Peering at the Internet's Frontier: A First Look at ISP Interconnectivity in Africa

Arpit Gupta[1], Matt Calder[2], Nick Feamster[1], Marshini Chetty[3,4], Enrico Calandro[4], and Ethan Katz-Bassett[2]

[1] Georgia Tech [2] University of Southern California
[3] University of Maryland [4] Research ICT Africa

Abstract. In developing regions, the performance to commonly visited destinations is dominated by the network latency, which in turn depends on the connectivity from ISPs in these regions to the locations that host popular sites and content. We take a first look at ISP interconnectivity between various regions in Africa and discover many circuitous Internet paths that should remain local often detour through Europe. We investigate the causes of circuitous Internet paths and evaluate the benefits of increased peering and better cache proxy placement for reducing latency to popular Internet sites.

1 Introduction

An Internet user's experience depends on having reliable paths that offer good performance to the set of sites that a user commonly visits. In the developed world, we have grown accustomed to rich peering and interconnection between a variety of ISPs, ranging from access networks to content providers. In developing regions, however, the story is more nuanced. The servers that host popular content (*e.g.*, Facebook, Google) may be distant, sometimes even on a different continent. Even when an Internet destination or content *is* nearby, a user's path to that destination may be circuitous, if two ISPs do not connect directly. For example, a user's path between South Africa and Kenya might "detour" through Europe if the African ISPs do not directly connect with one another. These pathologies can significantly degrade Internet performance and can affect decisions about how and where to place Internet content and services. For example, a content provider might deploy a cache in a particular ISP in one region, expecting to serve a large group of users in that region. Yet, if many Internet paths to that ISP from nearby users are circuitous, the cache will provide limited benefit to local users.

In this paper, we characterize the nature of interdomain Internet connectivity in Africa, focusing in particular on the connectivity at two major local Internet exchange points in Africa, JINX in Johannesburg and KIXP in Nairobi. To construct a view of Internet paths between various destinations in Africa, we perform continual traceroute measurements from BISmark routers in South Africa to the Measurement Lab servers deployed in South Africa, Kenya, and Tunisia. We use these traceroutes to explore the extent to which Internet paths that would otherwise remain local ultimately take a circuitous path through a remote exchange point (typically in Europe). We then use more detailed BGP routing information from RouteViews, Packet Clearing House (PCH), and Hurricane Electric, as well as information from PeeringDB and the IXP websites, to help explain *why* these Internet paths are circuitous.

We also quantify whether better peering or more extensive placement of caching proxies (or both) can reduce latency to popular services. We perform a trace-driven emulation to recommend specific steps for improving the performance of specific services in Africa. For example, we observe that the deployment of a Google cache node in a particular ISP will provide no benefit to Internet users in other local ISPs *unless* that ISP makes specific peering arrangements with ISPs hosting cache nodes. Based on this observation, we study how certain "surgical" peering arrangements within Africa could improve Internet performance by short-circuiting longer routes via European IXPs. Specifically, we study the following two questions:

- *What is the nature of interdomain Internet paths between locations in Africa? (Section 3)* We characterize interdomain peering between various ISPs in Africa. We measure the presence of local ISPs at various African IXPs, and the extent to which ISPs (or groups of ISPs) choose to interconnect at these exchanges. We find that 66.8% of the paths between residential access links and Google cache nodes in Africa leave the continent. Many circuitous paths result from the fact that local ISPs are often not present at local exchanges, and, when they are, they often do not peer with one another.
- *What can be done to reduce latency to Internet services in Africa? (Section 4)* We explore how both additional peering relationships and proxy cache server deployments could improve the performance of specific services for users in Africa by short-circuiting circuitous paths.

To facilitate follow-on work in this area, we have released the measurement and analysis code for the results in this study [6].

2 Measurements

To measure latency between endpoints and infer peering relationships at IXPs, we rely on 2 datasets: (1) BGP routing tables from RouteViews, PCH, and Hurricane Electric; (2) periodic traceroutes from BISmark routers deployed across South Africa to globally deployed Measurement Lab servers, IXP participants, and Google cache server deployments across Africa.

Caveats and limitations. Our data has several caveats. First, many of the Internet paths that we measure are either to or from locations in one country, South Africa. Thus, our measurements may not reflect the nature of paths in other African countries. We are expanding the deployment across other countries in Africa, and we hope that this study will encourage others to study similar phenomena. Second, peering in Africa is rapidly evolving, and the characteristics we observe from our current measurements will certainly change. Our study provides a snapshot of the current state of peering in parts of Africa and an evaluation of the benefits of improved peering.

2.1 Interdomain Routes: BGP Routing Tables

We also use several sources of BGP routing tables: RouteViews, PCH, and Hurricane Electric. We used the BGP AS path attribute in the routing tables to infer peering relationships between ASes at each IXP. Each of these data sets provides a complementary view into the connectivity between ASes in Africa. In the case of RouteViews and PCH, an AS will peer with a route server at each of these collection points, providing routes to all of its customer ASes, but not to its peers. Most of the ASes at each IXP do not provide routes to RouteViews or PCH, making it difficult to determine the complete set of peering relationships at an IXP. To gain a more complete picture of peering relationships at each IXP, we crawled the Hurricane Electric web portal. This portal allows us to see many additional inter-AS relationships that are not visible in other data sets. If two ASes are (1) adjacent in any AS path that we observe and (2) both present at an IXP, we assume that a peering relationship exists at that IXP. Unfortunately, none of these datasets allow us to see peering links between customer ASes in these BGP feeds, so our view of peering is still limited. To augment these measurements, we use traceroute measurements from BISmark nodes, as described below.

2.2 Router-Level Paths: BISmark Routers

BISmark (Broadband Internet Service Benchmark) is the combination of OpenWrt-based custom firmware and user-space packages [9]; we deploy the software on home gateway routers immediately downstream of the residential broadband access link. BISmark runs on any OpenWrt-capable device, but we have primarily deployed the software on the NetGear WNDR 3700 and 3800. Because the router is always on and connected directly to the provider, we can perform continuous measurements. The BISmark router deployment in South Africa provides the primary vantage points from access networks for this study. We deployed 17 BISmark routers across 7 ISPs and all 9 provinces in South Africa (Figure 1). We performed regular measurements from these BISmark routers to nine global Measurement Lab servers (Figure 1), including three locations in Africa: Tunis, Johannesburg, and Nairobi. We perform traceroutes between the BISmark nodes and the Measurement Lab servers in both directions every thirty minutes using Paris traceroute. These traceroutes expose sequence of ASes that these paths traverse. We use the latency information in each traceroute hop for clues as to when an Internet path may have left Africa (given the scale of these latency values, the coarse-grained latency measurements are sufficient to make such inferences) and the latency in the last hop of the traceroute to estimate the latency to the corresponding M-Lab server.

We use the BISmark vantage points to send traceroute probes to all the IXP participants to infer the peering relationships at these IXPs. If a router at the IXP appears in the path, we conclude that the peering link is present at the IXP. We discarded measurements if the traceroute had missing hop information at a transition between two ASes, since such missing data would prevent us from determining if an additional AS was present between the two that we observed. These measurements revealed 40 additional links at JINX and 14 additional links at KIXP. In cases where an IXP participant hosts a BISmark router, this additional visibility is significant. For example, in the case of

ISP	ASN	Routers
MTN Business Solutions	16637	6
Telkom-Internet	37457	3
Internet Solutions	3741	3
Cybersmart	36874	2
SAIX-NET	5713	1
DPBOL	11845	1
MWEB	10474	1

Table 1: *The ISPs that host BISmark routers in South Africa.*

Fig. 1: *The South African BISmark nodes continuously measure latency to nine global Measurement Lab servers.*

AS36874 (Cybersmart), which hosts two BISmark routers, we can verify 15 out of 20 visible peering links using these additional measurements.

In 2012–2013, Google dramatically expanded its infrastructure for supporting search, adding 1,200 new sites across 850 ASes, which more than doubled the number of countries in which it has a presence [3]. Using data from previous work [3], we identify the 31 Google cache sites in South Africa and Kenya as of early September 2013. From each BISmark router in South Africa, we issue periodic traceroutes to each of these 31 sites, as well as sites in London and Amsterdam. In South Africa, we perform traceroutes to the following ASes that also host Google cache servers: TENET, TICSA, IS, MWEB, Google, MTNNS, and Cell C. In Kenya, we perform traceroutes to the following ASes that host Google cache servers: SafariCom, KENET, AccessKenya, JTL, and Wananchi. Three Google sites are hosted in ASes that also contain BISmark routers: IS, SAIX, and MWEB. Two other sites that contain BISmark routers, Cybersmart and DataPro, have SAIX as a provider.

3 A First Look at ISP Interconnectivity in Africa

We first explore the prevalence of high-latency paths, as measured from the BISmark routers in South Africa to the global M-Lab server destinations (Section 3.1); we explore the nature of these paths in both directions. We then explore the *causes* of these circuitous paths (Section 3.2). Based on our findings, the subsequent sections make recommendations for improving performance to Internet services in Africa.

3.1 High-Latency Paths

Figure 3a shows the distribution of median network latencies observed from the BISmark routers in South Africa to various M-Lab servers around the world. We define the *latency penalty* as the ratio of the observed median latencies to the best-case propagation delay between South Africa and that city (determined by speed-of-light propagation). Figure 3b shows the latency penalties observed for each destination. For both figures, the cities are shown in increasing order of the geographic distance from South Africa. Due to the nature of peering relationships, the increase in latency does not correlate with geographic distance. For example, even though South Africa is closer geographically to Porto Allegre, Brazil than it is to London, latencies to Porto Allegre are

Fig. 2: *The 31 Google sites in sub-Saharan Africa.*

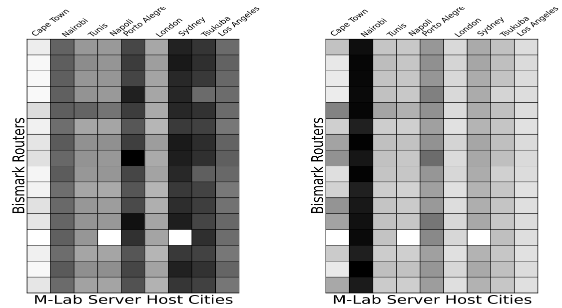

(a) *Absolute latencies.* **(b)** *Latencies normalized by geographic distance.*

Fig. 3: *Latency between BISmark routers and M-Lab servers in various cities. Cities are ordered by increasing distance from South Africa. Darker pixels represent larger values.*

higher, since the path between South Africa and Brazil traverses the London Internet Exchange (LINX). Similar anomalies are evident to other destinations, such as Nairobi, which is geographically close to South Africa but whose paths to South Africa traverse the LINX.

We find that ASes in Africa often do not peer with each other anywhere on the continent. As a result, many Internet paths "detour" through Europe. In the next section, we further explore the extent and causes of these high-latency circuitous paths.

3.2 The Cause of High Latency: Circuitous Paths

We define a *circuitous path* as one that traverses a geographic location that is far from the path created by taking the geographically shortest path between two endpoints. There are two common reasons for circuitous paths: (1) the ASes that provide connectivity along the Internet path between two endpoints are not physically present in a local Internet exchange point (IXP) that is close to the geographic shortest path; or (2) the ASes that provide connectivity are present at a geographically proximal IXP but do not have business relationships with one another or do not prefer that route.

The presence of a local IXP facilitates local peering between multiple ISPs and prevents local traffic from leaving the region. The existence of a local IXP is not enough to guarantee a low-latency path: local ISPs must also choose to connect at the local IXPs. When local ISPs do not connect at a local IXP, the resulting paths can be circuitous. For example, Liquid Telecom (AS 30844) connects at JINX and has a presence in Nairobi [8], but does not peer at KIXP. As a result, users in South Africa must reach must reach many Kenyan networks via LINX in London, significantly increasing the latency of these paths.

Observation: Local IXPs are often not present on local Internet paths. We analyzed the traceroutes between BISmark routers in South Africa and Measurement Lab server locations in Tunisia, Kenya, and South Africa to quantify prevalence of different IXPs

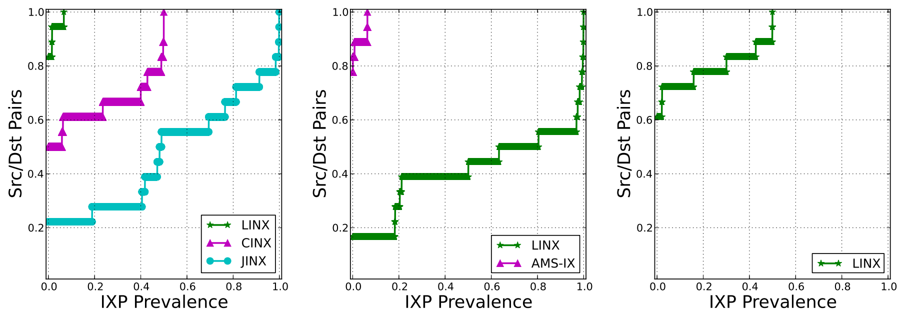

(a) *Johannesburg* **(b)** *Nairobi* **(c)** *Tunis*

Fig. 4: *Distribution of IXP prevalence for the paths from BISmark nodes in South Africa to M-Lab servers in three different cities.*

along Internet paths between end points within Africa. We confirmed that the routers at the IXPs responded to our traceroute probes; we used these responses to identify an IXP's presence on a particular path. We define *IXP prevalence*, which quantifies the pervasiveness of an IXP for various routing paths between the two end hosts. For a pair of end hosts with N observed routing paths, IXP prevalence for the IXP I, P_I, is defined as: $P_I = \sum_{i=1}^{N} x_i P_i, \forall x_i \in \{0, 1\}$, where P_i is the prevalence of the i^{th} routing path and $x_i = 1$ indicates IXP *I* is present for this route.

Local IXPs in South Africa keep local traffic within South Africa. In contrast, local IXPs are much less prevalent along paths between South Africa and other African countries. Figure 4 shows the distribution of IXP prevalence for paths between BISmark routers in South Africa and the three Measurement Lab locations in Africa. Figure 4a shows the IXP prevalence distribution for the Johannesburg M-Lab server; because most BISmark routers are located in South Africa, we observe most of the traffic to Johannesburg traverses IXPs in Johannesburg (JINX) and Cape Town (CINX). Figure 4b, on the other hand, shows a completely different story for paths between the BISmark routers in South Africa and the M-Lab server in Nairobi. The results show a lack of peering at local IXPs and we also did not observe any private peering. Interestingly, KIXP is not at all prevalent for these paths. Figure 4c shows that paths to the M-Lab server in Tunis do not traverse local IXPs in either Tunisia or South Africa.

Cause #1: ISPs do not connect to local IXPs. Sometimes, local ISPs do not connect at the local IXP at all. For example, we observed that Liquid Telecom (AS 30844) connects at JINX and has a fiber presence in Nairobi [8], but for some reason decides not to peer at KIXP, thus causing users in South Africa to take circuitous paths to destinations in Kenya. ISPs in Africa often prefer to interconnect at European exchange points such as LINX because of economy of scale. Most ISPs they need to peer with are present at LINX, not at the local exchanges, so connectivity at LINX is a requirement. Because African ISPs typically all connect at LINX anyhow, connecting at local IXPs simply represents an additional cost with limited additional benefit. Further, the absence of IP

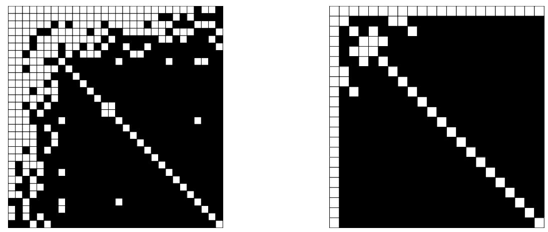

(a) *Johannesburg (JINX)* **(b)** *Nairobi (KIXP)*

Fig. 5: *Peering matrices for the ASes at two African IXPs. White squares represent the presence of peering between an local AS pair, which in many cases we can observe at the IXP itself using traceroute data. Black squares represent pairs of ASes for which we do not observe peering.*

traffic between African countries, such as between South Africa and Kenya, reduces the incentive of ISPs to connect locally. Deploying a cache server in one country to serve the users in another might increase traffic local to the African continent, but such a scenario introduces a catch-22: A service provider such as Google cannot improve performance for South African users by deploying a cache server in Kenya (or vice versa) until the local interconnectivity improves.

Cause #2: ISPs are present at the local IXP, but do not peer. In other cases, ISPs may be present at the same local IXP but may choose not to peer with one another. To study this phenomenon, we analyzed the *peering matrix* of several IXPs, which shows the IXP participants that peer with one another. We constructed peering matrices for the major IXPs in South Africa and Kenya (JINX and KIXP, respectively) using methods from previous work [2, 7]. We used both PeeringDB and the website of each IXP to enumerate the IXP participants. We then analyzed the BGP routing tables as described in Section 2.1 to infer peering relationships at each IXP.

Figure 5 shows the peering matrices for JINX and KIXP. We mapped 51 and 27 ASes for JINX and KIXP, respectively, but the figure includes only the ASes for which we could confirm at least one peering link for these peering matrices (30 ASes at JINX and 22 ASes at KIXP). Figure 5 assumes that if we observe a peering in any path between local ASes that the peering exists at the local IXP, even when we do not always directly observe the peering at the IXP itself. Our data sometimes prevents us from verifying the precise location of the peering link. When we use BGP AS paths, we cannot locate the peering link; we can only observe the existence of peering. In the case of our traceroute measurements, occasionally we see a direct peering without address space from the IXP, but such an observation does not mean a peering at the IXP does not exist. Peering may exist at the IXP but be numbered from one of the peer AS's address ranges, or the path through the IXP may be less preferred than another local private peering. We assumed that a relationship between two local ASes implies a peering relationship at the corresponding local IXP. Note that inferring peering links at an IXP is a hard problem [2] and even after combining multiple data sources, we were not able to infer all the peering links at these IXPs.

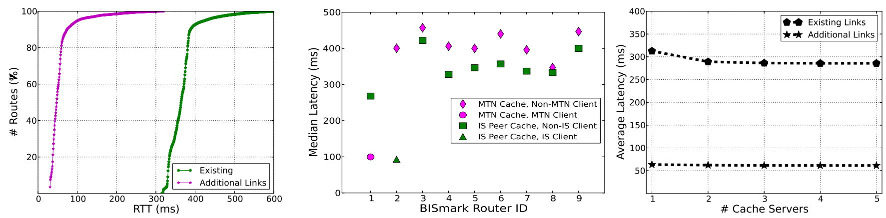

(a) *Adding more Links* **(b)** *Cache servers in local ISPs* **(c)** *Adding more cache nodes*

Fig. 6: *Performance benefits associated with adding additional peering links, placing caching servers in local ISPs, and adding more cache nodes in a region.*

Figure 5 suggests that the peering matrices at each of these IXPs may be sparse. Even when local ISPs are present at an IXP, they do not always peer with one another. When local ISPs do not peer with one another at these exchanges, paths between the local ISPs may be circuitous. Specific examples at KIXP are telling: AS 36914 (KENET) is present at KIXP but we only observed its peering with Ubuntunet and Jamii Telecom. Thus, most paths between KENET and South Africa take a circuitous path through LINX, even though several transit providers at KIXP have direct peering relationships with providers in South Africa (*e.g.*, AS 12556, Internet Solutions, and AS 16637, MTN, are both present at KIXP but do not peer with KENET).

4 Reducing Latencies to Popular Internet Sites and Services

We now evaluate the expected performance improvements that clients in Africa would experience as a result of increased peering at major local ISPs. We also evaluate the relative benefits of adding links versus deploying additional local cache nodes for improving the performance of distributed services using the recent Google cache expansion.

4.1 Add More Peering Links

We quantify the performance benefits of increasing peering at local IXPs to avoid circuitous routes between local ASes. We assume that any circuitous path to Europe could be avoided if the path includes two ASes in Africa that are both present at either JINX or KIXP. In these cases, we replace the delay associated with traversing a path through Europe with the propagation delay between JINX and KIXP, which is about 30 milliseconds. Figure 6a shows the distribution of existing latencies between South Africa and KIXP, and how that distribution would change if these circuitous paths could be avoided. Adding peering links between the ASes that are already present at these local exchanges can significantly improve performance.

4.2 Add More Local Caches

Figure 6b shows the median latencies (from measurements issued every ten minutes over three days) from BISmark routers in South Africa to two Google cache nodes, one in Kenya hosted by a peer of Internet Solutions, and one in Uganda hosted by MTN. Routers in an AS that hosts a cache node or in an AS that peers with an AS that hosts a cache node see low latency; on the other hand routers that are geographically nearby but not in the AS or one of its peers see significantly higher latencies (typically, more than 300 ms round-trip times). We expect that clients in a customer of an AS hosting a cache node would also see low latencies, but we lack such a vantage point.

This result demonstrates that Internet services such as Google can achieve significantly better performance by placing caches to serve local users in the caches' customer cones (or, in some cases, in their peers), even if the clients are in a different country from the caches. Performance from BISmark nodes that are geographically nearby but lack direct paths typically leave the continent and must traverse exchange points in Europe (*e.g.*, LINX). Even when direct paths do exist, the performance benefits may depend on cache placement, since caches typically serve only over customer links and not to providers and peers. Thus, in the absence of adequate interconnectivity, adding cache servers may not improve latency performance for local users who are outside the customer cone of any Google cache. If, on the other hand, a service provider adds caches *and* peers with local ISPs, latencies for local users can improve significantly, even if the service provider places only a single cache in a local ISP. Figure 6c shows the effects of adding additional Google cache nodes in Kenya (which we simulated by taking the minimum latency between a client among k Google cache nodes in Kenya), with and without additional local peering links. This result suggests that content providers should encourage local ISPs to connect at local exchanges, which might ultimately reduce the number of cache server deployments required to achieve a particular level of service.

5 Related Work

A recent study on "boomerang routing" [5] observed that many paths between ISPs in Canada take indirect paths through the United States. We observe similar phenomena for Internet paths that are located in Africa, with the exception that the boomerang is to Europe, as opposed to the United States (and the concern is performance, as opposed to security). Other recent work has studied the internal anatomy and interconnectivity of IXPs [1, 2] but do not share our focus on performance or connectivity in Africa. Other work has highlighted the importance of Internet exchange points for the development of Internet connectivity [4]. Policy work has highlighted the importance of self-organization to improve the efficiency of peering at IXPs in developing-world contexts [10], a behavior that we believe will become increasingly important as peering and interconnection increases in Africa in the coming years.

6 Conclusion

We have taken a first look at Internet paths between locations in Africa, focusing on paths between South Africa, Kenya, and Tunisia. Although this initial study does not represent connectivity across an entire continent, it highlights specific phenomena that deserve attention and further study. First, a significant fraction of local Internet paths in

Africa detour through Europe, resulting in latency penalties of several hundred milliseconds. For example, 66.8% of paths between BISmark routers and Google cache servers in Africa leave the continent. (Latency penalties to other global regions such as South America are also high.) Second, we find that local ISPs are often either (1) not present at the local exchanges; or (2) do not peer with one another at the local exchanges.

ISPs may or may not connect at specific IXPs or peer with one another at a given local IXP for many reasons. These reasons may be economic and political as much as technical, and this issue deserves further study. In contrast to ISPs in developed regions, ISPs in Africa must attain "backhaul" connectivity to large, distant IXPs in Europe, where they can achieve economies of scale with connectivity to other ISPs. Once an ISP connects to Internet destinations via Europe, it has less incentive to connect to local IXPs, which impose additional cost but no significant gains, particularly for ISPs where much traffic is remote. Some disincentives for local peering may relate to the absence of large volumes of traffic between local ISPs, yet we expect that the continued expansion of cache nodes into these regions (*e.g.*, from Google) may change this dynamic. In turn, the deployment of any single cache node may garner much more significant performance benefits in the presence of richer local peering arrangements. As more cache nodes are deployed and more traffic *could* remain local, the peering ecosystem may rapidly evolve to include more local peering links.

Acknowledgments. This work was partially supported by NSF Awards CNS-1059350 and CNS-1162088, and a Google Focused Research Award.

References

1. B. Ager, N. Chatzis, A. Feldmann, N. Sarrar, S. Uhlig, and W. Willinger. Anatomy of a Large European IXP. In *Proc. ACM SIGCOMM*, 2012.
2. B. Augustin, B. Krishnamurthy, and W. Willinger. IXPs: Mapped? In *Proceedings of the 9th ACM SIGCOMM Internet Measurement Conference*, IMC '09, pages 336–349, New York, NY, USA, 2009. ACM.
3. M. Calder, X. Fan, Z. Hu, E. Katz-Bassett, J. Heidemann, and R. Govindan. Mapping the Expansion of Google's Serving Infrastructure. In *Proceedings of the ACM Internet Measurement Conference (IMC '13)*, October 2013.
4. N. Chatzis, G. Smaragdakis, and A. Feldmann. On the importance of Internet eXchange Points for Today's Internet Ecosystem. *CoRR*, abs/1307.5264, 2013.
5. A. Clement and J. Obar. Internet Boomerang Routing: Surveillance, Privacy and Network Sovereignty in a North American Context. 2013.
6. Github: Peering-Africa. `https://github.com/agupta13/Peering-Africa`.
7. Y. He, G. Siganos, M. Faloutsos, and S. Krishnamurthy. Lord of the links: a framework for discovering missing links in the internet topology. *IEEE/ACM Trans. Netw.*, 17(2):391–404, April 2009.
8. Liquid Telecom Fiber. `http://www.liquidtelecom.com/fibre/fibre-map`.
9. S. Sundaresan, W. de Donato, N. Feamster, R. Teixeira, S. Crawford, and A. Pescapè. Measuring home broadband performance. *Commun. ACM*, 55(11):100–109, November 2012.
10. D. Weller and B. Woodcock. Internet traffic exchange: Market developments and policy challenges. Technical report, OECD Publishing, 2012.