

Reshaping the African Internet: From Scattered Islands to a Connected Continent

R. Fanou^{a,b,*}, F. Valera^b, P. Francois^c, A. Dhamdhere^d

^a*IMDEA Networks Institute, Madrid, Spain*

^b*Universidad Carlos III de Madrid (UC3M), Madrid, Spain*

^c*Cisco Systems, Inc, Trento, Italy*

^d*CAIDA/UCSD, San Diego, USA*

Abstract

There is an increasing awareness amongst developing regions on the importance of localizing Internet traffic in the quest for fast, affordable, and available Internet access. In this paper, we focus on Africa, where 37 IXPs are currently interconnecting local ISPs, but mostly at the country level. An option to enrich connectivity on the continent and incentivize content providers to establish presence in the region is to interconnect ISPs present at isolated IXPs by creating a distributed IXP layout spanning the continent. The goal of this paper is to investigate whether such IXP interconnection would be possible, and if successful, to estimate the best-case benefits that could be realized in terms of traffic localization and performance. Our hope is that quantitatively demonstrating the benefits will provide incentives for ISPs to intensify their peering relationships in the region. However, it is challenging to estimate this best-case scenario, due to numerous economic, political, and geographical factors influencing the region. Towards this end, we begin with a thorough analysis of the environment in Africa. We then investigate a naive approach to IXP interconnection, which shows that a theoretically optimal solution would be infeasible in practice due to the prevailing socio-economic conditions in the region. We therefore provide an innovative, realistic four-step interconnection scheme to achieve the distributed IXP layout that considers and parameterizes external socio-economic factors using publicly available datasets. We demonstrate that our constrained solution doubles the percentage of continental intra-African paths, reduces their lengths, and drastically decreases the median of their RTTs as well as RTTs to ASes hosting the top 10 global and top 10 regional Alexa websites. Our approach highlights how, given real-world constraints, a solution requires careful considerations in order to be practically realizable.

Keywords: African internet; Distributed IXP infrastructure; Interconnection scheme; Peering; Content providers

1. Introduction

The African continent, with a total of 1.2 billion inhabitants in its 54 countries, represents the next frontier in terms of end-users that are not connected to the Internet [42, 50, 62] — per ITU stats, only 23% of its population has access to the Internet as of June 2016 [48, 49]. The African Internet ecosystem is experiencing classic “growing pains”: A few Internet Service Providers (ISPs) currently operate in each country, and in many countries the ISP market is dominated by one or two large players. There are 37 local Internet eXchange Points (IXPs) as of March 2016 [32, 97]. However, only 29 of the 58 countries in the region (including nearby islands such as Sao Tome Principe, Mayotte, etc.) have at least one IXP, and the average number of IXP members is 16. While

local IXPs are being set up at a fast rate ¹ and prior studies have demonstrated the benefits that new IXPs can bring [29, 30], some local ISPs are hesitant to peer at those IXPs [43]. Adding to the difficulties, terrestrial fiber deployment remains fragmented [66, 91], since fewer technical and political hurdles make submarine fiber cheaper to build than inland fiber [12, 95].

A major reason behind the stunted growth of the African Internet ecosystem is that the region suffers from a lack of local content [31, 52, 63]. Content is mostly served from the United States (US) and Europe (EU), and even the most popular regional websites are hosted abroad, as investigated in [31]. Consequently, most local ISPs still doubt the value of peering at local IXPs. Those that peer locally are interconnected, but mostly at the country level. In developing regions, it is essential to not only localize traffic but also analyze existing infrastructures for opportunities to improve Internet services at an affordable cost

*Corresponding author

Email addresses: roderick.fanou@imdea.org (R. Fanou), fvalera@it.uc3m.es (F. Valera), pifranco@cisco.com (P. Francois), amogh@caida.org (A. Dhamdhere)

¹18 new IXPs were established in Africa from July 2014 to July 2015 [2, 72]

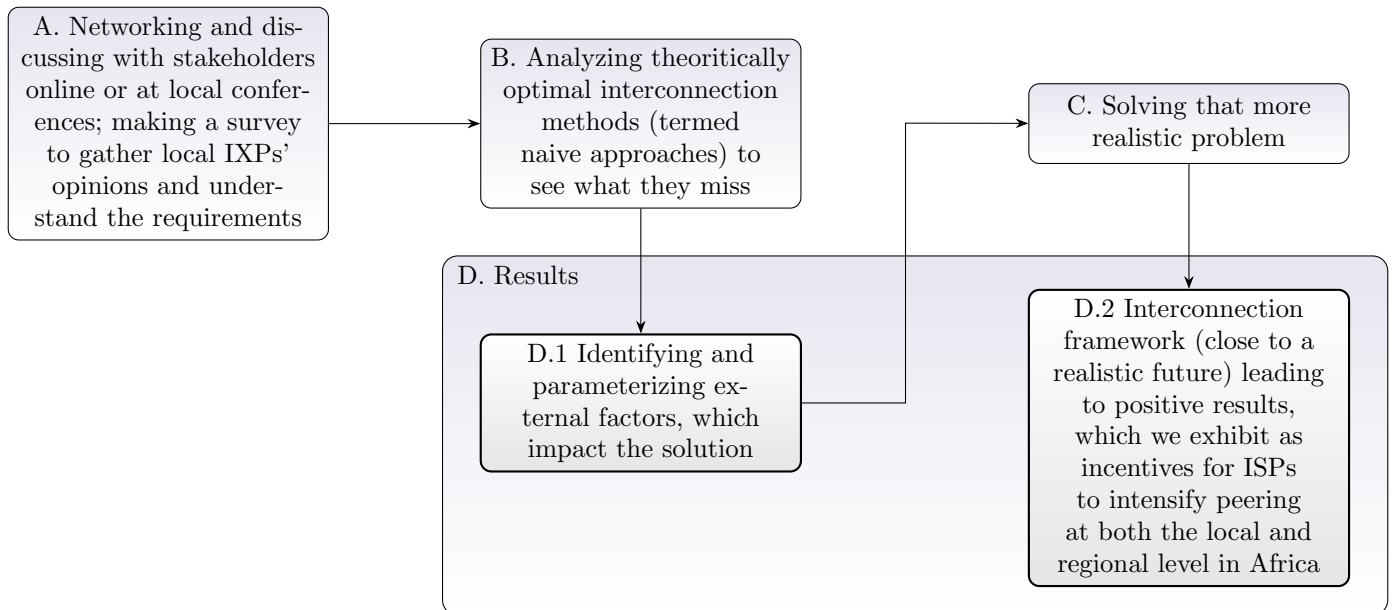


Figure 1: Block diagram of the methodology followed in this work.

[21, 42, 63]. An option that could be considered to enrich connectivity on the African continent and incentivize content providers to establish a presence in the region is to interconnect ISPs present at isolated IXPs by creating a distributed IXP layout spanning the continent. We are not the first to think about IXP interconnection as a way to achieve these goals [26, 27, 68, 70, 92, 94]. However, what is lacking is a concrete proposal for achieving IXP interconnection and a quantitative estimation of potential benefits from doing so. Our main goals in this work are to estimate the outcomes of this interconnection in the best possible scenario that can be realized. However, finding the best interconnection scheme is not straightforward, as this must be done considering all the economic, political, and geographic factors influencing the region.

Figure 1 shows an overview of the methodology we adopted to create the distributed IXP layout and quantitatively estimate its benefits. Similar to our previous studies [29, 30, 31], we have been working closely with local IXP operators and networks in Africa. First, we thoroughly analyzed the situation by means of extensive discussions with stakeholders and inspection of public datasets on the environment in Africa. Particularly for this paper, we conducted a survey of the 37 African IXP operators to get their opinions on the feasibility of IXP interconnections that we reported in §2.

We then explored two naive approaches to solving this problem as a reference point for the rest of this paper. They consist of interconnecting existing IXPs along the shortest (and possibly cheapest) paths, thereby creating a distributed IXP infrastructure spanning the continent. However, the analysis of these solutions revealed that they cannot be implemented due to external reasons such as political instability (including terrorist attacks, wars, riots, rebellions, etc.), lack of fiber, or investments in telecom in-

frastructure (§3). We, therefore, developed and evaluated a framework, which considers and parameterizes all these external factors using publicly available datasets (§5). Further, we used this framework to devise a constrained solution to IXP interconnection that aims to solve both the issue of poor traffic localization and the issue of poor access to popular content.

Our approach to building the distributed IXP structure consists of identifying *secure* local IXPs, selecting regional IXP hubs, connecting those local IXPs and regional hubs in a secure and economical manner, and finally, proposing strategic points where content providers could deploy caches (§4), as shown in Figure 2. Our approach is novel in the following respects: *(i)* we make design choices that ensure that the solution is realizable right away (§4), *(ii)* we incorporate constraints that ensure that the solution is realizable under the present-day geographical, political, and socio-economic realities of the African region (§5), *(iii)* we focus on a solution that requires the minimum possible investment in infrastructure (§6), and *(iv)* we suggest three options applicable within/across sub-regions, given the interests of the stakeholders, to realize the interconnection scheme (§8.2.2).

We use extensive simulations with the open source BGP routing solver C-BGP [79, 80, 81, 88, 64] to evaluate the proposed solution and to quantitatively demonstrate the benefits that would be realized at each step (§6). Specifically, we show that the fraction of continental intra-African paths would double from 40% to 92%, the mode of their lengths would decrease from 4 to 2, median RTTs on such paths would be roughly cut in half, and RTTs to the ASes of the top 10 global and top 10 regional Alexa² websites would decrease of more than their third. We hope

² The platform www.alexa.com [7] is well-known for ranking exist-

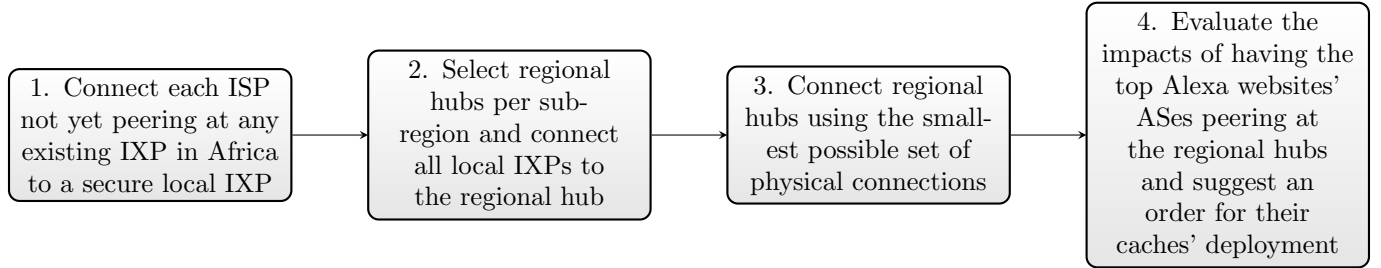


Figure 2: Overview of how we solved the problem (at step C of Figure 1): 4 steps of the proposed interconnection approach to build the distributed IXP layout.

such results will encourage local operators to increase peering and content providers to establish a presence in the region.

Our work makes the following scientific contributions: First, we show how to account for socio-economic realities as constraints in the topology optimization process and how to parameterize them using publicly available data. Note, obtaining data from African institutions or stakeholders on such key issues is difficult, since these are often not collected locally or categorized as classified information. Second, we present and evaluate a framework to build the distributed IXP infrastructure, ensuring that each step respects the practical constraints we have added. For instance, we characterize country stability to guide fiber deployment and justify it with a sensitivity analysis. A direct consequence of the implementation of this framework is that traffic between African countries, rather than traversing another continent, would be routed within Africa following a previously identified country path, through an hierarchically organized IXP substrate. Further, we demonstrate the quantitative benefits of the framework in terms of shorter AS paths, smaller RTTs, and traffic localization that could be realized from each step of the process, using data obtained from our previous measurements and extensive simulations in C-BGP. As an incentive for operators hesitating to invest in the region, we show with measurement data, simulations, and analysis that IXP interconnection has the potential to increase peering density and provide better QoS for intra-African paths and for paths going from African ASes to those hosting top global and regional content.

The remainder of this paper is structured as follows. In §2, we perform a broad analysis of the region that consists of related work and the results of our survey of local IXPs operators. In §3, we inspect naive approaches to the distributed IXP problem and briefly expose the reasons why they would not be feasible in practice. Next, we present in §4 an overview of our solution, a first attempt to interconnect existing IXPs in Africa. In §5, we present an overview of the data collection, the curation methodology, and the parameterization of the model. We then flesh out, in §6, each step of our approach and evaluate the benefits

as compared to the initial AS topology. We explore the sensitivity of our framework to variations in parameterization in §7, before discussing in §8 the limitations of our approach and its feasibility from a technical and political perspective. Finally, we conclude and present directions for future work in §9.

2. Broad analysis of the region

2.1. Background of the region

The 54 African countries can be classified into distinct *sub-regions* (North, West, East, Central, or Southern) as per the African Union [1, 3, 103]. Countries in the same sub-regions often share history, culture (e.g., Southern, North), official language (e.g., West, Central Africa), or currency (e.g., Central and West Africa). The concept of African sub-regions is important while planning infrastructure in the region. Since countries within a sub-region already agree and cooperate³ on various issues, this cooperation could be leveraged.

2.2. Related work on measuring performance on communications among local networks

Prior work has highlighted poor traffic localization and poor access to content from Africa. Gupta *et al.* [43] reported that 66.8% of paths from their vantage points toward Google caches, both in Africa, leave the continent, and often detour through Europe. Further studies underlined the reliance on ISPs based outside Africa for serving intra-continental traffic [29, 30], and found much web content being served from the US and Europe [31]. Despite an increase in the number of local IXPs and recognition of the positive impact of new IXPs on AS path lengths and delays [29, 30], the level of peering among local ISPs remains low due to the lack of participation at IXPs [51, 71]. To the authors' knowledge, there has been no previous research on the African Internet ecosystem such as the longitudinal study we released in [30], which suggests that local stakeholders intensify peering in the region.

ing websites worldwide and by region (as content providers can offer different services from a region to another).

³The sub-regional cooperation is bound under the Regional Economic Communities (RECs) to which the sub-regions belong i.e Economic Community of West African States (ECOWAS), East African Community (EAC), Southern African Development Community (SADC), etc. [3]

2.3. Related work on IXP interconnection

An IXP is a shared layer-2 switch fabric environment, with three or more participants, where new participation is not rigorously constrained, and over which the members peer with each other, exchanging customer routes [19, 28]. Its main business model consists of operating and managing a physical infrastructure to support public and private Internet interconnections [6]. Striking examples are those of NetNod, AMS-IX, and LINX, the managed non-profit IXPs, whose explicit mission is to work for “the good of the Internet” and whose worldwide success in the global IXP marketplace (as opposed to for-profit IXPs) is a result of their governance structure [19].

Intensifying peering in the African region [30] could be achieved by enabling ISPs present at any two isolated local IXPs to peer. A possible way to achieve this is to establish a link between the IXP infrastructures. We are not the first to propose IXP interconnection as a possible solution to the issues encountered by the African Internet. Indeed, the International Development Research Centre (IDRC) and the International Telecommunication Union (ITU) [47] showed in 2005 that establishing national and regional IXPs in the region would lead to monetary and bandwidth savings. They also stressed the need for an appropriate model of IXP interconnection. In 2006, while the number of national IXPs in Africa was standing at 14, Stucke emphasized the need for regional interconnection and listed the necessary pre-conditions for a regional IXP in [94]. In the same year, Pehrson *et al.* proposed [74] a fiber deployment scheme to meet the dual needs of supporting both a research network and IXPs interconnection. Ten years later, however, the required terrestrial fiber has not been established due to a host of economic and political reasons [66, 91, 95]. East African IXP operators proposed [27] to set up the East African Internet Exchange based on a full-mesh interconnection of all IXPs located in their sub-region. Among their guidelines were equal promotion of all IXPs and the absence of competition between IXPs and their members. In contrast, other community developments [26, 68, 92] prefer regional carriers to facilitate cross-border interconnection and provide transit between the various IXPs. But they did not define how to realize it during their meetings.

Interconnecting IXPs is a contentious issue since there are as many arguments for it as there are against it [19, 33, 67]. There are clear reasons why interconnection of IXPs has not gained traction in some cases where it has been implemented: for instance, between LyonIX and FranceIX, each member is limited to 100 Mbps on the interconnection link [33, 35]. Nipper [67] argues that an IXP should not go beyond its diameter since carriers (who are customers of IXPs) would lose revenue on local backhaul. He also advised an IXP operator who runs several IXPs not to interconnect them. However, Nipper also acknowledged that interconnecting smaller IXPs can contribute to gain more critical mass or better gravity. Fenioux [33] argues that IXP interconnection has the advantage of increasing

the attractiveness of an IXP as it facilitates connection of new members from each IXP. Indeed, IXP interconnection has, in the meantime, been achieved in some regions. In France for example, RezoPole operates 2 IXPs (LyonIX and GrenoblIX) that are interconnected. Moreover, these IXPs are interconnected to other IXPs in France or abroad such as FranceIX, Equinix, NetIX, SFINX, fr-IX (all in Paris, France), EuroGIX (in Strasbourg, France), TouIX (Toulouse, France), CIXP (Geneva, Switzerland), and Top-IX (Turin, Italy) [13, 19, 33, 41, 57, 67]. Other examples are those of FranceIX, which deployed interconnections with not only the above listed IXPs, but also LU-CIX in Luxemburg, enabling its members to peer with theirs [35], as well as InterLAN (Bucharest, Roumania) and BalkanIX (Sofia, Bulgaria) [19, 67].

2.4. Survey of African IXP operators

To understand the viewpoint of African IXP operators about IXP interconnection, we conducted a survey of the 37 local IXP operators, receiving 22 responses. Six respondents (27%) are against the idea of interconnecting IXPs. They are prevented by their current policy regime, or do not believe that it will have positive impact.

12 of the 22 responding IXPs (55%) are in favor of interconnecting IXPs. As an example, although the operator of an IXP in a nearby island thinks that its IXP would interconnect to others, he specified that “by nature of being located on an island, there are no other IXPs near enough geographically for it to be practical to connect”. Note that we propose a solution to this issue in §6.2.2. For the operator of one East African IXP in this category, the question is about the lack of a coherent interconnection policy regime among the ISPs, the lack of incentives for colocation services as well as the lack of incentives for local content creation and consumption. According to this IXP, a missing key enabler is that ISPs do not believe peering and interconnection will have positive impacts. In developing our proposed framework for IXP interconnection, we quantitatively show the benefits that can be achieved, to raise awareness about the benefits of peering and IXP interconnection. The said IXP operator further described 2 parameters as being essential to foster the development of peering and interconnection in the region: these are (i) the need of a program to interconnect ISPs operating in Africa at a local and regional level; (ii) the need to boost local content creation and consumption. We tackled the first parameter by proposing the 3 first steps of our framework, while to the second one is dedicated its fourth step. A second East African IXP was supportive of interconnection, even though they are aware of the arguments against it from others.

Four of the 22 responding IXPs (18%) are hesitant and unsure of the best way to proceed on IXP interconnection. For instance, one of the IXPs in Central Africa replied that it would be interested in interconnecting to other IXPs to improve its customers interconnection options, but further

specified that such interconnection would result in significant administrative or financial overhead. Such a fear is understandable: most local IXPs are non-profit entities run by volunteers whose equipment is donated by Internet developmental organizations. Two IXPs hosted in a country of the Eastern African sub-region described IXP interconnection as a controversial topic, as carrying bits over long distances is the business of IXPs participants (i.e. carriers), and it can be dangerous for IXPs to compete with them. The IXP operator also added that if there is no market offering for transport between two IXPs, or the price is very unreasonable, interconnecting the two IXPs as a time-limited measure can be useful to bootstrap the demand and competitive supply.

3. Naive approaches

One way to think about the problem of interconnecting IXPs is as a minimum spanning tree problem, which may be tempting to approach using standard graph algorithms.

We first present an approach in which we find the minimum spanning tree connecting all local IXPs. Let $G(V, E)$ be a graph in which each vertex in V corresponds to an IXP and each link in E , an interconnection between two IXPs. The weight of a link in E is defined as the distance between the two cities hosting the IXPs. Since optical fiber is generally deployed along the roadways or railways [25], we use the Google maps Distance Matrix API [39, 40] to compute the distance of the path between two cities along the shortest roadway that stays on the continent. When there is no path, we evaluate the distance as the crow flies between those two cities, by computing the great-circle distance between the GPS coordinates of the center of each city. We then applied the Kruskal algorithm to the resulting graph G to find the minimum spanning tree.

Next, we manually overlay the spanning tree solution produced by the Kruskal algorithm with known fiber maps [60, 66] to determine which physical links can be used to establish the spanning tree. Figure 3 illustrates the solution. It also shows the reasons why an “unconstrained” solution would be infeasible in practice. Vertices in red represent IXPs in “*unsecured countries*”, i.e. countries that experienced political instability (e.g., Ivory Coast, Egypt, Burkina Faso), rebellions (e.g., DR Congo, Nigeria, Burundi), or terrorists attacks (e.g., Sudan, Nigeria) over the last 5 years [1, 18, 34, 98, 100, 101, 102]. 32.4% of the IXPs are in such “*unsecured countries*”. It may be difficult to deploy fiber connecting these IXPs or to fix a fiber cut in those countries. In addition, if an IXP in an *unsecured country* goes offline, the graph could be partitioned, leading to outages such as those that occurred in Congo and Tchad in April 2016 [83, 84], or in Cameroon’s English speaking areas from January to April 2017 [4, 20].

Six links depicted in red cannot be established because one of the involved countries is *unsecured*. Five terrestrial links in orange could be used for interconnection, but

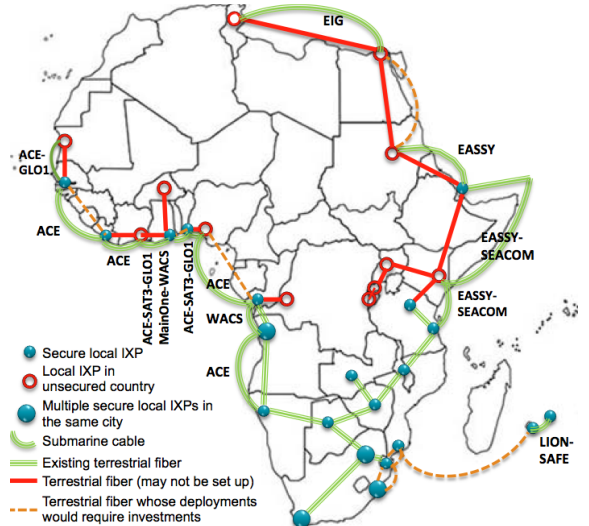


Figure 3: *Interconnecting IXPs in Africa along the minimum spanning tree would be infeasible due to “unsecured” IXPs and the difficulty of fiber deployments along some links; 32.4% of IXPs are in unsecured countries.*

do not currently exist due to various economic and political reasons. For example, DR Congo and Congo do not agree to let any fiber cross their common border; due to regulatory disagreements, optical fiber deployed 5 years ago through the Congo river to interconnect both countries has still not been switched on [85]. Four submarine cables in orange would also need to be deployed – these cables do not exist: none of the submarine cable landing in both countries belongs to the same cable operator. In contrast, green links currently exist and can be used; but these account for only 75% of links.

We also investigate a variant of the above solution where we compute, for each African sub-region, the minimum spanning tree connecting all IXPs within that sub-region. We then link the spanning tree in each sub-region to its 3 closest IXPs in different sub-regions, and manually overlay the interconnection scheme with fiber maps [60, 66]. We find that the result is quite similar to Figure 3 with the main difference being that Central Africa now plays the role of a hub. Still, many of the links within sub-regions cannot be established. In Central Africa, we end up with not only the problematic physical link between DR Congo and Congo, but also a terrestrial fiber between Kinshasa (DR Congo) and Bujumbura (Burundi). Pehrson et al. [74] showed, a decade ago, that the best way to connect the East to the West of Africa is to cross DR Congo with two optical fibers (in the North and the South). However, none of these links have been established until now, mainly because of insecurity in DR Congo and at its border with Rwanda [1, 18, 34].

In summary, we attempted to use standard graph algorithms to find an optimal way to interconnect all IXPs of the region. On inspecting the resulting solutions we find that they are unlikely to be realizable in practice. This analysis motivates the need to create realistic solu-

tions that account for socio-political and economic factors, which influence topology design in the region.

4. Overview of the approach

In this section, we present an overview of our four-step approach to achieve IXP interconnection in Africa. A key ingredient of this approach is that we incorporate geographic, socio-political, and economic realities as constraints in each step of the solution. Further, we discuss the feasibility of its implementation from both a technical and a political perspective in §8.2.

Step-1: Connect each ISP not yet peering at any existing IXP in Africa to a secure local IXP.

To protect their infrastructure investments from damage, destruction, non-usage, etc., it makes sense for ISPs to prefer peering at IXPs in *secure* countries, i.e., countries free of conflicts, terrorist attacks, and political instability. In the first step, we propose to connect each ISP to an IXP in the closest secure and easily accessible country. Note, this does not prevent an ISP from also peering at other IXPs in the world. Historically, there has been a long delay between an IXP setup in the region and wide participation at that IXP; therefore, we focus on connecting ISPs to local IXPs that are already established, rather than setting up new IXPs altogether. For cost-effectiveness, we choose the shortest interconnection paths using existing fiber where possible. We present the details of the interconnection from each ISP to its secure local IXP in §6.2.

Step-2: Select regional hubs per sub-region and connect all local IXPs to the regional hub.

This step leverages the well-known effect that an IXP with many members attracts new members [19, 22, 37]. In step-2, we select one IXP in each sub-region as the regional IXP hub. We then determine the best secure country path from each IXP to its regional hub. When local IXPs are connected to the regional hub, their members are able to peer with all ISPs reachable via the hub. This step incentivizes IXPs in “*unsecured*” countries to participate: (i) those IXPs are included in the framework regardless of the lack of security in their host countries, (ii) step-1 and step-2 are independent and run in parallel (i.e. step-2 proceeds without depending on the outcome of step-1, as explained in §6.3) to help avoiding negative consequences on IXPs located in “*unsecured*” countries.

Step-3: Connect regional hubs using the smallest possible set of physical connections.

In step-3, we connect the regional hubs themselves, using the smallest number of physical interconnections links as possible. By doing so, we ensure that the solution can be realized with minimum investment in additional infrastructure. We present the details of how to interconnect regional hubs in §6.4.

Step-4: Incentivize regional and global content providers to deploy caches at the regional hubs.

The final step consists of creating conditions for end-users in Africa to have access to local and global con-

tent with low latency and the best possible performance. In this work, we use the term *content providers* to refer to (web-based) service providers, which provide content (text, videos, websites, etc.) to end-users. We include traditional content distribution networks (CDNs) such as Akamai that serve third party content, as well as content providers such as Google and Netflix, which build and operate their own extensive networks. Content providers tend to deploy their cache servers within local ISPs, at IXP infrastructures, or within their own network [36] to be as close as possible to end-users. In step-4 of our proposed solution (§6.5), we investigate the benefits that could be achieved if content providers deploy their caches at the previously designated regional hubs, thereby allowing them to reach a large set of connected ISPs. In this step, we then order the regional hubs based on the number of end-users that would be reachable from each of them if they were used as locations for the content provider caches.

5. Data collection

We first discuss how we obtain data to parameterize external factors in our framework. After that, we describe how we build the Internet AS-level topology used for simulating our proposed solution and analyzing the impact on AS path lengths and RTTs.

5.1. Parameterizing geo-political and socio-economical contexts

Matrix of the African continent’s geography:

We define M_{geo} as a $N \times N$ matrix to represent whether two countries are neighbors, where N is the number of African countries (58 including all islands in the region). For instance, if a country A has a neighbor B, the entries A-B and B-A of M_{geo} are set to 1. All the entries of M_{geo} for which one of the countries is an island are set to 0.

Matrix depicting the existence of IXPs: We define M_{ixp} as a $N \times 1$ matrix to quantify the number of IXPs in a country. The value in the row of M_{ixp} corresponding to country c is the ratio of the number of IXPs hosted in c to the total number of IXPs in Africa.

Matrix of submarine cable deployment between African countries: We define M_{sfib} , a $N \times N$ matrix to denote whether one or more submarine cable systems belonging to the same operator land in a pair of countries. For example, the entry corresponding to the countries (Ghana, Ivory Coast) is 5, because 5 submarine cable systems land in both countries: GLO1, MainOne, WACS, SAT3, and ACE. The more common cable systems there are for two countries, the cheaper it is to lease wavelengths on them [95]. A country whose corresponding row in M_{sfib} contains at least one value higher than 0 is either a coastal country or an island. We use this matrix to find the most cost-effective secure country path between two countries in §6.2.2.

Matrix of terrestrial fiber deployment within or between countries: We define M_{tfib} , a $N \times N$ matrix

that captures the presence of terrestrial fiber within or between countries. Specifically, since terrestrial fiber is often deployed along roads [25], we compute for a pair of countries (A, B) the ratio of the length of fiber deployed between the cities hosting IXPs in A and B to the total distance of roadways linking those cities. We obtain these values from [39], following the road along which fiber is deployed [66, 91].

The diagonal elements of M_{tfib} capture the density of fiber deployment within the corresponding countries. To assign values to the diagonal elements, we proceed as follows: the only available datasets of fiber maps per country [66, 91] show that South Africa (ZA) has the highest ratio total length of terrestrial fiber to total distance of roadways. Still, fiber does not fully cover its roadway infrastructure; we estimate the coverage in ZA to be approximately 0.75 (i.e. 75%), the higher bound of the density of fiber deployment in African countries. We then assign to the remaining countries an estimated fraction from among the values 0.125 (denoting a really low fiber density), 0.25 (low fiber density), 0.5 (medium density), 0.75 (high density) depending on their respective deployment efforts [39, 66, 91]. We note that the relative values of these matrix entries are more important than absolute values. Moreover, the accuracy of these numbers may affect our simulation results only when M_{tfib} is involved in the selection of the best country path among two or more secure country paths of the same length (cf. Algorithm 1). We use M_{tfib} to find the most cost-effective secure path between two countries in §6.2.2.

Matrix of African security or political realities:

We define M_{pol} , a $N \times 1$ matrix that identifies countries that have experienced political issues, insecurity (wars, terrorist attacks, riots, rebellions), and disputes with their neighbors [1, 18, 34, 98, 100, 101, 102] during the last 5 years from 2016. The value for the row of M_{pol} corresponding to such countries is 1 and 0 for other countries. We use M_{pol} to identify secure local IXPs and to determine which cross-border fiber deployments are feasible in §6.2.

Matrix of African socio-economic conditions:

Investments in the telecommunications sector, and particularly in fiber deployments, depend on the environment set up by governments, regulators, and stakeholders. To characterize this, we define M_{se} , a $N \times 1$ matrix whose entries are populated with the ratio $R_{se} = I_T / (I_T + I_X + I_E)$, computed per country. In this formula, I_T , I_X , I_E represent the funds invested over the last 5 years by each country in the telecommunications, transport, and energy sectors, respectively [99]. We use the sum of the R_{se} values of countries traversed by a candidate path as a metric in the choice of the best secure country path in §6.2.2. Further, we use R_{se} values in the five-year threshold sensitivity analysis (§7).

5.2. Collecting the Internet AS-level topology

AS relationship dataset: We used the CAIDA AS-level topology snapshot from March 2016 [15], which contains

215,628 AS links and relationships among 53,537 ASes. CAIDA produces this dataset after running the AS-rank algorithm on BGP data from Routeviews and RIPE collectors, combined with traceroutes from Ark monitors toward randomly selected IP addresses in each routed /24 [14].

RTT distribution between ASes: To evaluate the proposed solution in terms of the benefits it can provide w.r.t. performance, we need to estimate the distribution of RTTs on AS links. To this end, we attempt to approximate the RTTs on AS-level links using multiple traceroute datasets. We used the Ark traceroutes data for the first two weeks of March 2016 [16]. This data contains traceroutes performed by 25 Ark probes (deployed worldwide) towards randomly selected IP addresses per (v4/v6) IP range. We also used the dataset collected in [30] composed, among others, of full mesh paris-traceroutes [11] that we performed every week between all or subsets of 238 active RIPE Atlas probes hosted in 136 African ASes in 35 countries from November 2014 to February 2015. To include data depicting access to content, we considered the top 10 global and the top 10 regional Alexa websites [7]: we added paris-traceroutes, previously collected in [31], performed during February - May 2015 from all RIPE Atlas probes in Africa to the front-ends of those top regional and global Alexa websites.

To estimate the delay on an AS link A-B, we computed from all traceroutes outputs in which AS A is followed by AS B, the RTT difference between the ingress point of AS A and that of AS B. This process aims at including the RTT to traverse AS A and reach AS B from AS A. While it is not expected to give us precise RTT values, we obtain several RTT samples for each AS link, which allows us to approximate the mean RTT and distribution of RTTs corresponding to that AS link. We termed this dataset the *AS link RTT dataset*.

IXP Colocation data: We gathered African IXP colocation information (IXP member lists, peering and management prefixes, as well as member ASNs) from PeeringDB [56, 73], PCH [71], TeleGeography Internet Exchange Map [77], and African IXP websites. We then asked local IXP operators to validate (§2) this dataset (from January to March 2016) for completeness, before using it in §6.

5.3. Geolocating ASes by country, by continent and African ASes by sub-region

We collected IPv4 address allocation data from delegation files published by the five Regional Internet Registries (RIRs) [5, 9, 10, 54, 86]. For each IPv4 address block, we geolocated the IPs in the block using the Netacuity Edge database [24]. We are well aware of the limitations of existing geolocation databases [45, 75]; however in this study we are interested in country-level accuracy, which the Netacuity database can provide. The output of this process is the number of IP addresses from a given address block that are geolocated to each country. Next, we

obtained the AS advertising each allocated IP block using Team Cymru’s IP-to-ASN mapping service [96] as of March 2016. For each AS, we thus obtained the number of IP addresses advertised by that AS in each country. We assume that an AS *primarily* operates (i.e., runs its business or is mostly present) in the country in which most of its IPs are geolocated. In total, we geolocated 28,333 ASes — 876 ASes operating primarily in Africa, 10,898 in Europe, 9965 in North America, 2281 in Asia, 3351 ASes in South America, and 773 in Australia. We further classified ASes operating in Africa into the 5 sub-regions: 199 ASes in West Africa, 296 in Southern Africa, 66 in Central Africa, 83 in North Africa, and 232 in East Africa. In this paper, we denote ASes that operate predominantly in the region as *African ASes*, while those operating predominantly outside the region are denoted *non-African ASes*.

5.4. Manual work vs. computational work in our data collection efforts

Collecting data that sheds light on the security situation prevalent in African countries, investments made by countries in different sectors, and mapping logical links to submarine cable maps involved some amount of manual effort, due to a lack of consolidated datasets that can be queried to obtain this type of information in an automated manner. We believe that as the documentation and access to existing datasets improves (for example, if those datasets were indexed in a queryable database), some of the required manual effort can be alleviated. Our results in the subsequent sections (§6 and §7) demonstrate, however, that the manual effort we invest here can have a large payoff in terms of the quality of the solution we obtain.

For some data such as IXP colocation, we combined automated collection from public datasets with a survey for completeness. In our survey, we asked African IXPs operators to validate and complete if necessary the inferred list of their IXP members obtained from publicly accessible datasets such as PeeringDB [56, 73] or PCH [71]. All other data collection tasks including collection of AS topology and relationships, IP geolocation, AS path inference from traceroute and inference of RTT distribution between ASes are automated. Our datasets are freely accessible in the technical report [32].

In summary, we first collected the data necessary to picture the African Internet and simulate our proposed approach. We then parameterized geographical, political, and socio-economic contexts, geolocated ASes by country and by continent, and geolocated African ASes by sub-region.

6. Building and evaluating the distributed IXP layout

In this section, we first construct and characterize our view of the current African AS topology. We then build the proposed solution step by step. At each stage, we evaluate

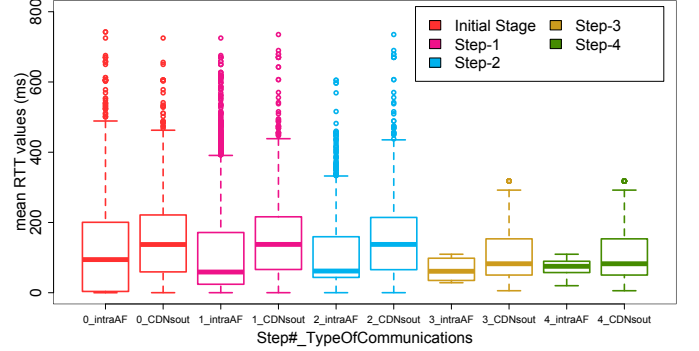


Figure 4: Boxplot of the estimated mean RTT distribution on AS paths at each step, depending on the type of path. The median and interquartile range of RTTs on both intra-African paths and paths towards ASes hosting popular content decrease progressively, as we execute each of the steps.

the resulting topology and quantitatively estimate the impact in terms of the following metrics: (i) fraction of continental paths, (ii) AS path lengths, and (iii) estimated path RTTs. We perform this characterization separately for intra-African paths, outside-African paths, and paths going from African networks to networks hosting top Alexa websites. Table 1 shows an overview of the metrics used to characterize the initial topology and the result of each step. Figure 4 shows the distribution of estimated path RTTs for the initial topology and after each step. We will refer to both Table 1 and Figure 4 throughout the remainder of the paper.

6.1. Building the initial AS topology

6.1.1. Downscaling the collected AS topology

To simulate the effect of interconnecting IXPs and adding peering links, we need a BGP solver, for which we use C-BGP [81]. However, simulating the entire AS-level Internet would be computationally infeasible, as C-BGP quickly becomes memory-bound for large topologies. We implemented the following procedure to scale down the topology to a size suitable for simulation.

We start from every African AS (as defined in §5.3) and traverse customer-provider links until we reach the clique of tier-1 providers [14]. We retain every AS visited in this manner as well as the peers of each visited AS. The retained topology contains ASes that predominantly operate in Africa and other ASes traversed on paths within, from, or towards the region, for a total of 1389 ASes and 10,756 AS links. We then add the prefixes advertised by these ASes to a set \mathcal{P} . Next, we use a list of the top 10 regional and top 10 global Alexa websites as measured by Fanou et al. [31], and obtain the ASes hosting those websites. This gives us 104 ASes hosting popular content, which we add to the subgraph. Note that 74% of those were already present in our retained subgraph. We also add the prefixes originated by these ASes to the set \mathcal{P} . Finally, we need to include prefixes originated by networks outside the previously extracted subgraph. To achieve this, we add to \mathcal{P} all

Table 1: Overview of topology characterization from each step of the process. The column “Initial Stage” reflects the initial topology before any optimizations. The number of continental AS paths, path lengths, and estimated path RTTs all improve progressively as we proceed with the 4 steps. As for the section “Sensitivity Analysis”, the values in the column “Initial Stage” represent the percentage of the 113 initially available secure country paths from any country source to any country destination that need to be recomputed due to the change of the five-year threshold. The values in the remaining columns represent the percentage of the selected best country paths affected by that change. In the last row, we estimated the distance of terrestrial fiber to deploy per step and the corresponding costs. N/A stands for non available.

Type of paths	Metrics	Initial Stage	Step-1	Step-2	Step-3	Step-4
Intra-African AS paths	% of continental AS paths	40%	51.2%	69.5%	94%	91.8%
	% of intercontinental AS paths	60%	48.8%	30.5%	6%	8.2%
	% of AS paths with length ≤ 4	56.9%	69.9%	83.5%	93.05%	93.05%
	% of AS paths with length of 2	1.47%	9.17%	24.8%	74.5%	74.5%
	Mode	4	4	3	2	2
	% of AS paths with mean RTT ≤ 100 ms	37%	59.2%	59.8%	87.5%	95.3%
	% of AS paths with maximum RTT ≤ 1000 ms	20%	47.4%	47%	100%	100%
	Median of mean RTTs (Quartile 2)	144.1ms	58.9 ms	61.75 ms	61.1 ms	75.2 ms
Paths from African ASes to non-African ASes	Interquartile range (Quartile 3 – Quartile 1)	162.1 ms	147.3 ms	115.9 ms	63.2 ms	32.1 ms
	% of AS paths with length ≤ 4	50.77%	53.93%	53.93%	54.38%	54.18%
	% of AS paths with length of 2	0.7%	1.53%	1.53%	2.08%	2.08%
	Mode	4	4	4	4	4
Paths from African ASes to African ASes hosting popular content	% of AS paths with length ≤ 4	61.21%	74.7%	86.3%	91.68%	91.68%
	% of AS paths with length of 2	2.71%	11.1%	31.6%	73.5%	73.5%
	Mode	4	2	2	2	2
Paths from African ASes to non-African ASes hosting popular content	% of AS paths with length ≤ 4	71.06%	70.19%	73.2%	73.79%	74.32%
	% of AS paths with length of 2	2.56%	2.89%	3.8%	4.79%	6.62%
	Mode	4	4	4	4	3-4
	% of AS paths with mean RTT ≤ 100 ms	30.57%	36.44%	37.42%	64.58%	65.67%
	% of AS paths with maximum RTT ≤ 1000 ms	22.82%	22.81%	23.03%	60.83%	87.5%
	Median of mean RTTs (Quartile 2)	137.27 ms	137.55 ms	137.54 ms	82.48 ms	82.49 ms
	Interquartile range (Quartile 3 – Quartile 1)	162.1 ms	150.2 ms	148.72 ms	103.1 ms	103.1ms
Sensitivity Analysis (% best country paths affected by the change of the “insecurity” threshold)	last year	4.42%	3.7%	6.9%	0%	N/A
	last 3 years	1.77%	3.7%	6.9%	0%	N/A
	last 10 years	4.42%	7%	6.9 %	33.3%	N/A
Estimation of minimum and maximum distances (km) for terrestrial fiber deployment in a country/ lower and higher boundaries of total costs (\$) needed at each step	Minimum distance of fiber needed in a country	N/A	173 km	72 km	0 km	0 km
	Maximum distance of fiber needed in a country	N/A	3026 km	72 km	0 km	0 km
	Total distance of fiber to be deployed	N/A	12,024 km	72 km	0 km	0 km
	Lower boundary of total costs needed	N/A	\$73.4 million	\$439,849	\$0	\$0
	Higher boundary of total costs needed	N/A	\$1,80 billion	\$11 million	\$0	\$0

the prefixes originated by the two ASes from each country, which originate the largest number of IPs geolocated to that country. The set \mathcal{P} thus contains 1725 prefixes.

Finally, we obtain the preferred path from each AS in the retained topology to prefixes in \mathcal{P} , by simulating in C-BGP ⁴ the whole AS graph, which consists of 53,537 ASes, 215,628 AS links from the *CAIDA AS relationship dataset* [15], and the set of prefixes from \mathcal{P} . In our C-BGP simulations, we model each AS as a single router, i.e., we do not model the internal topology of ASes. We believe that this is a reasonable simplification for the purposes of this work. We represent an IXP by the set of peers and the peering links found between them according to our data as described in §5.2.

6.1.2. Evaluating the predicted paths

As a sanity check, we then ensure that the C-BGP solver produces reasonable path predictions, by comparing the AS paths produced from the simulation with BGP data available in RouteViews. We first loaded the topology in C-BGP, but only propagated the prefixes of the 876 routers corresponding to ASes geolocated in Africa. We then extracted from the simulated RIBs all 32,486 AS paths starting from AS30844 (Liquid Telecom, one of the largest local networks that are connected to the JINX RouteViews collector) and all 263 starting from AS4558 (known to host the KIXP Routeviews collector).

The BGP data from the JINX and KIXP Routeviews collectors for the first 3 days of March 2016 contained 16,458,193 and 142,599 AS paths, respectively. After comparing both sets, we found 729 common paths for JINX and 48 for KIXP. The fact that we only propagate the prefixes of *African ASes* in this simulation is the reason why the number of simulated paths from JINX and KIXP is small. 82% of the common AS paths have the same predicted length as the actual BGP paths collected from the JINX Routeviews collector. For KIXP, 91% of paths are of the same length. We refer the reader to our technical report for more details [32].


6.1.3. Characterizing the initial topology

We define an *intra-African path* as an AS path, which originates and terminates at *African ASes* ⁵. An *outside-African path* is a path from an *African AS* to a *non-African AS* (cf. §5.3). *Continental paths* refer to AS paths that only traverse *African ASes*, while *intercontinental AS paths* are those, which traverse at least one *non-African AS*.


In the initial topology, intra-African AS paths are composed of 60% intercontinental paths, of which 31% traverse ASes predominantly operating in Europe (EU), 37% traverse ASes operating mostly in North America (NA), while 12% traverse both EU and North American ASes. Figure 5 shows the path length distribution for both intra-African AS paths and outside-African AS paths. We find

⁴C-BGP [79, 80, 81, 88, 64] is an open source routing solver that eases the investigation of changes in the routing or in the topology of large networks.

⁵African ASes are those that predominantly operate in Africa, as defined in §5.3



(a) Distribution of AS path lengths for intra-African paths and for paths between African ASes and non-African ASes.



(b) Distribution of AS path lengths for paths between African ASes and ASes hosting popular content.

Figure 5: In the initial topology, paths length distributions for intra-African paths, paths from African ASes to non-African ASes, as well as paths between African ASes to ASes hosting popular content. The path length distributions are similar in each case, with a mode of 4 AS hops.




Figure 6: In the initial topology, CDF of the mean, minimum and maximum RTT estimates on intra-African AS paths and paths from African ASes to non-African ASes hosting popular content. The CDFs for both types of paths have similar properties.

that the mode of path lengths is 4 in either case. 56.9% of intra-African AS paths have a length of 4 or less. AS paths used to access intercontinental ASes hosting popular content have similar properties. For every AS path, we estimate the mean, minimum, and maximum RTTs on that path by summing the mean, minimum, and maximum RTTs for each AS link on the path, respectively. Figure 6 shows the distribution of minimum, mean, and maximum RTTs on intra-African AS paths and paths between African ASes and intercontinental ASes hosting popular content. We find that the distributions are similar in both cases; for instance, 37% of intra-African AS pairs have a mean RTT of 100 ms or less, while this is 30% for paths to ASes hosting the top regional and global Alexa websites and operating outside Africa (popular content hosted outside Africa).

6.2. Step-1: Connecting each African ISP to its closest secure local IXP

The first step of our solution consists of connecting each ISP not yet peering at any existing IXP in Africa to its closest *secure local IXP*. For this purpose, we need to (i) identify secure local IXPs using M_{pol} and M_{ixp} , (ii) identify, using M_{pol} and M_{geo} , the best path from each country to the closest secure IXP such that the path only traverses

secure countries, and (iii) generate the new AS-level topology (by adding to the initial topology new peering links that can be established at this step) before simulating it in C-BGP.

6.2.1. Identifying secure IXPs and secure relationships between countries

We use the M_{ixp} and M_{pol} matrices (§5.1) to construct the matrix \bar{M}_{ixp} representing secure local IXPs. For any country A, if $M_{pol}[A]$ is 1 (labelled *not secure*), then we set $\bar{M}_{ixp}[A]$ to 0. Table 2 provides details about the 25 secure local IXPs in 18 secure countries covering four African sub-regions: North Africa does not have any secure IXP. We next use M_{geo} and M_{pol} (§5.1) to construct the matrix \bar{M}_{geo} , representing relationships between two secure countries: we discard all inbound relationships towards an unsecured country, but keep outbound relationships from unsecured countries, since ISPs in such countries need to exit them to reach their *closest secure IXPs*.

6.2.2. Identifying the country path from an African AS to its closest secure IXP

After identifying secure IXPs, we need to connect each *African AS* to its closest secure IXP. Suppose an AS A predominantly operates in country s . For this “source” country s , we need to choose a “destination” country d (hosting a secure IXP) for which (i) d is closest to s in terms of country-level hops, (ii) there exists a secure country path from s to d , and (iii) that path would be the most feasible to establish in terms of the real-world constraints specified by M_{sfib} , M_{tfib} , and M_{se} (availability of submarine cable, terrestrial fiber, and telecom investments by countries lying on the path, respectively). As a design principle, we prefer paths via submarine cables over terrestrial fiber: since there are fewer technical and political hurdles to overcome, submarine cables are more established and cheaper in the African region as compared to terrestrial fiber [12, 60, 66, 90, 91, 95].

We start by applying on \bar{M}_{geo} the Breadth-First search (BFS) algorithm to find all possible secure country paths from a “source” country s to a “destination” country d .

Table 2: List of (the 25) secure local IXPs in Africa with the number of members, classified by sub-region and country: the numbers in bold were validated by the corresponding IXPs, as their operators responded to our survey. Numbers in regular font were fetched from the IXP websites but could not be validated. The remaining correspond to IXPs, which neither have a website, nor responded to our survey. We thus collected, where possible, their number of ASNs from public datasets other than the IXP websites, or estimated it to the total number of ASNs operating in the IXP host country (numbers in *italics*).

African sub-region	Country	#IXPs	#Members ASNs
East Africa	Djibouti	1	5
	Mauritius	1	12
	Reunion	1	16
	Tanzania	2	33 - 6
Central Africa	Congo	1	8
Southern Africa	Angola	2	11 - 6
	Botswana	1	12
	Malawi	1	14
	Mozambique	1	11
	Namibia	1	5
	South Africa	6	56 - 37 - 17 141 - 83 - 29
	Swaziland	1	7
	Zambia	1	12
	Zimbabwe	1	8
West Africa	Benin	1	5
	Gambia	1	14
	Ghana	1	17
	Liberia	1	5
Total	18 countries	25 IXPs	

We then use Algorithm 1, which we describe briefly in the subsequent paragraphs, to select the best country path $s - d$ from among the available candidates.

For a “source” country s that is itself secure, the closest secure “destination” country is obviously itself; for all such countries, we trivially obtain the best country path. For s having only a single secure path to d , we retain that path $s - d$ as best country path. These two cases accounted for 25 source countries. For each of the 33 remaining countries, either there is no path, or there are at least two possible secure paths to destination countries. For 19 of the said countries, multiple paths have the same length: we, therefore, need a tie-breaker. Since our rationale for breaking ties is based on the fact that submarine cables are preferred to terrestrial cables, we first try to find the best possible path via submarine cables.

To tie-break among paths of length l , we examine all paths $s - d$ that can be established using only submarine cables. The following parameters are computed for each such path: $A_s = \sum_{c \in C} (M_{sfib}[c]/|C|)$ and $C = M_{ixp}[d]$, where C is the set of countries lying on $s - d$. While A_s is a measure of the total number of common submarine cable operators to any two consecutive countries on the path, C is a measure of the number of IXPs in d at which a network could peer. If there is a country path of length l for which A_s and C are both highest, we label that path $s - d$ as the best country path. Otherwise, we retain the path for which A_s is highest. As an example, we prefer the country path Togo - Ghana (via GLO1 or WACS submarine cables) over the path Togo - Benin (via only GLO1). We also prefer the path DR Congo to Angola (via WACS and ACE and

Algorithm 1: Identification of the best country path from a country to the closest secured IXP

Data: Set P of all possible country paths p from a given country c towards any reachable secure country d , M_{sfib} , M_{ixp} , M_{tfib} , M_{se}

Result: Set Pb of best paths from any country towards its closest secure country

```

1  $Pb = \{\}$  /* Initialization of Pb */
/* Label as best any unique country path */
2 for  $c \in P.keys()$  do
3   if  $len(P[c]) = 1$  then  $Pb[c] = P[c]$ 
/* Identify the best path for the rest */
4  $current\_country\_path\_len = 2$ 
5 while  $current\_country\_path\_len < 58$  do
6   for  $c \in P.keys() \mid c \notin Pb.keys()$  do
7     /* Can we use submarine cables ? */
 $A_s = \{\}$  /* Sum # of common types of
8       submarine cables per path */
 $C = \{\}$  /* Percentage of African IXPs
9       in destination country */
10    for  $p \in P[c]$  do
11       $i = 0$ 
12      while  $i < len(p)-1$  do
13         $A_s[p] += M_{sfib}[p[i], p[i+1]]$ 
14         $i += 1$ 
15       $C[p] += M_{ixp}[p[i]]$ 
16    if  $\exists p \mid A_s[p] = \arg \max A_s(x)$  and  $C[p] = \arg \max C(x)$  then  $Pb[c] = p$ 
17    else if  $\exists p \mid A_s[p] = \arg \max A_s(x)$  then
18       $Pb[c] = p$ 
19    /* What about terrestrial fiber ? */
20     $A_t = \{\}$  /* Ratios of terrestrial
21      cables deployment per path */
22     $B_t = \{\}$  /* Investments in the
23      countries on each path */
24     $C = \{\}$  /* Percentage of African IXPs
25      in destination country */
26    for  $p \in P[c]$  do
27       $i = 0$ 
28      while  $i < len(p)-1$  do
29         $A_t[p] += M_{tfib}[p[i], p[i+1]]$ 
30         $B_t[p] += M_{se}[p[i]]$ 
31         $i += 1$ 
32       $B_t[p] += M_{se}[p[i]]$ 
33       $C[p] += M_{ixp}[p[i]]$ 
34    if  $\exists p \mid A_t[p] = \arg \max A_t(x)$  and  $B_t[p] = \arg \max B_t(x)$  and  $C[p] = \arg \max C(x)$  then  $Pb[c] = p$ 
35    else if  $\exists p \mid A_t[p] = \arg \max A_t(x)$  and
36       $B_t[p] = \arg \max B_t(x)$  then  $Pb[c] = p$ 
37    else if  $\exists p \mid A_t[p] = \arg \max A_t(x)$  and  $C[p] = \arg \max C(x)$  then  $Pb[c] = p$ 
38     $current\_country\_path\_len += 1$ 

```

toward two IXPs) to the path DR Congo to Congo (via only WACS and toward 1 IXP).

If there is no path of length l from s via submarine cables, we then look for a path using terrestrial cables. The following parameters are computed for each secure country path originating from s : $A_t = \sum_{c \in \mathcal{C}} (M_{tfib}[c]/|\mathcal{C}|)$, $B_t =$

$\sum_{c \in \mathcal{C}} R_{se}$, and $C = M_{ixp}[d]$, where \mathcal{C} is the set of countries on the path $s - d$. A_t is a measure of the terrestrial fiber that exists on the path, B_t is a measure of the investment in telecoms for all countries on the path, and C is a measure of the number of IXPs in d at which a network could peer. If to a path of length l correspond the maximum values of A_t , B_t , and C , we select that path as the best country path⁶. These are, for instance, the cases of Rwanda-Tanzania, Uganda-Tanzania (through terrestrial fiber and toward two IXPs). Otherwise, if we find a path with the maximum values for A_t and B_t , we select that path⁷. Otherwise, if to a path correspond the maximum values for A_t and C , we select that path⁸. As an example, the country path Burkina Faso - Ghana is preferred to Burkina Faso - Benin, because A_t is higher for the former and both Benin and Ghana have one IXP.

If we cannot find a path of length l after these steps, we repeat the process starting with submarine cable paths of length $l + 1$. Exploring all country paths of length l before moving to paths of length $l + 1$ aims at preferring paths whose destination countries are close, rather than paths traversing those countries to reach countries far away. As a consequence, ISPs in 66.7% of unsecured countries have their best paths destined to a neighboring country.

After the previous steps, we have assigned a best path to 44 countries out of 58. The remainder corresponds to islands without IXPs (e.g., Comoros, Saint Helena, Cape Verde, etc.) or countries for which all neighbors are labelled unsecured (Libya, Egypt, etc.). For these, we identify the closest secure country hosting an IXP and sharing submarine cables run by the same operator. For instance, ISPs in Comoros need to connect to Mauritius via LION, while those in Egypt connect to Djibouti via EASSY and SEACOM. At the end of this step, all countries are assigned a best path, as depicted in Figure 7.

6.2.3. Connecting ISPs to their closest secure IXPs

We next simulate ASes peering at their closest secure local IXP. We assume that an AS only peers with networks that are not in its customer cone [23, 69], as it has no incentive to peer with networks it can reach via customer links. This consideration is consistent with our goal

⁶preference for country paths with considerable terrestrial fiber deployment, larger investments in telecoms, and more diversity in IXPs at the destination

⁷preference for country paths with considerable terrestrial fiber deployment and characterized by larger investments in telecoms

⁸preference for country paths with terrestrial fiber deployment and more diversity in IXPs at the destination

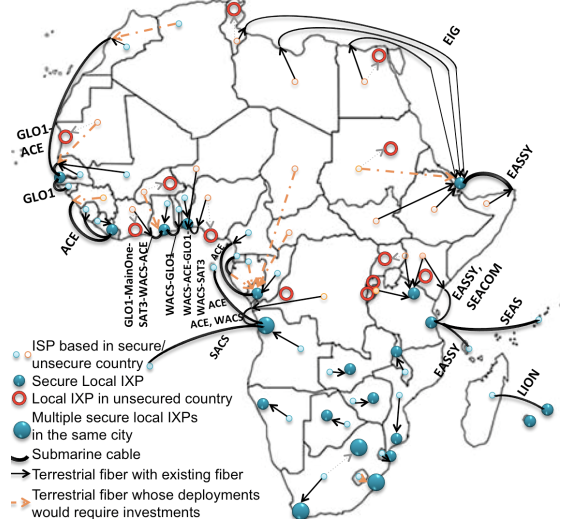


Figure 7: Result of step-1, where each ISP connects to its closest secure IXP.


to evaluate the best possible scenario that can be realized and its impacts on AS path lengths and performance. In §8.2 we discuss the inherent complexities of peering economics, which may cause an ISP to prefer another country path/IXP than the one proposed, or to connect to more than one IXP, or to selectively peer with a subset of ISPs present at an IXP.

We simulate peering at IXPs applying the customer-cone constraint based on the customer cone of each AS from the March 2016 AS relationship data [15]. We add 56,863 peering links to the initial topology at the completion of step-1. The average number of members per IXP doubles from 18 in the initial topology to 37 after step-1. The biggest IXP that emerges is NAPAfrica Johannesburg (JB) with 240 peers.


To estimate the RTTs on newly created interconnection links, we first compute the geographic distance $C_h(s, d)$ between the IXPs at which the interconnecting networks are present. When the interconnection occurs via two or more terrestrial fibers, we sum the distances of those fibers as per [39]. When the interconnection occurs via one terrestrial and one submarine fiber, we sum the length of the terrestrial fiber with the distance as the crow flies between the two cities connected via the submarine cable. Since light travels about 1/3 slower through optical fiber than through a vacuum [76, 82], the $RTT(s, d)$ over the established link can be estimated as: $RTT(s, d) = \frac{2 * C_h(s, d)}{2/3c} = \frac{3 * C_h(s, d)}{c}$, where $C_h(s, d)$ is the distance (km) between the cities following roads/railways [39], and c the speed of light in vacuum. Finally, to connect an AS to a secure local IXP located in the same country, we estimate the RTT on the newly established links as the mean of all RTTs among ASes operating in that country.

6.2.4. Characterizing the resulting topology

To simulate the effect of step-1, we propagate the 1725 prefixes in C-BGP on a topology where all the new peering



(a) Distribution of AS path lengths for intra-African paths and for paths between African ASes and non-African ASes.



(b) Distribution of AS path lengths for paths between African ASes and ASes hosting popular content.

Figure 8: After step-1, paths length distributions for intra-African paths, for paths between African ASes and non-African ASes, as well as for paths between African ASes and ASes hosting popular content. While the mode of the path length distribution is still 4, the fraction of paths with length 4 or fewer is higher than in the initial topology.




Figure 9: After step-1, CDF of the mean, minimum, and maximum RTT estimates on intra-African AS paths and paths from African ASes to non-African ASes hosting popular content. The median of mean RTTs for intra-African paths decreases from 144 ms in the initial topology to 59 ms after step-1.

links have been established as described in §6.2.3. Compared to the initial stage, the percentage of continental intra-African paths increases from 40% to 51.2%. Still, 26% of the intercontinental intra-African paths traverse ASes operating predominantly in EU, 31% traverse ASes in North America, and 9.75% traverse ASes operating predominantly in both regions.

Figure 8(a) shows that the mode of intra-African AS paths lengths is still 4. The percentage of such paths having a length of 4 or fewer increases from 56.9% to 69.9%, however. Similarly, Figure 9 shows that the percentage of intra-African AS pairs with a mean RTT of 100 ms or less has increased from 37% to 59.2%. The metric *median of mean RTTs* refers to the median of the estimated mean RTTs across all paths of a certain type (intra-African or outside-African path). Interestingly, the median of mean RTTs for intra-African paths has declined from 144.1 ms (with an interquartile range of 162.1 ms) in the initial topology to 58.92 ms (with an interquartile range of 147.3 ms) after step-1, as shown in Figures 4 and 9.

Unsurprisingly, we observe no change in the distribution of AS path lengths or RTTs for paths between African ASes and non-African ASes. Indeed, since step-1 increases peering among African networks, which are often leaves of the topology, we did not expect those metrics to improve.

The median of mean RTTs values remains steady: 137.5 ms with an interquartile range of 150.2 ms. The number of AS paths of length 2, from all African ASes to those hosting popular content triples as compared to the initial stage.

6.3. Step-2: selecting regional IXP hubs

In step-2 of our proposed approach, we select a regional hub from among the secure IXPs in each sub-region. Recall that a high-level objective of our optimization is that it should be realizable at the present time. Consequently, we would like step-2 to proceed without depending on the outcome of step-1. We, therefore, select as the regional hub, the secure IXP of a sub-region, which *currently has the most member ASes*. From Table 2, the regional IXP hubs are TIX in Tanzania (33 members) for East Africa, CGIX in Congo (8) for Central Africa, NAPAfrica JB in South Africa (141) for Southern Africa, and GIXA in Ghana (17) for West Africa.

Next, we need to connect each of the 33 remaining local IXPs⁹ to its regional hub. This involves finding the best secure path from the country of the local IXP to that of the regional hub. Towards this end, we only consider the secure paths $s - d$ (computed as in §6.2.2) going from any “source” country s hosting a local IXP, towards the destination country d that hosts the regional hub. Again, we tie-break among paths of the same length according to Algorithm 1, using parameters A_s, A_t, B_s, B_t, C , and preferring submarine cables over inland fiber. This gives us the best path for 26 of the 33 IXPs.

The remaining IXPs can be classified into 3 categories: First, among IXPs in North Africa such as CAIX (Egypt), RIMIX (Mauritania), TUNIXP (Tunisia), and SIXP (Sudan), no secure local IXP was found, hence no regional hub could be selected. We connect such IXPs to those in their best destination country as per step-1 (§6.2). Second, we found KINIX (DR Congo) to have no secure path to its regional hub: again, we use the best country path from

⁹37 African IXPs minus the 4 regional hubs

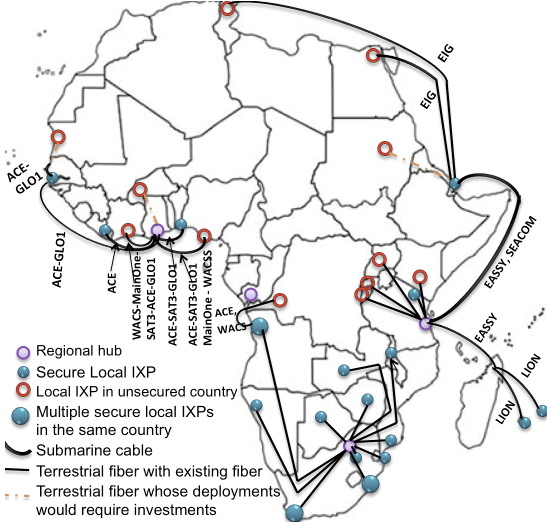


Figure 10: Result of step-2, where each IXP connects to the regional hub selected among the secure IXPs of each region.

step-1. Finally, we connected IXPs located on islands, such as Mauritius-IX (Mauritius) and Renater-IX (Reunion) to their closest regional hub (TIX) via submarine cables. The results are shown in Figure 10.

At the end of step-2, the average number of IXP members in Africa increases from 37 to 50, when compared to step-1. The biggest IXPs are now NAPAfrica JB (382 peers), TIX (334 peers), and GIXA (239 peers), each having at least twice the number of their peers after step-1.

6.3.1. Characterizing the topology after step-2

In step-2, we add 89,709 peering links to the topology out of 110,416 possible links (81%). This makes the percentage of continental intra-African AS paths increase from 51.2% to 69.5%. After this step, the AS path length distribution of such AS paths has a mode of 3. Moreover, 83.5% of the continental intra-African AS paths have a length of 4 or less. The percentage of continental intra-African paths having a length of 2 increases from 9.2% to 24.8%. The median of the mean RTT values is slightly higher (61.75 ms) than that of step-1 with an interquartile range of 115.9 ms, reduced of 31 ms. In addition, AS paths to African ASes hosting popular content see an improvement: the mode of their length is now 2, and 86% of these paths have a length below 5 (see Table 1). AS paths towards non-African ASes, however, still have a mode of 4. Meanwhile, AS paths for accessing any of the non-African ASes hosting popular content, towards which users are often redirected [31], have kept the same distribution as the initial stage.

6.4. Step-3: interconnecting regional IXP hubs

After step-2, we are left with 4 regional IXP hubs: NAPAfrica JB, TIX, GIXA, and CGIX located in South Africa (ZA), Tanzania (TZ), Ghana (GH), and Congo (CG), respectively. Since the next step is to interconnect these hubs, we leverage, once again, Algorithm 1 to find

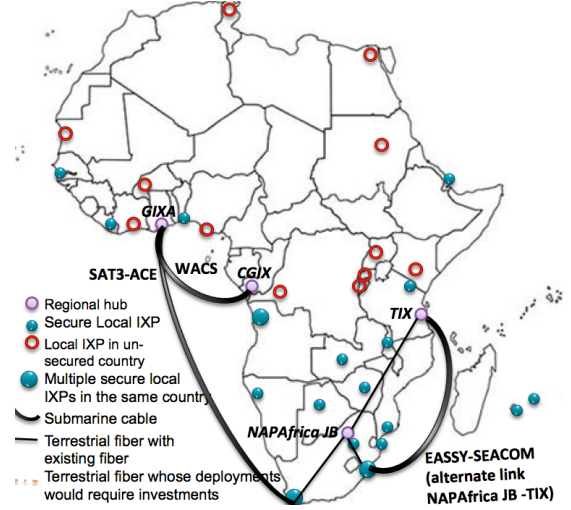



Figure 11: Result of step-3, where regional IXPs are interconnected with a minimum number of links.

the best country path as in §6.2. In this case, however, instead of using the full M_{ixp} matrix as input, we use a sub-matrix of M_{ixp} composed of the rows and columns corresponding to GH, ZA, TZ, and CG. The country-path algorithm gives us the set of physical links that could be used to establish connections between the regional hubs. TZ and ZA appear as the closest secure countries for each other, and the preferred link between them is an existing terrestrial fiber passing through Mozambique, although using a link via EASSY or SEACOM submarine cable is also possible. Meanwhile, the preferred link from ZA to GH is either the submarine cable SAT3 or the submarine cable ACE. No terrestrial fiber is found in this case (ZA – GH) since Nigeria, labeled as unsecured country, cannot be traversed. Finally, we find a link from CG to GH (via WACS). Further, there is no link of any type from CG to ZA or TZ, making any attempt to interconnect all regional hubs with a full-mesh practically impossible: DR Congo, labeled as unsecured, cannot be traversed by the terrestrial fiber and no functional submarine cable lands in both ZA and CG. Given that a full-mesh of links between the regional hubs would be practically impossible, we choose instead to find the smallest set of links that could be used to interconnect all IXPs.


6.4.1. Choosing the smallest set of physical links

To select from among the possible physical links that can be set up to link the regional hubs, we use a greedy approach. At each iteration, we connect the pair of regional hubs, which would result in the largest number of potential new peering links. We repeat this process until all regional hubs are interconnected.

Figure 11 summarizes the results. We find that 3 links are needed: the link between NAPAfrica JB to TIX via Mozambique (with only 72 km of terrestrial fiber to be deployed), the link between NAPAfrica JB and GIXA via SAT3 or ACE, and the link between GIXA and CGIX via



(a) Intra-African communications and communications between African ASes and non-African ASes.



(b) Communications between African ASes and ASes hosting popular content.

Figure 12: After step-3, paths length distributions for intra-African paths, paths from African ASes to non-African ASes, as well as for paths between African ASes to ASes hosting popular content. The mode of intra-African AS path lengths is now 2; 74.5% intra-African paths have a length of 2. AS paths between African ASes and African ASes hosting popular content also have a mode of 2.

WACS. If these links were established, 299,740 (64% of possible) new peering links would be added to the topology. The distributed IXP thus created would have in total 964 unique members.

A natural question that may arise is which entities may have the incentive and the capability to provide links between regional hubs? This is a complex issue that involves not only economics but also the business interests and strategies of various stakeholders. We discuss the issue in depth in §8.2.2.

6.4.2. Characterizing the topology after step-3

After step-3, we find that 94% of the intra-African paths are now continental paths. The remainder traverse ASes that predominantly operate in another continent: 5% traverse ASes predominantly in EU, 1.6% traverse ASes in North America, and 0.6% traverse ASes in both regions. In terms of AS path lengths (Figure 12), we find that this step changes the mode of the intra-African path length distribution to 2. In fact, 74.5% of intra-African paths have a length of 2. Further, AS paths between African ASes and African ASes hosting popular content also have a mode of 2. But the distribution of AS paths from African networks to non-African ASes remains unchanged. Specifically for AS paths going from African ASes to non-African ASes hosting popular content, the mean RTT values have, however, decreased to a median of 82.48 ms with an interquartile range of 103.1 ms, as compared to 137.54 ms with an interquartile range of 103.1 ms for step-2. 64.6% of the AS paths for accessing content hosted in non-African ASes now experience a mean RTT of 100 ms or less (Table 1).

6.5. Step-4: Incentivizing regional and global content providers to deploy caches at the regional IXP hubs

The previous steps produce a hierarchy in the African IXP substrate: ISPs – local IXPs – regional IXPs. To trigger the interests of content providers (as defined in §4) to contribute to its realization, we aim at emphasizing in this section what they might gain from participating in it. A typical content provider controls a hierarchy of servers,

using its back-end servers to efficiently ensure the distribution of content within its infrastructure, and its front-end servers to handle user-server communications. This infrastructure replicates content at multiple locations across the Internet [93]. While content providers can vary in their technical operation (e.g., whether they operate their own backbone network or not), we leverage the fact that all content providers would be interested in establishing a presence (either deploying caches or peering infrastructure) at a few locations that can have the most impact in terms of performance. The regional hubs selected in step-2, which later constitute the core of the distributed IXP framework, serve as natural points where content providers could establish a presence to serve end-users of each African sub-region with their content popular in each of them.

In step-4, we evaluate the outcomes in terms of AS path length, end-to-end delay, and number of end-users whose performance may be improved, if ASes hosting the top global and regional ¹⁰ Alexa websites [7] mapped in [31] were to peer with networks present at the 4 regional hubs. Note here that we simulate a specific mode of operation wherein the content provider network peers with other networks present at the IXP. We find that this peering would create 12,339 (85.33%) new links, out of the possible 14,460 peering links, since some of them already exist. The properties of the resulting topology are similar to those after step-3. Most noticeably, 95.3% of intra-African AS pairs now have a mean RTT of 100 ms or less as compared to 87.5% for step-3. The median of mean RTTs on intra-African paths increases from 61.1 ms to 75.2 ms with a halved interquartile range (32.15 ms), as shown in Table 1. Meanwhile, the median of RTTs from African ISPs to popular content hosted outside Africa stays at 82.5 ms (Figure 4). These similarities are expected, as adding the presence of content providers at strategic locations does not significantly change the properties of the macroscopic topology, but instead influences the performance of paths used to access their content.

¹⁰Content providers can offer different services from one region to another.

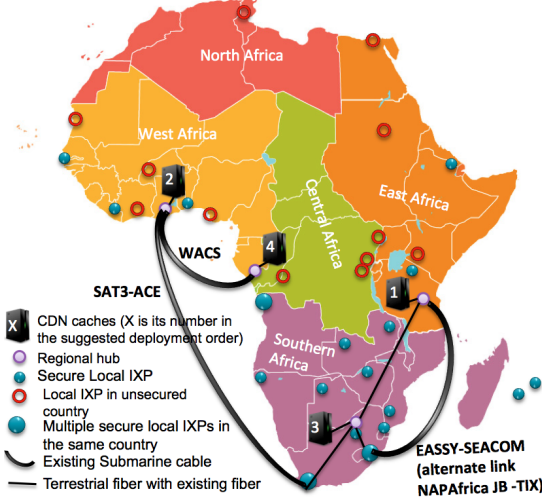


Figure 13: Result of step-4, where we suggest an order of content providers’ caches deployment within the infrastructures of the strategic points represented by regional IXPs.

While establishing content providers’ presence at all regional hubs will have the most impact, the cost of doing so at each regional hub may be prohibitive. We, therefore, suggest an order of deployment by estimating the number of end-users (as a percentage of the Internet population in the region) that are reachable from each regional hub. To determine the size of the user population in Africa, we consider all ASes operating in the region and sum their estimated number of users, as per the APNIC labs measurement project [8]; we obtain an estimated total of 331,428,949 end-users in Africa. We then consider each of the regional hubs from step-2, and compute the total number of end-users reachable from that hub by adding the estimated user base of each AS connected to that hub. With 334 peers after step-2, TIX serves an estimated 132,571,579 end-users corresponding to 40% of the end-user population in Africa. GIXA (239 peers) corresponds to 39%, NAPAfrica JB (382 peers) corresponds to 16%, and CGIX (43 peers) to 3.17%. Interestingly, while NAPAfrica JB has the largest number of peers among the regional hubs, it is third in terms of the number of end-users served. Thus, we suggest that to incrementally establish presence at the regional hubs, content providers should proceed in the order TIX, GIXA, NAPAfrica JB, and finally CGIX to have the largest impact (Figure 13).

7. Sensitivity analysis

An important consideration that drives the construction of the distributed IXP layout proposed in this paper is the notion that a country is labelled “secure” or “unsecured” due to geo-political factors: in §5.1 we chose a period of 5 years without conflicts, riots, rebellions, or security issues to decide whether or not a country is “unsecured”. Given that this parameter can impact the resulting topology, we perform a sensitivity analysis of the

“insecurity” threshold to determine whether a different value of this threshold qualitatively changes our results. It turns out that while the number of unsecured countries is 23 for the last 5 years, it is 20 for the last one year or the last 3 years, and 27 for the last 10 years. Despite this difference, the list of secure local IXPs does not vary. Moreover, the regional hubs (selected at step-2 §6.3) remain identical for any of these thresholds. However, the number of secure country paths initially available at each step, and hence the best country paths selected for the interconnection links may also change if a threshold different from the five-year period were preferred. We evaluate, per step, the percentage of secure paths that would be affected and summarize the results in Table 1. We find that choosing a threshold of one year or three years has a small impact: at most 6.9% of the selected best country paths are different. For the ten-year period, however, this percentage reaches 33.3% for step-3, because an unsecured country is now traversed by one of the three country paths selected to interconnect the regional hubs. Interestingly, whenever the country path previously selected for a five-year period threshold now traverses an “unsecured” country, Algorithm 1 ensures that an alternative path is selected.

We also use a five-year period for computing per country the ratios R_{se} with which we populate the matrix M_{se} (§5.1). To evaluate how varying the threshold would affect our results, we compute M_{se} for a one-year, three-year, and ten-year period. We then quantify the correlation between their respective values R_{se} and the values R_{se} registered for the five-year period. We found the correlation coefficient r to be 0.9721 for $(M_{se}(5years), M_{se}(1year))$, 0.9795 for $(M_{se}(5years), M_{se}(3years))$, and 0.8692 for $(M_{se}(5years), M_{se}(10years))$. Figure 14 shows these correlations, in terms of the strength and direction of the relationship. This analysis shows that selecting a threshold different from the five-year period used in this paper to compute the values R_{se} , will not qualitatively change our results. In other words, choices operated with $M_{se}(5years)$ will not differ significantly from those with $M_{se}(1year)$, $M_{se}(3years)$, or $M_{se}(10years)$.

8. Discussions

8.1. Limitations of the current approach

We discuss in this section the limitations of our work. First, we acknowledge that socio-economic conditions are quite unstable and constantly evolve. While we showed with our sensitivity analysis that our results are robust to changes in these parameters over a few years, we recognize that this analysis needs to be repeated periodically with fresh data in order to accurately reflect real conditions. Second, we recognize that accounting for socio-economic and political factors is complex, and there are many factors beyond the ones we considered in this paper (§5.1) that could affect the realization of the distributed infrastructure we propose. Nonetheless, this paper was a first




Figure 14: Sensitivity analysis: Correlation between Ratios R_{se} of the matrix M_{se} evaluated for different thresholds is found to be 0.9721 for $(M_{se}(5years), M_{se}(1year))$, 0.9795 for $(M_{se}(5years), M_{se}(3years))$, and 0.8692 for the pair $(M_{se}(5years), M_{se}(10years))$.

attempt to incorporate such factors into a distributed infrastructure design. Future work may identify further factors, which must be accounted for in order to reach a practical solution. Our framework allows additional factors to be plugged in as long as they can be parameterized from publicly available datasets. Third, we modeled each AS as a single router in our simulations and did not consider the internal topology of ASes, since the micro-factors that influence intradomain topology and routing are not the focus of this paper. We are instead interested in showing how increasing peering facilitated by our framework will impact the macroscopic properties of the topology (AS path lengths) and performance (distribution of the estimated RTT among ASes). We kept this focus while designing our C-BGP model in order not to deviate from our primary goal. Finally, we did not include traffic data in our model, due to the lack of publicly available datasets about inter-domain traffic patterns. However, our topology design and simulation framework does not preclude using traffic data if it becomes available in the future; in fact, the availability of traffic data would allow us to quantify the benefits of the distributed IXP layout in terms of the amount of traffic that would be routed over shorter paths or with smaller RTTs. All these leave room for possible improvements if additional datasets and inputs become available in the future.

8.2. Feasibility of this approach from a technical and political perspective

8.2.1. Peering economics

In designing the distributed IXP layout, we did not at any stage suggest that ISPs present at an IXP should be regulated or mandated to interconnect with other ISPs; we are well aware that past examples of mandated peering have resulted in failure and have been abandoned in favour of a more market-driven approach. We instead assumed that two ISPs peer if one is not in the customer cone of the other. We recognize that there are numerous economic considerations beyond the customer-cone rule that impact real-world peering economics. Our goal was to investigate a best-case, yet realistic scenario, so as to quantitatively

demonstrate the benefits of IXP interconnection. In the real-world where business aspects, costs, and competition determine peering decisions, the number of peering links added at each step will likely be less than what we estimate.

Further, we emphasize that there are certain preconditions for our approach to be successful: ISPs in Africa need to be more open to participation at IXPs and interconnection with other local networks. Second, countries should encourage cross-border fiber deployment to enable the growth of the Internet ecosystem in the region. The quantitative framework we have developed can play a role here; specifically, demonstrating the impact that IXP interconnection could have on performance can be the biggest incentive for ISPs to join IXPs, for countries to invest in fiber crossing their borders, and for content providers to establish a presence in the region.

8.2.2. Suggested options for the feasibility of IXP interconnection

After discussing with local IXPs operators and stakeholders, we suggest the following options to build the proposed distributed IXP layout and achieve the ultimate goal of intensifying peering in the region. These alternatives involve different entities that are responsible for moving packets between IXPs. The options can possibly be combined, wherever needed (within and across sub-regions), given the interests of the IXP members.

1. First, an ISP carrier present at most local IXPs of a sub-region and at the regional hub could provide transport from local IXPs to the regional hub [17, 46]. Similarly, an ISP carrier can also provide transport between two regional hubs. Examples include Liquid Telecom [55], SEACOM [87], and MainOne [61] that have already built their own optical fiber network.
2. The set of ISPs that participate in the interconnection framework at each IXP could collectively lease wavelengths on dark fiber that already exists, and share the costs.

3. A regional carrier, both IXPs together, or a content provider with interests in the region (e.g. Google) could also invest in facilitating the interconnection. In this third category can be classified Google’s effort for the last mile internet connectivity problem [38, 78].

The goal of this study was to mostly focus on the technical aspects of the feasibility of the IXPs interconnection in the region. Investigating the sustainability IXP interconnection and investigating the feasibility of the proposed alternatives involves complex economic analysis, which is out of the scope of this paper. We leave a detailed analysis for a future work that will be focused solely on the economics of IXP interconnection, and conclude the feasibility study by providing a back-of-the-envelope cost estimate for our proposed scheme. To setup the IXP interconnection, new investments are only required in terrestrial fiber. In Africa, inland fiber deployment costs are mostly a function of labour costs; other costs, e.g. permits, rights of way, regulation, and whether the build is trans-national or metro can also add to the cost. A per-km build cost varies between US \$6,109 and US \$150,000, when all the various factors are considered, given the costs of fiber laying projects in Africa from 2011 to 2017 [44, 53, 65, 89]. With this estimate, between US \$73.9 million and US \$1.8 billion may be spent in the establishment of the backbones required for the framework (Table 1). Details on the computations are available in the technical report [32]. Almost all (99%) of the budget corresponds to step-1, in which 27 countries are involved. According to the projection, the total amount will be spent in step-1 and step-2. By the time step-3 is performed, all needed physical links will already be deployed in the two first steps.

While a detailed analysis and discussion of how this build-out cost should be supported is out of scope for this paper, we provide a few initial suggestions next. ISPs operating in the involved countries could carry the costs corresponding to their countries, since this will allow their networks to connect to the regional hub through the local IXP. They may also be (technically, financially, or politically) supported by regional fiber networks (Liquid Telecom [55], SEACOM [87], MainOne [61], etc.), large content providers such as Google [38, 78], local governments, Internet developmental institutions, or through regional projects setup by the African Union. As for the costs of infrastructure operation, we suggest that ISPs on both sides of each physical link share the operational costs based on the amount of traffic they transport over the link [58, 59].

In the long run, stakeholders should consider making the proposed infrastructure redundant to improve its robustness to outages [25]. The first step would be to complete the set of links between the regional hubs (recall that the solution in step-3 is a spanning tree and thus does not provide redundancy) so that it becomes a ring or a full-mesh for redundancy. Next, backup regional hubs could be selected. Finally, IXPs in countries that become secure

could be progressively integrated as well.

9. Conclusions and Future Work

This paper proposes a solution to the need for the African region to better localize its Internet traffic for offering affordable and better performing Internet access to end-users. As shown in previous studies [29, 30, 31, 43, 51], the African Internet suffers from significant performance problems due to a number of systemic issues including low peering density in the region and a lack of local content. However, prior proposals to address these issues (e.g., by interconnecting IXPs [27, 68, 70, 92]) are not always realizable due to the prevailing external factors. We also showed in §3 how naive approaches that do not take into account prevailing socio-economic realities of the region are infeasible in practice.

In this study, we first introduced an innovative framework that acknowledges the existence of geographical, political, and socio-economic realities that affect infrastructure design, and incorporates them as constraints in the design problem. As an example, our proposed approach relies on available cables to minimize investments and make its realization faster; it accounts for the presence of “secure” and “unsecured” countries in the region that dictate how physical infrastructure should be established in order to be feasible. A direct consequence of the implementation of this framework would be that paths from one African country to another, rather than traversing a different continent, are routed within Africa through an hierarchical IXP substrate: ISP source – local IXP (– regional IXP hub – local IXP) – ISP destination.

Next, we evaluated the proposed layout and quantified the benefits using extensive simulations with C-BGP. Our results show that our proposed solution doubles the percentage of continental intra-African paths, reduces their lengths, and drastically decreases the median of their RTTs as well as RTTs to ASes hosting top global and regional Alexa websites. Our evaluation demonstrates that it is possible to obtain shorter AS paths and better performance than what we currently have, if local ISPs intensify peering and content providers were to deploy caches at the designated regional hubs. By doing so, we highlight the potential for cross-border, sub-regional, and continental interconnection as opportunities that can be seized by a partnership between the diverse actors.

Furthermore, we identified three options to realize the proposed IXP infrastructure in 8.2.2, amongst which stakeholders of each sub-region may select given their interests. The three options differ in terms of the key entity that would be responsible for moving traffic between the IXPs. Finally, given the costs of fiber deployment projects from 2011 to 2017 [44, 53, 65, 89], we estimated the costs of inland fiber laying required to implement our approach to vary between US\$73.9 million and US\$1.8 billion, and provided some initial suggestions for how this cost could be supported.

Our proposed solution and obtained results may encourage stakeholders in other developing regions to consider similar infrastructure designs; however, we emphasize that our solution is based on numerous factors related to the nature of the existing and developing African infrastructure that may not prevail in those regions. Performing a similar analysis for other regions, while feasible, will require a careful consideration of the unique factors inherent to those regions, significant domain knowledge about the region, and focused data collection. As for our future work, we are aware that there may be further socio-economic factors beyond the ones captured in this work that influence connectivity in the African region. We plan to engage further with stakeholders in the region to discover those parameters and capture them in our framework. We also plan to reach out to local ISPs in the African region to obtain traffic data to augment our C-BGP simulations. The addition of traffic data promises to make the evaluation of the proposed approach more insightful, as it will augment estimates of path length and RTT with estimates of the traffic volume carried by those paths. Finally, including terrorist attacks and riots in the identification of unsecured countries may eliminate countries that do not appear safe but where companies are investing anyway, as cables are extensively deployed within/at their borders or their governments implement a policy environment that attracts those investments. We plan to not account for those two phenomena while identifying unsecured countries, and assess the impact of this methodological change on the proposed interconnection scheme.

10. Acknowledgements

We are grateful to the comments and advice provided by Bradley Huffaker, Vasileios Giotsas, all the CAIDA team, and Michuki Mwangi. We are thankful to all African IXPs which replied to our survey. The maps of the African Undersea Cables and terrestrial fiber networks provided by Steve Song as well as his feedback were extremely helpful in our data collection process. We thank the anonymous reviewers that helped us improve this work. The work done by Rod  rick Fanou was partially supported by IMDEA Networks Institute, US NSF grant CNS-1414177, and the project BRADE (P2013/ICE-2958) from the Directorate General of Universities and Research, Board of Education, Madrid Regional Government. Francisco Valera was partially funded by the European Commission under FP7 project LEONE (FP7-317647). Amogh Dhamdhere was partially funded by US NSF grants CNS-1414177 and CNS-1513847.

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. African Internet eXchange System. <http://pages.au.int/axis>, 2017.
- [3] African Union Commission and New Zealand Ministry. African Union Handbook 2017. Technical report, African Union, 2017.
- [4] Africanews. Cameroon restores internet in 2 anglophone regions after 93-days offline. <http://www.africanews.com/2017/04/20/cameroon-lifts-internet-blackout-in-2-anglophone-regions-after-93-days-offline/>, 2017.
- [5] AFRINIC database. <ftp://ftp.afrinic.net/>, 2017.
- [6] B. Ager, N. Chatzis, A. Feldmann, N. Sarrar, S. Uhlig, and W. Willinger. Anatomy of a Large European IXP. In *Proceedings of the ACM SIGCOMM 2012 conference on Applications, technologies, architectures, and protocols for computer communication*, pages 163–174. ACM, 2012.
- [7] Alexa. The top 500 sites on the web. <http://www.alexa.com/topsites>, 2017.
- [8] APNIC. Visible ASNs: Customer Populations (Est.). <http://stats.labs.apnic.net/cgi-bin/aspop?c=>, 2017.
- [9] APNIC database. <ftp://ftp.apnic.net/pub/stats/apnic/>, 2017.
- [10] ARIN database. <ftp://ftp.arin.net/pub/stats/arin/>, 2017.
- [11] B. Augustin, X. Cuvelier, B. Orgogozo, F. Viger, T. Friedman, M. Latapy, C. Magnien, and R. Teixeira. Avoiding traceroute anomalies with Paris traceroute. In *Proceedings of the 6th ACM SIGCOMM conference on Internet measurement (IMC)*, pages 153–158. ACM, 2006.
- [12] P. Bates. Submarine cables in Sub-Saharan Africa: Terrestrial Networks Need to Keep up. Technical report, Analysys mason, 2014.
- [13] M. Bruyere. TOUSIX first OpenFlow European IXP. <https://www.youtube.com/watch?v=VLgJcz-ops&feature=youtu.be>, 2017.
- [14] CAIDA. Automated Autonomous System (AS) ranking. Research Project. <http://as-rank.caida.org>, 2017.
- [15] CAIDA. CAIDA AS Relationships Data. Research Project. <http://data.caida.org/datasets/as-relationships/>, 2017.
- [16] CAIDA. Internet Topology Datasets Collected on the Archipelago (Ark) Infrastructure. http://www.caida.org/projects/ark/topo_datasets.xml, 2017.
- [17] I. Castro, J. C. Cardona, S. Gorinsky, and P. Francois. Remote peering: More peering without internet flattening. In *ACM CoNEXT*, pages 185–198. ACM, 2014.
- [18] Central Intelligence Agency (CIA). The World Factbook. <https://www.cia.gov/library/publications/the-world-factbook/docs/guidetowfbook.html>, 2017.
- [19] N. Chatzis, G. Smaragdakis, A. Feldmann, and W. Willinger. There is more to IXPs than meets the eye. *ACM SIGCOMM Computer Communication Review*, (5):19–28, 2013.
- [20] CNN. Cameroon goes offline after Anglophone revolt. <http://edition.cnn.com/2017/02/03/africa/internet-shutdown-cameroon/index.html>, 2017.
- [21] J. Crowcroft, A. Wolisz, and A. Sathiaselalan. Towards an Affordable Internet Access for Everyone: The Quest for Enabling Universal Service Commitment (Dagstuhl Seminar 14471). 2015.
- [22] T. Denton. Hurricane Electric’s Observations on Internet Exchange Point Management. https://conference.apnic.net/data/41/apricot-he-presentation-1_1456014430.pdf, 2016.
- [23] A. Dhamdhere and C. Dovrolis. The Internet is Flat: Modeling the Transition from a Transit Hierarchy to a Peering Mesh. In *ACM Proceedings of the 6th International Conference*, page 21. ACM, 2010.
- [24] Digital Element. Netacuity. http://www.digital-element.net/ip_intelligence/ip_intelligence.html, 2017.
- [25] R. Durairajan, P. Barford, J. Sommers, and W. Willinger. InterTubes: A Study of the US Long-haul Fiber-optic Infrastructure. pages 565–578, 2015.
- [26] East African Communications Organisation (EACO). Conclusions of the AU-Eastern Africa Regional Internet Exchange Point (RIXP) and Regional Internet Carrier (RIC) Workshop 26-30 May 2014. Technical report, 2014.

- [27] East African Communications Organisation (EACO). Report of the Meeting of East Africa Internet Exchange (EAIX) Steering Committee Held from 5-7 October 2015, Kampala-Uganda. http://www.eaco.int/docs/WGsReports/EAPIC_REPORT.docx, 2015.
- [28] European Internet Exchange Association 2012. Report on European IXPs. Technical report, 2012.
- [29] R. Fanou, P. Francois, and E. Aben. On the Diversity of Interdomain Routing in Africa. In *International Conference on Passive and Active Network Measurement*, 2015.
- [30] 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*, 106:117–135, 2017.
- [31] R. Fanou, G. Tyson, P. Francois, and A. Sathiaselan. Pushing the Frontier: Exploring the African Web Ecosystem. In *Proceedings of the 25th International Conference on World Wide Web (WWW)*, 2016.
- [32] R. Fanou, F. Valera, P. Francois, and A. Dhamdhare. Reshaping the African Internet: From Scattered Islands to a Connected Continent (Technical Report). https://fourier.networks.imdea.org/external/techrep_resaping/index, 2017.
- [33] A. Fenioux. Why and How to Interconnect IXP. <https://es.slideshare.net/InternetSociety/why-and-how-to-interconnect-ixp>, 2015.
- [34] France Diplomatie. France Diplomatie: Country Files. <http://www.diplomatie.gouv.fr/fr/conseils-aux-voyageurs/conseils-par-pays/>, 2017.
- [35] FranceIX. FranceIX: Interconnection with other IXPs. <http://www.franceix.net/en/solutions/interconnection/>, 2017.
- [36] B. Frank, I. Poesse, G. Smaragdakis, V. Aggarwal, A. Feldmann, S. Uhlig, B. Maggs, and F. Schneider. Collaboration opportunities for content delivery and network infrastructures. *ACM SIGCOMM ebook on Recent Advances in Networking*, 1(1):305–377, 2013.
- [37] H. Galperin. Connectivity in Latin America and the Caribbean: The role of Internet Exchange Points (IXPs). *Internet Society*, November, 2013.
- [38] Google. Google, Convergence Partners, International Finance Corporation, and Mitsui co-invest in CSquared. <http://www.csquared.com/>, 2017.
- [39] Google. Google Maps. <https://www.google.com/maps>, 2017.
- [40] Google Developers. Google Maps APIs: Distance Matrix API. <https://developers.google.com/maps/documentation/distance-matrix/start>, 2017.
- [41] GrenobliX. GrenobliX: Infrastructure. <http://www.grenoblix.net/en/infrastructure>, 2017.
- [42] GSMA Association. The Mobile Economy 2017. Technical report, GSMA, 2017.
- [43] 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*, 2014.
- [44] Hamilton Research Ltd 2017. Africa Bandwidth Maps. <http://www.africabandwidthmaps.com/?p=5348>, 2017.
- [45] B. Huffaker, M. Fomenkov, and K. Claffy. Geocompare: a comparison of public and commercial geolocation databases. *Proc. NMMC*, pages 1–12, 2011.
- [46] G. Huston. Interconnection, peering, and settlements. In *proc. INET (Vol. 9)*, volume 9, 1999.
- [47] IDRC and ITU. Via Africa: Creating local and regional IXPs to save money and bandwidth. Technical report, 2005.
- [48] International Telecommunication Union (ITU). ICT Facts & Figures 2016. <https://www.itu.int/en/ITU-D/Statistics/Documents/facts/ICTFactsFigures2016.pdf>, 2016.
- [49] International Telecommunication Union (ITU). Key ICT indicators for developed and developing countries and the world (total and penetration rates). <https://idp.nz/Global-Rankings/ITU-Key-ICT-Indicators/6mef-ytg6/>, 2016.
- [50] Internet Society (ISOC). Global Internet Maps: Global Internet Penetration. <http://www.internetsociety.org/map/global-internet-report/>, 2017.
- [51] 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*, pages 64–73. Springer, 2014.
- [52] M. Kende and B. Quast. Promoting Content in Africa. Technical report, Internet Society (ISOC), 2016.
- [53] E. Kisambira. East Africa invests a combined US\$400m in fiber. <http://www.techadvisor.co.uk/feature/network-wifi/east-africa-invests-a-combined-us400m-in-fiber-3282889/>, 2011.
- [54] LACNIC database. <ftp://ftp.lacnic.net/pub/stats/lacnic/>, 2017.
- [55] Liquid Telecom. A network like no other. <http://liquidtelecom.com/about-us/network-map>, 2017.
- [56] A. Lodhi, N. Larson, A. Dhamdhare, C. Dovrolis, and K. Claffy. Using PeeringDB to Understand the Peering Ecosystem. *ACM SIGCOMM Computer Communication Review*, 2014.
- [57] LyonIX. LyonIX: Infrastructure. <http://www.lyonix.net/en/infrastructure-2>, 2017.
- [58] R. T. Ma, D. M. Chiu, J. Lui, V. Misra, and D. Rubenstein. Internet Economics: The use of Shapley value for ISP settlement. *IEEE/ACM Transactions on Networking (TON)*, 18(3):775–787, 2010.
- [59] R. T. Ma, D. M. Chiu, J. C. Lui, V. Misra, and D. Rubenstein. On cooperative settlement between content, transit, and eyeball internet service providers. *IEEE/ACM Transactions on Networking (TON)*, 19(3):802–815, 2011.
- [60] G. Mahlknecht. Greg’s Cable Map. <http://www.cablemap.info/>, 2017.
- [61] MainOne. MainOne: Our Network. <http://www.mainone.net/our-network/coverage/>, 2017.
- [62] J. Manyika, A. Cabral, L. Moodley, Yeboah-Amankwah, S. Moraje, M. Chui, J. Anthonyrajah, and A. Leke. Lions go digital: The Internet’s transformative potential in Africa. Technical Report 17, McKinsey&Company, 2013.
- [63] McKinsey&Company. Offline and falling behind: Barriers to Internet adoption. Technical report, 2014.
- [64] W. Mühlbauer, A. Feldmann, O. Maennel, M. Roughan, and S. Uhlig. Building an AS-Topology Model that Captures Route Diversity. *ACM SIGCOMM Computer Communication Review*, 36(4):195–206, 2006.
- [65] M. Ncube and C. L. Lufumpa. *Infrastructure in Africa: Lessons for Future Development*. Policy Press, 2017.
- [66] Network Startup Resource Center (NSRC). Mapping Undersea and Terrestrial Fibre Optic Cables. <https://afterfibre.nsrc.org/>, 2017.
- [67] A. Nipper. Interconnecting IXPs: Pros and Cons. https://www.apricot.net/apricot2012/_data/assets/pdf_file/0003/45561/e-an-20120228-apricot2012-interconnecting-ixp_1330679773.pdf, 2012.
- [68] Northern African Community. Conclusions de l’Atelier de L’Union Africaine (UA) - Afrique du Nord sur le Point D’Echanges Internet Regional (RIXP) et la Fourniture Regionale d’Internet (RIC). Technical report, 2014.
- [69] W. B. Norton. *The 2014 Internet Peering Playbook: Connecting to the Core of the Internet*. DrPeering Press, 2014.
- [70] A. Owens. Africa: Regional Insights. https://ripe72.ripe.net/presentations/21-Ripe_2016v5_final.pdf, 2016.
- [71] Packet Clearing House (PCH). Internet Exchange Directory: Download datasets. https://prefix.pch.net/applications/ixpdir/menu_download.php, 2017.
- [72] Packet Clearing House (PCH). IXP growth per region. <https://prefix.pch.net/applications/ixpdir/summary/growth-region/>, 2017.
- [73] PeeringDB. http://www.peeringdb.com/private/exchange_list.php, 2017.

- [74] B. Pehrson and M. Ngwira. Optical Fibre for Education and Research Networks in Eastern and Southern Africa. <https://www.ubuntunet.net/sites/default/files/Sarua-fibre-final-report-draft-2006-03-04.pdf>, 2006.
- [75] I. Poese, S. Uhlig, M. A. Kaafar, B. Donnet, and B. Gueye. IP geolocation databases: Unreliable? *ACM SIGCOMM Computer Communication Review*, 41(2):53–56, 2011.
- [76] F. Poletti, N. Wheeler, M. Petrovich, N. Baddela, E. N. Fokoua, J. Hayes, D. Gray, Z. Li, R. Slavík, and D. Richardson. Towards high-capacity fibre-optic communications at the speed of light in vacuum. *Nature Photonics*, 7(4):279–284, 2013.
- [77] PriMetrica. TeleGeography Internet Exchange Map. <http://www.internetexchangemap.com/>, 2017.
- [78] Quartz Africa. Google may have a solution to Africa's last-mile internet connectivity problem. <http://qz.com/514919/google-may-have-a-solution-to-africas-last-mile-internet-connectivity-problem/>, 2015.
- [79] B. Quoitin, C. Pelsser, O. Bonaventure, and S. Uhlig. A performance evaluation of BGP-based traffic engineering. *International journal of network management*, 15(3):177–191, 2005.
- [80] B. Quoitin and S. Tandel. C-BGP user's guide. *CSE Departement, UCL, Belgique*, 2004.
- [81] B. Quoitin and S. Uhlig. Modeling the Routing of an Autonomous System with C-BGP. *IEEE Network Magazine*, 2005.
- [82] R. Ramaswami, K. Sivarajan, and G. Sasaki. *Optical networks: a practical perspective*. Morgan Kaufmann, 2009.
- [83] Reporters without Borders. Media obstructed during Chad's presidential election. <https://rsf.org/en/news/media-obstructed-during-chads-presidential-election>, 2016.
- [84] Reporters without Borders. RSF decries phone and Internet blackout during Congo election. <https://rsf.org/en/news/rsf-decries-phone-and-internet-blackout-during-congo-election>, 2016.
- [85] Réseau fibre optique: Brazzaville et Kinshasa bientôt interconnectées. Réseau fibre optique : Brazzaville et Kinshasa bientôt interconnectées. <http://adiac-congo.com/content/reseau-fibre-optique-brazzaville-et-kinshasa-bientot-interconnectees-31551>, 2016.
- [86] RIPE NCC database. <ftp://ftp.ripe.net/ripe/stats/>, 2017.
- [87] Seacom. Seacom Fiber Network. <http://seacom.mu/network-map/>, 2017.
- [88] A. Sehgal and J. Schönwälder. Getting Familiar with the C-BGP Simulator. <http://cnds.eecs.jacobs-university.de/courses/np-2011/p3-cbgp-tutorial.pdf>, 2011.
- [89] S. Song. Africa Telecoms Infrastructure in 2015. <https://manypossibilities.net/2016/01/africa-telecoms-infrastructure-in-2015/>, 2016.
- [90] S. Song. African Undersea Cables. <https://manypossibilities.net/african-undersea-cables/>, 2016.
- [91] S. Song. AfTerFiber Mapping African Terrestrial Fibre Optic Infrastructure. <https://www.flickr.com/photos/ssong/albums/72157627195113720>, 2017.
- [92] South African Development Community. Conclusions of the AU-SADC Regional Internet Exchange Point (RIXP) and Regional Internet Carrier (RIC) Workshop 03-07 February 2014. Technical report, 2014.
- [93] V. Stocker, G. Smaragdakis, W. Lehr, and S. Bauer. The growing complexity of content delivery networks: Challenges and implications for the Internet ecosystem. *Elsevier Telecommunications Policy*, 2017.
- [94] W. Stucke. Regional IXPs: the need for regional interconnection in Africa. *Information & communications technologies*, 2006.
- [95] Submarine Telecoms Forum. Submarine telecoms industry report 2016 (5th anniversary edition). Technical report, 2016.
- [96] Team Cymru. Team Cymru Services. <https://www.team-cymru.com/>, 2017.
- [97] The African IXP Association (Af-IX). List of Active Internet eXchange Points in Africa. <http://www.af-ix.net/ixps-list>, 2017.
- [98] The History Guy. Current Wars in Africa. http://www.historyguy.com/wars_of_africa_current.html, 2017.
- [99] The World Bank. The World Bank: Data By Topic. <http://data.worldbank.org/topic>, 2017.
- [100] Wikipedia. List of active rebel groups. https://en.wikipedia.org/wiki/List_of_active_rebel_groups, 2017.
- [101] Wikipedia. List of conflicts in Africa. en.wikipedia.org/wiki/List_of_conflicts_in_Africa, 2017.
- [102] Wikipedia. List of non-state terrorist incidents. https://en.wikipedia.org/wiki/List_of_terrorist_incidents, 2017.
- [103] Wikipedia. Regions of the African Union. https://en.wikipedia.org/wiki/Regions_of_the_African_Union, 2017.



Roderick Fanou graduated from Polytechnics of Abomey-Calavi (EPAC, University of Abomey Calavi), Benin Republic, in 2011 as a Design Engineer in Computer Science, Networks, and Telecommunications. He worked as network engineer at EUPHORBIA Sarl, a local Internet Service Provider (ISP) before joining IMDEA Networks Institute as a PhD student in 2013. He received a MSc. degree in Telematics Engineering from Universidad Carlos III de Madrid (UC3M), Spain, in 2014 and enrolled in the PhD program of the said university. His current research interests are impacts of Internet eXchange Points in the African region, Interdomain routing, Internet measurement, and CDNs.



Prof. Dr. **Francisco Valera** received the Telecommunication Engineering degree in 1998 from the Technical University of Madrid (UPM), and the Ph.D. in Telecommunications in 2002 from UC3M, where he is currently a tenured associate professor. He has been involved in several national and international research projects and contracts related with experimental facilities, protocol design, inter-domain routing, protocol engineering, next generation networks and multimedia systems, serving there as PI, work package leader and also as coordinator. Some of the recent research projects in which he has participated, are: IST Trilogy, IST LEONE, DRO-NE, IST MUSE or IST E-NEXT. Dr. Valera has published over 50 papers in the field of advanced communications in magazines and conferences. He has also participated in the scientific committee, organization, and technical review in different national and international conferences.



Dr. **Pierre Francois** obtained his Ph.D degree in computer science from Université Catholique de Louvain, Belgium, in 2007. He is an Engineering manager at Cisco Systems since 2015. Before that, he was an associate researcher at IMDEA Networks where he led a research team focused on Interdomain and Intradomain. His work includes several papers published in top conferences and journals within the networking field, as well as multiple Internet Engineering Task Force (IETF) Working Group documents and RFCs, in various working groups of the IETF Routing and IETF Operations and Management areas.



Dr. **Amogh Dhamdhare** is a Research Scientist at the Cooperative Association for Internet Data Analysis (CAIDA), based at the University of California, San Diego. Prior to joining CAIDA in 2009, He obtained a Ph.D. in Computer Science from the College of Computing at Georgia Tech. His research mainly focuses on measurement and modeling of Internet topology, traffic, and economics. More recently, He has worked on measurement of IPv6 adoption, and building systems for localizing and quantifying interdomain congestion.