often related to hidden events that are not made public. We found congestion on a link used to update Google caches hosted at the IXP, on a link used by an ISP to peer at the IXP

(e.g., VP1), or on an under-provisioned link connecting a Google cache to one of the IXP peers (e.g., VP4). High demand appears to be the main cause in the last scenario. In the two first cases, congestion was sustained; in the case of the link GIXA–GHANATEL, there was a dispute between the two parties, while in the case of the GIXA–KNET link, the low packet loss on the link likely meant that end-users were not severely impacted and hence the ISP did not upgrade the link.

As for implications for network/IXP operators, we learned that (i) when considering links at IXPs, links used to access content are susceptible to congestion; hence, they need to be monitored more carefully, and (ii) local IXP operators willing to host content caches must be aware that they need transit services to be functional; such a situation may lead to dispute with the provider if not well managed; e.g., in case of increase in the demand without any update of the Service level Agreement or if demand increase is combined with a free provision of transit services.

8 CONCLUSIONS AND FUTURE WORK

We investigated the causes of congestion and measured its impact on the African IXP substrate using vantage points deployed at six strategically selected IXPs in Africa. While we detected cases of congestion at four IXPs, we did not find evidence of widespread congestion; only 2.2% of the discovered links experienced congestion during our measurement period. We then detailed the most interesting case studies and discussed the implications for both research and network operations. Although our findings regarding the causes of congestion at IXPs may apply to IXPs in other regions, we prefer not to attempt to generalize them beyond what we could directly observe and validate with the operators.

Since an IXP only monitors ports sizes/traffic or ensures upgrades upon requests from ISPs, it is important that ISPs carefully monitor their peering links at IXPs to avoid or to quickly mitigate congestion (as noticed for VP4). We plan to continue deploying additional Ark probes at networks and IXPs operating in developed and developing regions, including Africa, to increase our coverage of the African sub-regions that have not received much attention so far. Meanwhile, we plan to keep analyzing collected TSLP data to delve into the dynamics and causes of congestion at IXP infrastructure, and compare the results with those obtained in this work. Finally, it will be interesting to correlate our observations from TSLP measurements with data from the IXP operators. To this end, we are working on strengthening our relationship with operators in the region to make such a study feasible in the future.

9 ACKNOWLEDGEMENTS

We thank our shepherd Cristel Pelsser and the anonymous reviewers for their insightful comments. We are grateful to all the probe hosts and to those who participated in our interviews. Rod  rick Fanou was partially supported by IMDEA Networks Institute, US NSF grant CNS-1414177, and the project BRADE (P2013/ICE-2958) from the Board of Education, Madrid Regional Government. Amogh Dhamdhare was partially funded by US NSF grant CNS-1414177. Francisco Valera was partially funded by the European Commission under FP7 project LEONE (FP7-317647).

REFERENCES

- [1] Africa Program. African Regional and Sub-Regional Organizations: Assessing Their Contributions to Economic Integration and Conflict Management. Technical report, Woodrow Wilson International Center for Scholars, 2008.
- [2] African Union (AU). African Internet eXchange System. www.au.int/web/en/axis, 2017.
- [3] African Union Commission and New Zealand Ministry. African Union Handbook 2017. Technical report, African Union (AU), 2017.
- [4] AfriNIC. AfriNIC Database. <ftp://ftp.afrinic.net/>, 2017.
- [5] APNIC. APNIC database. <ftp://ftp.apnic.net/pub/stats/apnic/>, 2017.
- [6] APNIC. Visible ASNs: Customer Populations (Estimation). <http://stats.labs.apnic.net/cgi-bin/aspop?c=>, 2017.
- [7] ARIN. ARIN database. <ftp://ftp.arin.net/pub/stats/arin/>, 2017.
- [8] CAIDA. Automated Autonomous System (AS) Ranking. Research Project. <http://as-rank.caida.org>, 2015.
- [9] CAIDA. Archipelago (Ark) Measurement Infrastructure. <http://www.caida.org/projects/ark/>, 2017.
- [10] CAIDA. Border Mapping (bdrmap) Dataset. http://www.caida.org/data/active/bdrmap_dataset.xml, 2017.
- [11] M. Chetty, S. Sundaresan, S. Muckaden, N. Feamster, and E. Calandro. Measuring Broadband Performance in South Africa. In *Proceedings of the 4th Annual Symposium on Computing for Development*. ACM, 2013.
- [12] Digital Element. Netacuity. http://www.digital-element.net/ip_intelligence/ip_intelligence.html, 2017.
- [13] R. Fanou, P. Francois, and E. Aben. On the Diversity of Interdomain Routing in Africa. In *International Conference on Passive and Active Network Measurement (PAM)*, 2015.
- [14] R. Fanou, P. Francois, E. Aben, M. Mwangi, N. Goburdhan, and F. Valera. Four Years Tracking Unrevealed Topological Changes in the African Interdomain. *Computer Communications*, 2017.
- [15] R. Fanou, V. Sánchez-Agüero, F. Valera, M. Mwangi, and J. Coffin. African Route-collector Data Analyzer (ARDA). <https://arda.af-ix.net/>, 2017.
- [16] D. Genin and J. Splett. Where in the Internet is Congestion? <http://arxiv.org/abs/1307.3696>, 2013.
- [17] Ghana Internet Exchange Association (GIXA). Ghana Internet Exchange Association (GIXA) Website. www.gixa.org.gh/, 2017.
- [18] A. Gupta, M. Calder, N. Feamster, M. Chetty, E. Calandro, and E. Katz-Bassett. Peering at the Internet's Frontier: A First Look at ISP Interconnectivity in Africa. In *International Conference on Passive and Active Network Measurement (PAM)*, 2014.
- [19] B. Huffaker, M. Fomenkov, and K. Claffy. Geocompare: a Comparison of Public and Commercial Geolocation Databases. *Proc. NMMC*, 2011.
- [20] Hurricane Electric (HE). Hurricane Electric Internet Services: BGP Toolkit Home. <http://bgp.he.net/>, 2017.
- [21] Internet Service Providers' Association (ISPA). Internet Exchange. <http://ispa.org.za/inx/>, 2017.
- [22] C. Josiah, 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. In *e-Infrastructure and e-Services for Developing Countries*. Springer, 2014.
- [23] K-NET. K-NET. <http://www.knetgh.com>, 2017.
- [24] E. Katz-Bassett, H. V. Madhyastha, V. K. Adhikari, C. Scott, J. Sherry, P. Van Wesep, T. E. Anderson, and A. Krishnamurthy. Reverse Traceroute. In *NSDI*, volume 10, 2010.
- [25] M. Kende and C. Hurpy. Assessment of the Impact of Internet Exchange Points (IXPs) - Empirical Study of Kenya and Nigeria. *Internet Society (ISOC)*, (59), 2012.
- [26] LACNIC. LACNIC database. <ftp://ftp.lacnic.net/pub/stats/lacnic/>, 2017.
- [27] M. Luckie. Scamper: a Scalable and Extensible Packet Prober for Active Measurement of the Internet. In *Proceedings of ACM SIGCOMM Internet Measurement Conference (IMC)*, 2010.
- [28] M. Luckie, A. Dhamdhere, D. Clark, B. Huffaker, and K. Claffy. Challenges in Inferring Internet Interdomain Congestion. In *Proceedings of ACM SIGCOMM Internet Measurement Conference (IMC)*, 2014.
- [29] M. Luckie, A. Dhamdhere, B. Huffaker, D. Clark, and K. Claffy. bdrmap: Inference of Borders Between IP Networks. In *Proceedings of ACM SIGCOMM Internet Measurement Conference (IMC)*, 2016.
- [30] D. Mayer. University of Oregon RouteViews Archive Project. <http://routeviews.org>, 2017.
- [31] Packet Clearing House (PCH). <https://www.pch.net/>, 2017.
- [32] Packet Clearing House (PCH). PCH IXP directory. http://prefix.pch.net/images/applications/ixpdir/ip_asn_mapping.txt, 2017.
- [33] PeeringDB. <https://www.peeringdb.com/>, 2017.
- [34] I. Poesse, S. Uhlig, M. A. Kaafar, B. Donnet, and B. Gueye. IP Geolocation Databases: Unreliable? *ACM SIGCOMM Computer Communication Review*, 41(2), 2011.
- [35] RIPE NCC. RIPE NCC database. <ftp://ftp.ripe.net/ripe/stats/>, 2017.
- [36] RIPE NCC. RIPE RIS. <https://www.ripe.net/analyse/internet-measurements/routing-information-service-ris/>, 2017.
- [37] RIPE NCC. RIPE Stats. <https://stat.ripe.net/>, 2017.
- [38] Serekunda Internet Exchange Point (SIXP). <http://www.sixp.gm/>, 2017.
- [39] Tanzania Internet Service Providers Association (TISPA). Tanzania Internet eXchange - TIX. <https://www.tix.or.tz/>, 2017.
- [40] W. A. Taylor. Change-Point Analysis: A Powerful New Tool for Detecting Changes. <http://www.variation.com/cpa/tech/changepoint.html>, 2000.
- [41] Telecommunications Service Providers of Kenya (TESPOK). Kenya Internet Exchange Point (KIXP). <https://www.tespok.co.ke/>, 2017.
- [42] The African IXP Association (Af-IX). List of Active Internet eXchange Points in Africa. <http://www.af-ix.net/ixps-list>, 2017.