uc3m | Universidad Carlos III de Madrid

UNIVERSIDAD CARLOS III DE MADRID

TESIS DOCTORAL

METHODS FOR REVEALING AND RESHAPING THE AFRICAN INTERNET ECOSYSTEM AS A CASE STUDY FOR DEVELOPING REGIONS: FROM ISOLATED NETWORKS TO A CONNECTED CONTINENT

Autor: Rodérick Fanou, IMDEA Networks Institute and
Universidad Carlos III de Madrid

Director[/es]: Prof. Dr. Francisco Valera, Universidad Carlos III de Madrid
Dr. Pierre Francois

DEPARTAMENTO DE INGENIERÍA TELEMÁTICA

Leganés (Madrid), Diciembre de 2017

uc3m | Universidad Carlos III de Madrid

UNIVERSIDAD CARLOS III DE MADRID

PH.D. THESIS

METHODS FOR REVEALING AND RESHAPING THE AFRICAN INTERNET ECOSYSTEM AS A CASE STUDY FOR DEVELOPING REGIONS: FROM ISOLATED NETWORKS TO A CONNECTED CONTINENT

Autor: Rodérick Fanou, IMDEA Networks Institute and
Universidad Carlos III de Madrid

Director[/s]: Prof. Dr. Francisco Valera, Universidad Carlos III de Madrid
Dr. Pierre Francois

DEPARTMENT OF TELEMATIC ENGINEERING

Leganés (Madrid), December 2017

Methods for Revealing and Reshaping the African Internet Ecosystem as a Case Study for Developing Regions: From Isolated Networks to a Connected Continent

A dissertation submitted in partial fulfillment of the requirements for the degree of Doctor of Philosophy

Prepared by
Rodérick Fanou, IMDEA Networks Institute and Universidad Carlos III de Madrid

Under the advice of
Prof. Dr. Francisco Valera, Universidad Carlos III de Madrid
Dr. Pierre Francois

Departamento de Ingeniería Telemática, Universidad Carlos III de Madrid

---

Date: December, 2017

Web: http://people.networks.imdea.org/~roderick_fanou/

Contact: roderick.fanou@imdea.org

This work has been supported by IMDEA Networks Institute.

TESIS DOCTORAL

METHODS FOR REVEALING AND RESHAPING THE AFRICAN INTERNET ECOSYSTEM AS A CASE STUDY FOR DEVELOPING REGIONS: FROM ISOLATED NETWORKS TO A CONNECTED CONTINENT

Autor: Rodérick Fanou, IMDEA Networks Institute and Universidad Carlos III de Madrid

Director[/es]: Prof. Dr. Francisco Valera, Universidad Carlos III de Madrid
Dr. Pierre Francois

Firma del tribunal calificador:

Presidente:

Vocal:

Secretario:

Calificación:

Leganés, de de

Acknowledgements

Before getting to the heart of the matter, I would like to express my gratitude to my advisors for contributing in guiding my steps in both the Industry and Research communities. In particular, I would like to thank Dr. Pierre Francois, for his support during my first years in Europe, for letting my inspiration flow and for sowing while advising me, the seeds of what will later become this thesis. I would also like to address my profound gratitude to Prof Dr. Francisco Valera for taking the floor so smoothly: apart from the results-oriented research focus, situations management, patience, tact, and good writing are the top aspects that working at your side helped me to develop. I thank you for advising me until the end so that I can make this dream a reality. I am also thankful to Dr. Amogh Dhamdhere for guiding me both during and after my internship. I will not forget how my technical capabilities were strengthened at your side. Needless to remember that after a meeting with each of you, I have always seen difficulties as opportunities.

I am deeply grateful to IMDEA Networks Institute and especially its Director Prof. Dr. Arturo Azcorra and its Deputy Director Prof. Dr. Albert Banchs for the opportunity of pursuing my academic studies with top Professors and top Researchers in a highly competitive environment and for their unwavering support for this project. I would like to thank Dr. Antonio Fernández Anta, Dr. Paolo Casari, Dr. Vincenzo Mancuzo, Dr. Marcelo Bagnulo, Dr. Andra Lutu, Dr. Narseo Vallina Rodriguez for their encouraging comments and pieces of advice during these years. I will not forget the availability of the IMDEA admin, especially Brian Dunne, Flora Quintans, and Sonia Balaguer, who helped me provide on time the unbelievable amount of necessary documents for my different visa processes or that of Dr. Jose Felix Kukielka and his team for always doing their best to satisfy my continually growing needs of more technical resources. I would like to express my highest appreciation to Rebeca de Miguel for her tireless efforts, ideas, and pieces of advice to give more visibility to the results of this work and draw the attention of the community.

Following that, I wish to thank all my co-authors notably Víctor Sánchez-Agüero, Eder Leao Fernandes, Dr. Gareth Tyson, and Dr. Arjuna Sathiaseelan for their priceless contributions to the work presented in this thesis. It has been a privilege to work with you.

Much appreciation goes to those I have taken for granted, as they have always been available for supporting my initiatives along the way. I express my gratitude to Michuki Mwangi, Mathieu Paonessa, Nishal Goburdhan, Frank Habicht, Dr. Dawit Bekele, and Jane Coffin, for introducing me to the Africa Peering and Interconnection Forum (AfPIF) community. Discussions with

Michuki, Nishal, and Frank have been insightful, their comments constructive, and their pieces of advice knowledgeable. Special thanks also go to them for not hesitating in introducing me to potential probe hosts, for simultaneously deploying probes notably in Southern and Eastern Africa, or for their numerous invitations at organized African Internet eXchange System (AXIS) workshops and operators meetings, which I could unfortunately not all honor with my presence. I am deeply grateful to Michuki Mwangi, Jane Coffin, and the Internet Society (ISOC) for their invaluable support on the African Route-collectors Data Analyzer (ARDA) project. Nor will I forget, how the logistics for my trips in the context of the AXIS project were handled or the promptitude to end my retention at a couple of airports for visa issues by Marsema Tariku, Betel Hailu, and Hisham Ibrahim. I thank you so much. Further, I owe my profound gratitude to Vesna Manojlovic, Emile Aben, Massimo Candela, Robert Kisteleki, Lia Hestina, Alun Davies, and the RIPE Atlas team, for their tireless efforts to support our initiatives, our constant needs for more measurements capabilities with the RIPE Atlas network, their constant pieces of advice, and so one. To Dr. Bradley Huffaker, Dr. Matthew Luckie, Paul Hick, Josh Polterock, Daniel Andersen, Dr. Alberto Dainotti, and the CAIDA team, it was a real pleasure working alongside you and learning from your experience: my stay in your premises was knowledgeable partly thanks to your comments, pieces of advice, and availability.

Also, I gratefully thank the experts Dr. Dawit Bekele and Prof. Dr. Kc Claffy for their detailed evaluation of this thesis on time, thus supporting in an invaluable way the process leading to its defense. In particular, I appreciate the additional efforts of Prof. Dr. Kc Claffy for providing insightful comments and a very detailed review of this doctoral thesis, which contributed to significantly improving the final document. I also wish to thank Raquel Moreno for her availability and flexibility, which were more than helpful during the process leading to the Ph.D. defense.

I wish to thank ISPs engineers, stakeholders, probe hosts, etc., who trusted us during this trip, kept the probes online within their network, or answered to our surveys despite their busy schedule. Nor will I forget those with whom I extensively discussed trending topics through online meetings/interviews. Although I can not list them all, I would particularly like to thank Emmanuel Kwarteng, Mohamed Faye, Michael Otieno, Victor Oyetola, and John Aoga.

Special thanks are expressed to my family for their incredible support in times of difficulties. Above all, I learned that my stamina and competitiveness is also related to the diversity in my work environment and the good relationships with my amazing friends. I spent my best and unforgettable moments in Spain in the company of Allison, Amr, Ander, Antonio, Arash, Aymen, Camilo, Christian, Dario, Elli, Evgenia, Foivos, Ginés, Guido, Guillermo, Hany, Ignacio, Joan, Jona, Luisfo, Margherita, Maurizio, Noelia, Pablo, Roberto, Rosa, Sofia, Victor, to only name a few. For all those I could unfortunately not list due to the lack of room, I thank you so much.

Abstract

While connecting end-users worldwide, the Internet increasingly promotes local development by making challenges much simpler to overcome, regardless of the field in which it is used: governance, economy, education, health, etc. However, African Network Information Centre (AfriNIC), the Regional Internet Registry (RIR) of Africa, is characterized by the lowest Internet penetration: 28.6 % as of March 2017 compared to an average of 49.7 % worldwide according to the International Telecommunication Union (ITU) estimates [139]. Moreover, end-users experience a poor Quality of Service (QoS) provided at high costs. It is thus of interest to enlarge the Internet footprint in such under-connected regions and determine where the situation can be improved. Along these lines, this doctoral thesis thoroughly inspects, using both active and passive data analysis, the critical aspects of the African Internet ecosystem and outlines the milestones of a methodology that could be adopted for achieving similar purposes in other developing regions.

The thesis first presents our efforts to help build measurements infrastructures for alleviating the shortage of a diversified range of Vantage Points (VPs) in the region, as we cannot improve what we can not measure. It then unveils our timely and longitudinal inspection of the African *interdomain routing* using the enhanced RIPE Atlas measurements infrastructure for filling the lack of knowledge of both IPv4 and IPv6 topologies interconnecting local Internet Service Providers (ISPs). It notably proposes reproducible data analysis techniques suitable for the treatment of any set of similar measurements to infer the behavior of ISPs in the region. The results show a large variety of transit habits, which depend on socio-economic factors such as the language, the currency area, or the geographic location of the country in which the ISP operates. They indicate the prevailing dominance of ISPs based outside Africa for the provision of intra-continental paths, but also shed light on the efforts of stakeholders for *traffic localization*.

Next, the thesis investigates the causes and impacts of *congestion* in the African *IXP substrate*, as the prevalence of this endemic phenomenon in local Internet markets may hinder their growth. Towards this end, Ark monitors were deployed at six strategically selected local Internet eXchange Points (IXPs) and used for collecting Time-Sequence Latency Probes (TSLP) measurements during a whole year. The analysis of these datasets reveals no evidence of widespread congestion: only 2.2 % of the monitored links experienced noticeable indication of congestion, thus promoting *peering*. The causes of these events were identified during IXP operator interviews, showing how essential collaboration with stakeholders is to understanding the causes of

performance degradations.

As part of the Internet Society (ISOC) strategy to allow the Internet community to profile the IXPs of a particular region and monitor their evolution, a *route-collector data analyzer* was then developed and afterward, it was deployed and tested in AfriNIC. This open source web platform titled the "African" Route-collectors Data Analyzer (ARDA) provides metrics, which picture in real-time the status of interconnection at different levels, using public routing information available at local route-collectors with a peering viewpoint of the Internet. The results highlight that a small proportion of Autonomous System Numbers (ASNs) assigned by AfriNIC (17 %) are peering in the region, a fraction that remained static from April to September 2017 despite the significant growth of IXPs in some countries. They show how ARDA can help detect the impact of a policy on the IXP substrate and help ISPs worldwide identify new interconnection opportunities in Africa, the targeted region.

Since broadening the underlying network is not useful without appropriately provisioned services to exploit it, the thesis then delves into the availability and utilization of the *web infrastructure* serving the continent. Towards this end, a comprehensive measurement methodology is applied to collect data from various sources. A focus on Google reveals that its content infrastructure in Africa is, indeed, expanding; nevertheless, much of its web content is still served from the United States (US) and Europe, although being the most popular content source in many African countries. Further, the same analysis is repeated across top global and regional websites, showing that even top African websites prefer to host their content abroad. Following that, the primary bottlenecks faced by Content Providers (CPs) in the region such as the lack of peering between the networks hosting our probes and poorly configured DNS resolvers are explored to outline proposals for further ISP and CP deployments.

Considering the above, an option to enrich connectivity and incentivize CPs to establish a presence in the region is to interconnect ISPs present at isolated IXPs by creating a *distributed IXP layout* spanning the continent. In this respect, the thesis finally provides a four-step interconnection scheme, which parameterizes socio-economic, geographical, and political factors using public datasets. It demonstrates that this constrained solution doubles the percentage of continental intra-African paths, reduces their length, and drastically decreases the median of their Round Trip Times (RTTs) as well as RTTs to ASes hosting the top 10 global and top 10 regional Alexa websites. We hope that quantitatively demonstrating the benefits of this framework will incentivize ISPs to intensify peering and CPs to increase their presence, for enabling fast, affordable, and available access at the Internet frontier.

Table of Contents

List of Tables

List of Figures

List of Acronyms

ADSL Asymmetric Digital Subscriber Line

Af-IX The African IXP association

AfPIF Africa Peering and Interconnection Forum

AfriNIC African Network Information Centre

AIMS the Workshop on Active Internet Measurements

AO Angola

APNIC Asia-Pacific Network Information Center

ARDA African Route-collectors Data Analyzer

ARIN American Registry for Internet Numbers

AS Autonomous System

ASes Autonomous Systems

ASN Autonomous System Number

AU African Union

AXIS African Internet eXchange System

BE Belgium

BF Burkina Faso

BFS Breadth-First Search

BG Bulgaria

BGP Border Gateway Protocol

BH Bahrain

BI Burundi

BJ Benin

BR Brazil

BW Botswana

c2p customer-to-provider

CA Canada

CAf Central Africa

CAIDA Center for Applied Internet Data Analysis

CC Country Code

CD DR Congo

CDN Content Delivery Network

CG Congo

CGIX Congo Internet eXchange

CH Switzerland

CI Ivory Coast

CM Cameroon

CN China

CORE Computing Research & Education

CP Content Provider

CV Cabo Verde

DB Database

DE Germany

DJ Djibouti

DNS Domain Name Server

DS Data Source

DSes Data Sources

DZ Algeria

EAf East Africa

EDNS0 Extension mechanisms for DNS

EG Egypt

EH Western Sahara

ER Eritrea

ES Spain

ET Ethiopia

EU Europe

FI Finland

FR France

GA Gabon

GCB Google Course Builder

GDP Gross Domestic Product

GGC Google Cache

GH Ghana

GIXA Ghana Internet eXchange Association

GM Gambia

GN Guinea

GPS Global Positioning System

GQ Equatorial Guinea

GW Guinea-Bissau

HTTP Hypertext Transfer Protocol

IANA Internet Assigned Numbers Authority

ICMP Internet Control Message Protocol

ICT Information and Communications Technology

IDRC International Development Research Center

IE Ireland

IN India

Int Intercontinental

IP Internet Protocol

IQR Interquartile Range

ISOC Internet Society

ISP Internet Service Provider

IT Internet Technology

ITU International Telecommunication Union

IXP Internet eXchange Point

IXPN Internet eXchange Point of Nigeria

JB Johannesburg

JINX Johannesburg Internet eXchange

KE Kenya

KIXP Kenya Internet eXchange Point

KM Comoros

LAC Latin America and the Caribbean

LACNIC Latin America & Caribbean Network Information Center

LAN Local Area Network

LR Liberia

LS Lesotho

LY Libya

MA Morocco

MARWAN Moroccan NREN

MG Madagascar

MGIX Madagascar Internet eXchange

ML Mali

MM Maxmind

MR Mauritania

MU Mauritius

MW Malawi

MY Malaysia

MZ Mozambique

NA Namibia

NAf North Africa

NAm North America

NAT Network Address Translation

NDA Non-Disclosure Agreement

NE Niger

NG Nigeria

NL Netherlands

NOC Network Operating Center

NREN National Research and Education Network

OSI Open Systems Interconnection

OUA Organization of African Unity

PCH Packet Clearing House

p2p peer-to-peer

PE Peru

QoS Quality of Service

RE Reunion

RECs Regional Economic Communities

RINEX Rwanda Internet eXchange

RIPE Réseaux IP Européens

RIPE NCC Réseaux IP Européens Network Coordination Center

RIR Regional Internet Registry

RO Romania

RR Record-routes

RTT Round Trip Time

RW Rwanda

s2s siblings-to-siblings

SAf Southern Africa

SC Seychelles

SD Sudan

SDN Software Defined Networks

SE Sweden

SH Saint Helena

SIXP Serekunda IXP

SL Sierra Leone

SLA Service Level Agreement

SN Senegal

SO Somalia

SS South Sudan

SSH Secured SHell

ST Sao Tome and Principe

SZ Swaziland

TC Team Cymru

TCP Transmission Control Protocol

TD Chad

TG Togo

TIX Tanzania IXP

TN Tunisia

TSLP Time-Sequence Latency Probes

TUREN Tunisia NREN

TZ Tanzania

UC3M Universidad Carlos III de Madrid

UDP User Datagram Protocol

UG Uganda

UK United Kingdom

UN United Nations

US United States

VP Vantage Point

VPN Virtual Private Network

WACREN West and Central African Research and Education Network

WAf West Africa

YT Mayotte

ZA South Africa

ZM Zambia

ZW Zimbabwe

Chapter 1

Introduction and Background

To provide the reader with the necessary background information for comprehending this study and better understanding our results, we first present the context in which this research has been performed. We then detail our motivations and highlight our targeted objectives. Next, we shed light on the study process and outline the structure of the remainder of the thesis. Finally, we enumerate the publications on which this work is based and briefly detail how we contributed to each of them.

1.1. Context of the study

Apart from bringing people closer and removing boundaries among nations, the Internet nowadays plays the role of a knowledge sharing space that considerably impacts the local development. There exist countless examples of its positive effects wherever it is accessible and affordable: when applied to communications, social relationships, governance, economy, education, health, etc., it makes challenges much simpler to overcome and helps achieve efficient and effective results [51, 75, 181, 270, 295]. In this respect, the need for a better Internet access for everyone and particularly in developing regions is the primary motivation of this work.

1.1.1. Internet and Internet number resources

Technically speaking, the Internet is a set of networks interconnected and collaborating with one another so that end-to-end communications can take place. It is composed of Autonomous Systems (ASes) namely Internet Service Providers (ISPs),[1] universities, and private companies. Each of them is a group of networks with the same routing policy and a single routing protocol, often owned by the same entity and operating under a sole administrative control [121]. A given Autonomous System (AS) has a globally unique identifier termed ASN,[2] which is used for the

[1] The term ISP refers to an AS that provides for a fee, telecommunications and Internet services to individuals, universities, or companies, etc., which are registered as its customers.

[2] Note that in this thesis, we link an abbreviation with its definition in the section "List of acronyms."

exchange of exterior routing information (between neighboring ASes) [238, 272]. As of August 2017, the Internet comprises over 58,000 ASes [21, 178], through which 49.7 % of the worldwide population are connected according to the ITU estimates [139] of March 2017.

ASes exchange routing information using Border Gateway Protocol (BGP) version 4 [33, 238, 272]. They are connected to one another through business relationships such as typically *customer-to-provider* (*c2p*), *peer-to-peer* (*p2p*), or *siblings-to-siblings* (*s2s*) [102, 129, 130]. To these three traditional approaches of modeling relationships between ASes, more complex AS relationships that appear as special cases of *p2p* and *c2p*, can be added: *hybrid relationships* and *partial transit* [92, 107, 169]. These AS relationships, which are an essential aspect of the Internet structure, result from commercial agreements among administrative domains; they enable the traffic flow, which is always in the opposite direction of the flow of routing information. *c2p* links are used by transit providers (large ASes) to provide services for a fee to their customers (smaller ASes) [102, 130]. *p2p* links or settlement-free peering links are established between ASes, which agree to exchange traffic between their customers, free of charge [102, 130]. They can be set up as direct point-to-point links or set up at a public peering point called Internet eXchange Point (IXP) [13, 48]. An IXP is a shared layer-2 switch fabric environment, with three or more members (ASes), and over which the members peer with each other, exchanging customer routes [48, 257]. An *s2s* relationship is established between two ASes owned by the same entity, which may mutually provide transit to one another.

The Internet can be divided into two main parts: the edge *i.e.*, end-systems or hosts, and the core of the network *i.e.*, the routers or the "dumb" network. To deal with its complexity, enable easy maintenance and system update, as well as build a model for discussions, etc., diverse layered reference models have been introduced. These include the Open Systems Interconnection (OSI) model, and finally the Transmission Control Protocol (TCP)/IP model on which the Internet is based [75, 290]. In the TCP/IP model, the IP or network layer has a unifying effect, as it works in the same way on top of any type of physical link (*e.g.*, ethernet link, radio link, etc.). It ensures packets forwarding hop by hop from the source IP to the destination IP. This task is performed by routers based on their knowledge of the network topology. In this thesis, we break down the IP networking in the African region bearing in mind that it affects any underlying network.

Similarly to the unique identification of an AS, each host on the Internet is identified by a unique public IP address (obviously, this is not valid for the IP anycast architecture or for a Network Address Translation (NAT) gateway behind which many hosts are attributed private IP addresses). In fact, IP packet headers contain the source and the destination IP addresses of the packets traversing the network and are used by routers to forward the packets towards the destination. IP addresses and ASNs are termed *Internet number resources*. These are uniquely delegated by the Internet Assigned Numbers Authority (IANA) [206] to organizations commonly known as RIRs. There are five Internet regions around the world (Figure 1.1) for which the Internet number resources are administrated by their respective RIRs: Réseaux IP Européens Network

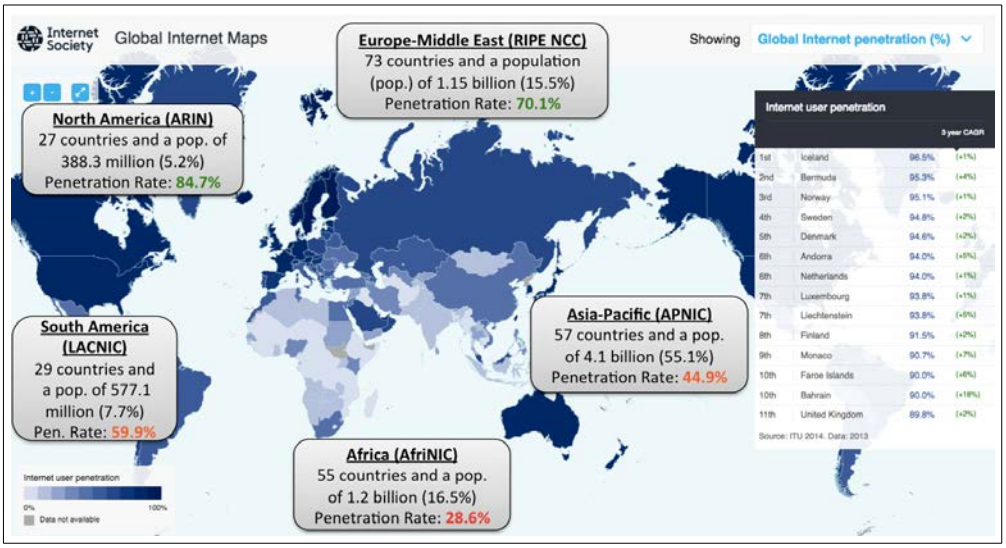

Figure 1.1: Population and Internet penetration rate per Internet region as of March 2017 computed based on data collected from the public data sources [139, 146, 149, 205]

Coordination Center (RIPE NCC),[3] Asia-Pacific Network Information Center (APNIC),[4] American Registry for Internet Numbers (ARIN),[5] Latin America & Caribbean Network Information Center (LACNIC),[6] and the AfriNIC,[7] listed in the order of their establishment [206]. Each RIR then uniquely allocates its Internet number resources to local ISPs and large organizations, and each ISP or the IT department of each company, in turn, assigns them to end-users.

Now that the Internet has been penetrating most of these regions [146] and that Internet access has turned out to be a rights enabler in connected areas, there is a growing interest from the community in understanding the barriers to Internet adoption and narrowing the digital divide between the developed and developing world [61, 181, 222, 270]. Figure 1.1 presents, based on data collected from [139, 146, 149, 205], both the population and the Internet penetration rate per region. A comparison of these values indicates that the AfriNIC region, with a total population of 1.2 billion (16.6 % of the world population), registers the lowest penetration rate. In fact, only 28.6 % of its inhabitants (by March 2017) can access the Internet, a percentage slightly higher than the half of that of the world (*i.e.*, 49.7 %), as per the ITU estimates [137, 139, 146].

1.1.2. A continent struggling to eradicate poverty

Africa is the second-largest continent after Asia in size and population [236, 310]. In 2010, Krause showed, in an infographic [159] aiming at revealing the geographical dimensions of Africa, that it can contain the entirety of the US,[8] China (CN), India (IN), as well as Japan and

[3] for Europe, the Middle East, and parts of Central Asia, created in 1992
[4] for the Asia-Pacific region, created in 1993
[5] for Canada (CA), the United States (US), several Caribbean and North Atlantic islands, created in 1997
[6] for the Latin American and Caribbean regions, 2002
[7] for Africa, created in 2005
[8] We refer to countries using ISO 2-letter Country Codes (CCs), which are also listed among the acronyms.

most parts of Europe (EU), all combined. These have been discussed in [301] and can be verified using the freely accessible tool *True Size Of* [110]. In this regard, Table 1.1 lists all African territories, with their area, number of inhabitants, and official languages. It also specifies their coastal or inland (*i.e.*, landlocked) status and currencies; N/A stands for "Non-Available," Coast., for "coastal countries," and Inl., for "inland countries." Table 1.1 indicates that the population of Africa is lower than that of China or India (1.3 billion each), but higher than that of North America (363.2 million) or EU (822.7 million) [149, 312]. The table also shows that the area of Africa is of 30.5 million km^2, equivalent to 23.6 % of the world land [312]. These explain the high differences noticed throughout the continent, especially regarding telecoms services provision when moving from an African country to another.

The African continent is composed of 55 territories of which one (Western Sahara) is considered as self-proclaimed (for reasons that are out of the scope of this work), and is thus not recognized by the United Nations (UN). Since Western Sahara is also governed by AfriNIC [205] and receives its Internet number resources from that RIR, we list it in Table 1.1 as well. To adopt a neutral position, we will use the term *African territories* in the remainder of this thesis, when Western Sahara is included in the set of considered territories and the term *African countries* whenever it is not. 30 (54.5 %) of the African territories are coastal ones, while the remaining ones are inland. In Chapter 5 for instance, we include in our analysis neighboring islands such as Mayotte (YT), Reunion (RE), Saint Helena (SH), etc., territories that are not parts of the AfriNIC region, but that are geographically close to Africa.

Most African countries are underdeveloped and struggling to eradicate poverty. To achieve this goal, they set up altogether, in 2001, the African Union (AU) on the ashes of the Organization of African Unity (OUA), as the umbrella organization that gathers all countries of the region and coordinates decisions aiming at promoting their development and the wellbeing of their citizens. Still, Deaton *et al.* noticed in 2005 that measured poverty had fallen less rapidly than it appears warranted by growth measured in developing countries [63]. Using the World Bank's past estimates of global poverty combined with better data, Chen *et al.* [52] later showed that a quarter of the population of the developing world was living below the international line of US$1.25 a day in 2005 prices. They also demonstrated that the poverty rate stayed at 50 % in Sub-Saharan Africa over a period of 25 years. Next, Fosu [95] suggested in 2015 that recent progress on poverty had been substantial contrary to that registered in the 90s. He underlined that in the meantime, however, the low Gross Domestic Product (GDP) per capita inhibited the effectiveness of income and inequality improvements in reducing poverty in many African countries [68, 95, 237].

The African territories can be classified into five distinct *sub-regions*: *North Africa* (*NAf*), *West Africa* (*WAf*), *East Africa* (*EAf*), *Central Africa* (*CAf*), and *Southern Africa* (*SAf*). Table 1.1 also illustrates this classification. The sets of countries that constitute those *sub-regions* are respectively identical to those of the five geographic regions – defined by the OUA in 1976 (CM/Res.464QCXVI) – in which are divided the AU member states [8, 9, 309]. Moreover, countries in the same sub-regions often share history, culture (*e.g.*, *SAf*, *NAf*), official language (*e.g.*,

Table 1.1: List of the 55 territories of the African continent, gathered by sub-regions as per the African Union (AU) [7–9, 297].

African Sub-regions (commonalities)	Country (CC)	Area (in km^2)	Number of inhabitants	Int. Pen. Rate	Main official language(s)	Coast. or inl.	Currency
North Africa (Official language, culture)	Algeria (DZ)	2,381,740	41,063,753	45.2 %	Arabic, French	Coast.	DZD
	Egypt (EG)	1,001,451	95,215,102	36.5 %	Arabic	Coast.	EGP
	Libya (LY)	1,759,540	6,408,742	43.7 %	Arabic	Coast.	LYD
	Mauritania (MR)	1,030,700	4,266,448	16.7 %	Arabic, French	Coast.	MRO
	Morocco (MA)	710,850	35,241,418	57.3 %	Arabic, French	Coast.	MAD
	Tunisia (TN)	163,610	11,494,760	50.4 %	Arabic	Coast.	TND
	Western Sahara (EH) (proclaimed)	281,000	596,021	4.5 %	N/A	Coast.	N/A
Southern Africa (Official language, culture)	Angola (AO)	1,246,700	26,655,513	22.3 %	Bantu, Portuguese	Coast.	AOA
	Botswana (BW)	581,726	2,343,981	29.4 %	English, Setswana	Inl.	BWP
	Lesotho (LS)	30,355	2,185,159	20.3 %	Sesotho, English	Inl.	LSL
	Malawi (MW)	118,484	18,298,679	9.1 %	English, Nyanja	Inl.	MWK
	Mozambique (MZ)	801,590	29,537,914	6.2 %	Portuguese	Coast.	MZN
	Namibia (NA)	824,116	2,568,569	20.2 %	English	Coast.	NAD
	South Africa (ZA)	1,221,037	55,436,360	51.5 %	Afrikaans, English	Coast.	ZAR
	Swaziland (SZ)	17,364	1,320,356	33.1 %	English, siSwati	Inl.	SZL
	Zambia (ZM)	752,618	17,237,931	30.1 %	English	Inl.	ZMW
	Zimbabwe (ZW)	390,757	16,337,760	41.1 %	English	Inl.	ZBN
East Africa (Official language)	Comoros (KM)	2,235	825,920	7.3 %	Arabic, French	Coast.	KMF
	Djibouti (DJ)	23,200	911,382	16.4 %	French, Arabic	Coast.	DJF
	Ethiopia (ET)	1,104,300	104,344,901	11.1 %	Amharic	Inl.	ETB
	Eritrea (ER)	117,600	5,481,906	1.3 %	Tigrinya, Arabic, English	Coast.	ERN
	Kenya (KE)	580,367	48,466,928	81.8 %	English, Kiswahili	Coast.	KES
	Madagascar (MG)	587,041	25,612,972	5.1 %	French, Malagasy	Coast.	MGA
	Mauritius (MU)	2,040	1,281,353	62.7 %	English, French	Coast.	MUR
	Rwanda (RW)	26,798	12,159,586	30.6 %	Rwanda French, English	Inl.	RWF
	Seychelles (SC)	451	97,539	57.6 %	English, French	Coast.	SCR
	Somalia (SO)	637,661	11,391,962	5.8 %	Somali	Coast.	SOS
	South Sudan (SS)	619,745	13,096,190	16.6 %	Arabic	Inl.	SSP
	Sudan (SD)	1,886,068	42,166,323	25.8 %	Arabic	Coast.	SDG
	Tanzania, United Republic of (TZ)	945,087	56,877,529	6.5 %	Kiswahili, Kiunguju, English	Coast.	TZS
	Uganda (UG)	236,040	41,652,938	31.2 %	English	Inl.	UGX
West Africa Official language, currency XOF – XAF CFA franc for most french-speaking countries)	Benin (Benin (BJ))	112,622	11,458,611	10.7 %	French	Coast.	XOF
	Burkina Faso (BF)	274,000	19,173,322	11.2 %	French	Coast.	XOF
	Cabo Verde (CV)	4,033	533,468	44.1 %	Portuguese	Coast.	CVE
	Gambia (GM)	10,380	2,120,418	17.6 %	English	Coast.	GMD
	Ghana (GH)	238,534	28,656,723	27.8 %	English	Coast.	GHS
	Guinea (GN)	245,857	13,290,659	7.1 %	French	Coast.	GNF
	Guinea-Bissau (GW)	36,125	1,932,871	4.3 %	Portuguese	Coast.	XOF
	Ivory Coast (CI)	322,462	23,815,886	21.9 %	French	Coast.	XOF
	Liberia (LR)	111,369	4,730,437	8.3 %	Liberia	Coast.	LRD
	Mali (ML)	1,240,192	18,689,966	11.8 %	French	Inl.	XOF
	Niger (NE)	1,267,000	21,563,607	2 %	French	Inl.	XOF
	Nigeria (NG)	923,768	191,835,936	48.8 %	English	Coast.	NGN
	Senegal (SN)	196,723	16,054,275	22.7 %	French	Coast.	XOF
	Sierra Leone (SL)	71,740	6,732,899	4.6 %	English	Coast.	SLL
	Togo (TG)	56,785	7,691,915	7.1 %	French	Coast.	XOF
Central Africa (Official language, currency XAF CFA franc for most french-speaking countries)	Burundi (BI)	27,834	11,936,481	4.4 %	Kirundi, French	Inl.	BIF
	Cameroon (CM)	475,442	24,513,689	20 %	English, French	Coast.	XAF
	Central African Republic (CF)	622,984	5,098,826	4.8 %	French, Sangho	Inl.	XAF
	Chad (TD)	1,284,000	14,965,482	2.6 %	French, Arabic	Inl.	XAF
	Congo (CG)	342,000	4,866,243	8.2 %	French	Coast.	XAF
	Congo, Democratic Republic of the (CD)	2,345,409	82,242,685	3.8 %	French	Coast.	CDF
	Equatorial Guinea (GQ)	28,051	894,464	20.3 %	Spanish, French	Coast.	XAF
	Gabon (GA)	267,667	1,801,232	37.2 %	French	Coast.	XAF
	Sao Tome and Principe (ST)	964	198,481	25 %	Portuguese	Coast.	STD
5 African sub-regions	55 territories	30,588,212	1,245,374,471	28.6 %			

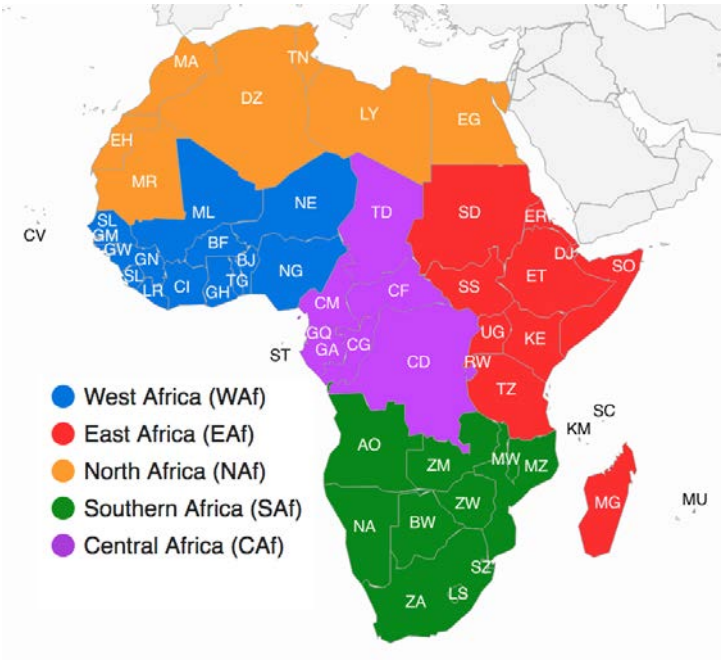

(a) *Map of the African continent with its territories classified by sub-regions. Each territory is labelled with its ISO 2-letter country code.*

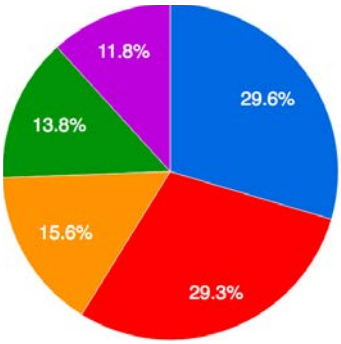

(b) *Percentage of the total inhabitants in African countries per sub-region.*

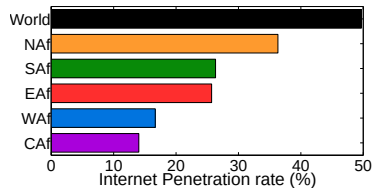

(c) *Average Internet Penetration rate in each African sub-region compared to the average rate worldwide.*

Figure 1.2: The African continent, its territories, and its sub-regions along with their respective average Internet penetration rates

WAf, *CAf*), or currency (*e.g.*, *CAf*, *WAf*). Furthermore, the 8 Regional Economic Communities (RECs)[9] of the AU [7–9, 297] are based on these sub-regions; they bring together countries that share commonalities and are willing to cooperate on certain aspects.

By taking those shared values (*cf.* Table 1.1) into consideration and putting in perspective the success of this cooperation, it goes without saying that there is a need for fast, cheap, stable, and high-quality communication within countries, among countries of the same sub-region, within the continent, as well as from the continent to other ones. Such needs are essential, especially for those developing nations: fulfilling them will open the door to sharing competencies, knowledge, know-how, and experiences among one another and with the rest of the world, necessary to overcome their common enemies: illiteracy and poverty. In a 21st century characterized by a rapid pace in the cross-sectional field of telecommunications, which often contributes to boosting all other areas, the Internet appears as the most effective way to achieve the goals mentioned above. In light of this, the AU recognizes, according to [207], Information and Communications Technology (ICT) as a significant sector that can help achieve its Agenda 2063 aspirations.

[9] *e.g.*, Economic Community of West African States (ECOWAS), East African Community (EAC), Southern African Development Community (SADC), Community of Sahel-Saharan States (CEN-SAD), etc.

1.1.3. Motivations

As a native of Benin (BJ, *WAf*) having grown, studied, worked, and traveled around the West African sub-region during my first 24 years, I was able to experience (notably from 2000 to 2012) the low penetration and the poor QoS offered, at high costs, by local network operators. To share my personal experience, while doing the paperwork for pursuing my studies in Spain five years ago, I was living in a neighborhood poorly covered by local ISP infrastructures and located at 27 km from Cotonou, the economic capital city of BJ [109]. Anytime I had to communicate with people or institutions overseas (*i.e.*, perform basic operations such as consulting my mailbox, respond to inquiries, send personal details or other information, etc.), I had to travel roughly 10 km to reach the closest cybercafe.[10] Figure 1.3 shows some examples of indoor and outdoor cybercafes in BJ, highlighting the cabling and the type of Internet access they rely on. In case of power cut or when the quality was worse than usual, I had to travel to the capital city. Such constraints, of course, led to an inefficiency from my side during the process, as I spent more time than I should.

Further, the costs varied between US$0.5 and US$0.9 on average per hour [197]: in other words, browsing for 24 hours in a cybercafe was equivalent to spending between roughly 19 % and 38 % the average per-month income of a Beninese end-user as per the World Bank [311]. Given these realities (that are common in developing regions [212, 222, 258, 270]), Internet access is considered in my country as a luxury. Although being ranked as such a commodity, the QoS provided to those who can afford it is far from that enjoyed in the West. For example, I recall that during my studies in Université d'Abomey Calavi (UAC, BJ), the need for a better QoS than that experienced in the day often forced some students to stay until late in the evenings, periods during which few users were sharing the low overall bandwidth. The fortunate ones preferred times of power outages in the day since their entities were ones of the few powered by a generator.

Marked by such a life experience, I applied to IMDEA Networks Institute with the hope that researching in the field of internetworking will help reverse the trend. As my application was accepted, I came to Spain motivated by *"the need to quantify this experienced QoS for the whole region, to identify the weaknesses of the African Internet ecosystem, and to suggest suitable actions to be taken for addressing this situation."* Together with Dr. Pierre Francois, my first advisor at IMDEA Networks Institute, that I previously met in one of my classes at UAC, we decided to undertake research studies that aim to help improve the QoS perceived by end-users while assisting both local ISPs and end-users save on costs, first in West Africa in particular [80, 131, 132], and then in Africa [77, 81, 82, 89, 310]. Later on, the focus of this work was sharpened to Africa and expanded to developing regions in general under the supervision of Prof. Dr. Francisco Valera, my second advisor at Universidad Carlos III de Madrid (UC3M, Spain, ES) [78, 79, 85, 87, 133, 194].

[10] The term *cybercafe* refers to a business usually run by a private company, which consists in commercializing Internet access so as to allow anyone to connect upon the payment of a fee per hour.

(a) *Cybercafe relying on an Internet access served through ADSL*

(b) *Indoor cybercafe (ADSL connection): image highlighting the cabling towards/from the router*

(c) *Indoor cybercafe with a better cabling: the router and access point can also be distinguished*

(d) *Outdoor cybercafe with the access distributed via wireless Access points*

Figure 1.3: Indoor/outdoor cybercafes in Cotonou (Benin (BJ)), November 2016

Having experienced the contrast of studying with better Internet access and aware of the myriad opportunities that it could give to the poor, I did my best so that this work can become a decisive turning point in the quest for a better and affordable Internet for all. Altogether with my advisors and co-authors, we strongly hope it will contribute to change the lives of the billions of people that are disconnected or suffering from this lack of Internet access or good connectivity.

More technically, in January 2013, at the beginning of this project, much less was known about the interdomain routing in Africa or even the African Internet, although several research studies [102, 119, 182, 183, 281] have analyzed the whole interdomain routing in detail (Chapter 2). We found very few previous measurement projects with African focus [106, 222, 223]. This lack of data on IP networks, known to be critical for ISPs and telecoms operators, paved the way for an inefficient routing and traffic engineering (*e.g.*, Figure 1.4). Meanwhile, the QoS of the Internet access is not worth the price. On the one hand, coastal countries are characterized by an underutilization of the optical fiber, a low Internet access or QoS, and an unfavorable ratio QoS/price. On the other hand, inland countries, whose data often traverse coastal countries, experience a deficient Internet access with a detrimental QoS/price ratio. Besides, most rural areas are not yet reached by technology, ISP infrastructures, or power. Despite the rapid penetration of mobile devices and the increase of mobile broadband access accounts, 84 % of Africa's inhabitants could not access the Internet [135, 207, 270]. We believe this statistic represents an opportunity, which may be seized by not only African network operators but also companies (around the world) planning to offer telecoms services in the region to help raise the bar. But a fundamental question is which parts of the network or the continent are favorable to such investments? Due to

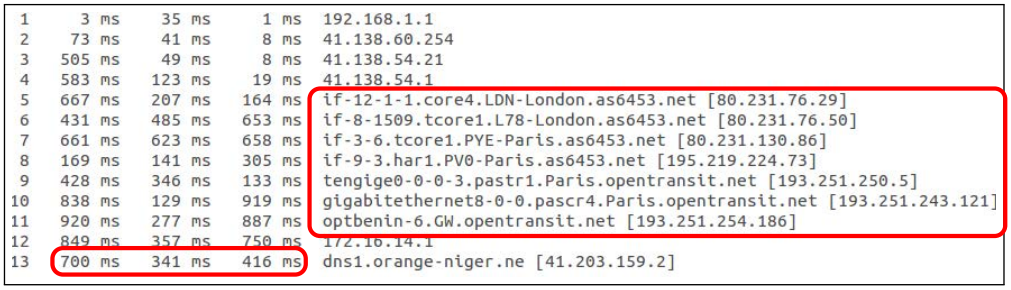

(a) *Traceroute performed from an end-user of AS37385 (SONITEL, NE) towards the DNS of AS37233 (ORANGE-NIGER, NE)*

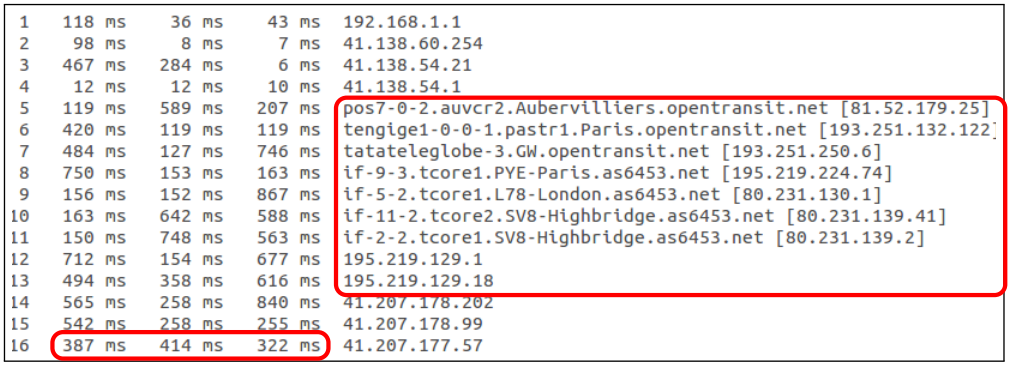

(b) *Traceroute run from an end-user of SONITEL towards a public IP of AS24691 (TOGO TELECOM, TG)*

Figure 1.4: Traceroute between adjacent ISPs in the same country (Niger (NE)) and in the same sub-region (*West Africa* (*WAf*)) on July 17, 2013

the lack of studies on the issues that are undermining the African Internet ecosystem, few insights into this topic were also available to either the industry or the research communities.

Based on ITU and Akamai estimates [15, 16] of the average connection speed per country, the project "Internet Speed in Africa" [148] estimates the average Internet speed on the continent to range from 694 kbps in Mali (ML) to 11,299 kbps in Rwanda (RW). This open data initiative [148] found it to be on average 2,439.3 kbps. This percentage can be broken down into 3,635.6 kbps in *EAf*, 3,052.6 kbps in *SAf*, 2,550.7 kbps in *NAf*, 2,102 kbps in *WAf*, and 1,768.6 kbps in *CAf*. Meanwhile, the Internet penetration rate in the region can be split into 36.3 % in *NAf*, 26.3 % in *SAf*, 25.7 % in *EAf*, 16.7 % in *WAf*, and 14 % in *CAf*.

A follow-up question is about the cost at which the privileged 16 % of Africa's inhabitants accesses the Internet. Oyelaran-Oyeyinka *et al.* [212] inspected in 2004 the pattern of Internet adoption in Kenya (KE) and Nigeria (NG), showing that Internet usage is constrained by both structural and cost factors. They involved empirical data notably collected through the use of questionnaires and in-depth interviews of over 200 academics in 10 universities of KE and NG. Next, Pejovic *et al.* [222] identified in 2012 the high cost of Internet access as one of the significant barriers to further Internet penetration in rural Sub-Saharan Africa. They gave as an example the fact that satellite access in Macha cost US$1,200 per month, while the average monthly income was about US$30. Similarly, Akue [17] noticed a year later that in the region, a mobile subscription cost up to seven times more than that of a telephone subscription. He then highlighted,

using tariffs for access to wired broadband applied in different countries of Sub-Saharan Africa, that Internet connection is expensive in both absolute and relative terms. Indeed, the ten lowest tariffs per month for Internet access ranged from US$26.3 in South Africa (ZA) to US$58.2 in ML, equivalent to a relatively large differential of US$32.1 between the prices in both countries. By contrast, the ten highest tariffs, according to his dataset, ranged from US$170 in Uganda (UG) to over US$1,000 in Burkina Faso (BF). These tariffs were all applied to leased lines intended mainly for businesses and were considered very high. Akue [17] also identified as reasons for the high cost of Internet connection, the inadequate investment in telecoms, the unfavorable economic market conditions, the lack of competitions, and the international Internet connection costs.

Along these lines, a study conducted in 2015 by the Internet Society (ISOC) [207] concluded that Internet access in Africa could cost 30 or 40 times more than in developed countries. [207] found, for instance, that 60.2 % of Ethiopia's GDP per capita is required for broadband access, while it is 31 % in Uganda (UG) and 7.4 % in Sudan (SD). Meanwhile, 15.7 % of KE's average GDP per capita is required for the same purposes, compared to 6.1 % in ZA and less than 2 % in most of Europe. [207] also remarked that while the increase in undersea cables [198, 264, 265] has reduced international transit costs, prices remain significantly higher than those for developed countries. Moreover, an African Union-supported study from 2008 showed that Africa spends per year between US$400 and US$600 million in transit fees for intra-African traffic, which gets routed through expensive transit links [5], and these costs have been increasing over time [141, 227]. A preliminary study [156] then reported in 2012 that IXPs would help improve QoS for local traffic and reduce these expensive transit fees paid by those developing countries. Initiatives such as the AXIS project were thus launched by the ISOC to promote the creation of local IXPs and enable cross-border interconnection [6].

In a nutshell, the quality of Internet access in Africa is an obstacle to its development. It is clear that without investigating connectivity between African networks and from African networks to Content Providers (CPs), we cannot find where this situation can be improved. These facts motivated us to inspect the state of African Internet, reveal its topology, and identify its weaknesses, especially the most critical ones. They also prompted us to observe, through extensive measurement studies, its evolution from 2013 to 2017 for shedding light on ISP habits that could be corrected or encouraged and suggesting ways to boost the Internet growth in the region.

1.2. Main contributions and organization of the thesis

1.2.1. Objectives

In the light of recent political moves supposed to push for African ISPs interconnection and IXPs establishment, the main objective of this thesis is to observe, periodically perform measurements, regularly make a status check of the QoS over the years, as well as analyze the effects on the African Internet ecosystem of the deployments of new interconnection facilities and peering links among local ISPs. We expect this work will have a positive impact in the region, in

particular:

1. Informing development organizations and policymakers on gaps and state of interconnection (*for the Internet community*).
2. Investigating the causes of congestion in the African IXP substrate (*for researchers and network operators*).
3. Supporting business investment decisions and opportunities in the region (*for Internet business development and CPs*).
4. Improving the QoS in data and voice transmissions (*for the Internet community*). Said measurement study also aims to archive reliable data on the Internet evolution and provide in real time *network operators* with supporting information for inferring profitable decisions about IXPs establishment and peering through IXPs.
5. Minimizing transmission delays and costs paid to transit operators, namely at two levels: intra-African communications as well as access to CPs (*for researchers and network operators*).
6. Boosting the integration of web services in local customs (*for the Internet community).*
7. Helping *researchers* undertake interconnection studies by making our measurement datasets available and freely accessible.
8. Suggesting a realistic framework (based on an optimal utilization of the existing – submarines and terrestrial – optical fibers and local IXPs and considering external factors influencing the Internet in the region) for improving the Internet infrastructure as well as opening the way for a better traffic localization (*for the Internet community*).

We believe this radiography of the African Internet ecosystem will serve as a case study for developing countries to build methods for better understanding and revealing the reasons behind their poor Internet penetration, their high access costs, or their poor QoS. In addition, it may help achieve efficient routing and reverse the trend to benefit from the full potential of the Internet. Further, we hope that thanks to this thesis, a traceroute from Africa to Africa would not be going through any other continent shortly.

1.2.2. Elaborate research problem and process of the study

Science is extremely driven by curiosity. Added to that, we were motivated by the context depicted in Section 1.1. However, although being necessary, these are not enough to achieve the above-listed objectives (Section 1.2.1). An essential task is to translate these goals into an elaborate research problem with more detailed sub-goals and a well-defined plan. I did so over the last five years (January 2013–December 2017), under the guidance of my advisors, while

working closely with the industry, as explained further. Figure 1.5 depicts, under the format of a Gantt chart, the planning of the project. On that figure, our work is split into items numbered from 1 to 10 and colored given their category. We distinguish five categories: *(i)* Related work analysis, teaching activity, and dissemination (blue); *(ii)* Active measurements in the African interdomain routing (pink); *(ii)* Passive measurements in the African interdomain routing (tan); *(iii)* Measuring the African web ecosystem (umber); and *(iv)* Topology and infrastructure (brown). We explain Figure 1.5 in the subsequent paragraphs and refer to it throughout the remainder of this thesis.

In January 2013, we began by studying the previous work related to the interdomain routing and specifically to the African interdomain routing (Item 1 of Figure 1.5). Such an analysis was then performed before our exploration of the causes of congestion in the IXP substrate, route-collectors data, the web ecosystem, and the topology and infrastructure for the region. Note, we took into account results published during the period of the project in our discussions, to have a related work up-to-date (Chapter 2). This task, therefore, lasted until September 2017.

Since we can not achieve our objectives without accessing some of the local ISPs technical data, and because the measurement infrastructure was quasi-inexistent in the region [54], we started by building it (Item 2). While traveling around the continent, actively leading AXIS capacity building workshops, performing teaching activities, and meeting or discussing with stakeholders (Item 3), we tried to set up our own raspi-based infrastructure in vain. Nevertheless, we successfully deployed Réseaux IP Européens (RIPE) Atlas probes [80, 248] and later, Center for Applied Internet Data Analysis (CAIDA)'s Ark probes [40] within African networks. These activities, which we present in Chapter 3, lasted until August 2016. It is worth mentioning that the deployed probes are still within the host networks and can be used by future measurement projects.

With our early measurements and obtained analysis results, we could characterize [77] the state of the interconnection in Africa (Item 4). We proved by analyzing data gathered with both deployed and existing RIPE Atlas probes that local ASes mostly rely on intercontinental ASes for intra-African communications. We also highlighted the first impacts of efforts to set up local IXPs [81]. Note, this work has been conducted jointly with the RIPE NCC.

In collaboration with not only the RIPE NCC but also members of other Internet development institutions (ISOC and Packet Clearing House (PCH)), we then continued to observe the interdomain routing. We performed from 2013 to 2016 several targeted full-mesh measurements among the probes in Africa, in the US, and in Europe (Item 4), while helping enlarge the platform in the African region and disseminating our first analysis results (Item 10). Our longitudinal study [85] (*cf.* Chapter 3), however, revealed the remaining dominance of ASes based outside Africa to provide intra-African communications. We showed that AS paths are significantly longer when considering communications among probes hosted in African networks than those among probes in the US or European networks. Further, we underlined the sustained efforts made by local networks to deploy IXPs for traffic localization especially in some countries. Next, we evaluated the impact of IXP setups on the delay among local ASes, finding that the setup of new IXPs

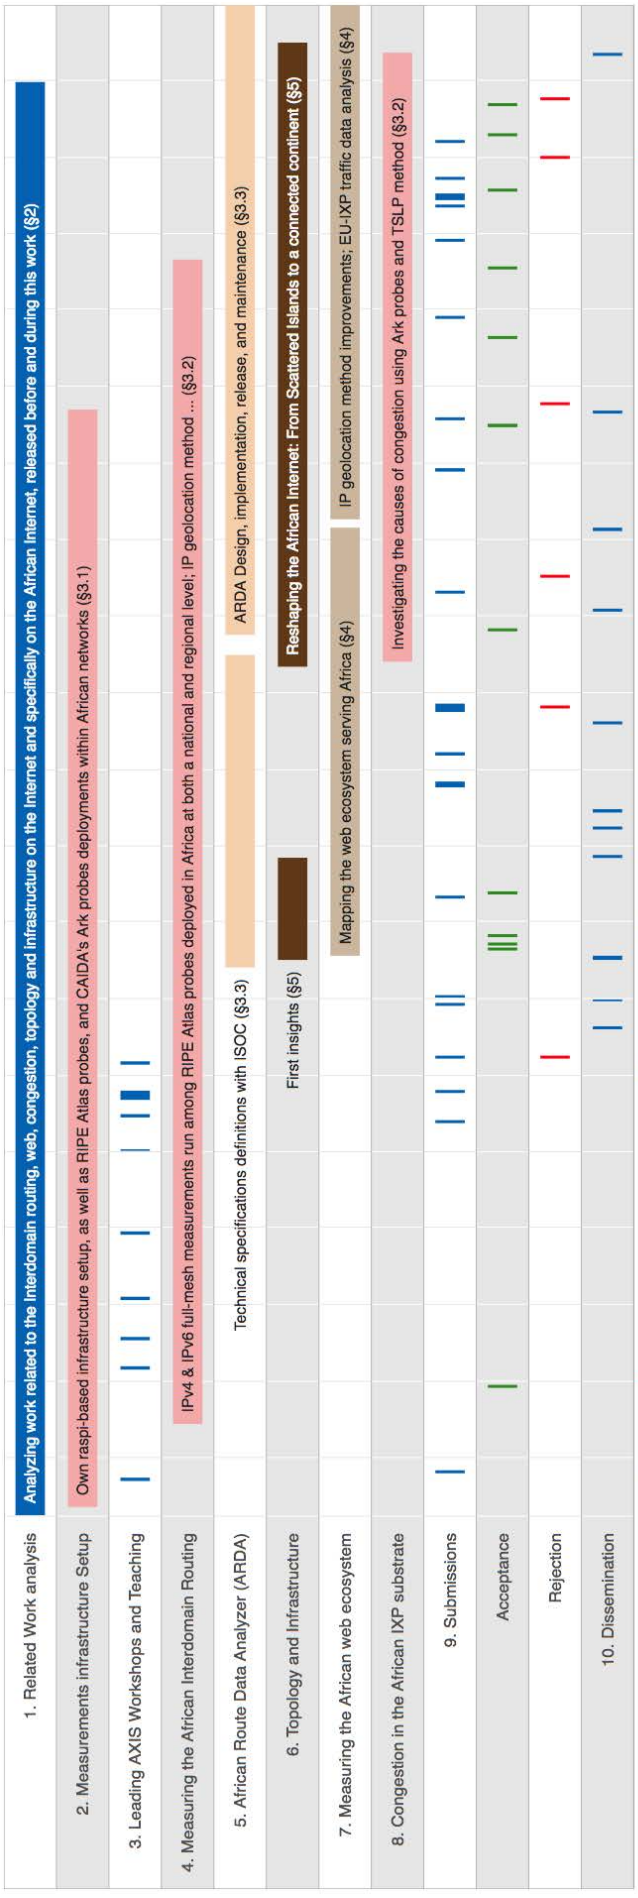

Figure 1.5: Gantt chart of the project covering the period January 2013 to December 2017.

contributed to reducing the latency among adjacent ISPs that agreed to peer with one another.

In this part of our thesis, our methodology has consisted of targeting, like Spring *et al.* [176], a restricted set of ISPs (*i.e.*, African networks) with extensive measurements. We combined data gathered from 6 [77, 81] and 10 [85] public Data Sources (DSes) with latency-based measurements to geolocate discovered IPs and infer the country paths traversed by our traceroute outputs. We then characterized the lengths of AS-paths conveying intra-African communications and compared them to those of AS paths among US and European networks. We inspected techno-economic insights on the routing trends and the impact of transit localization on end-to-end delay. Contrarily to Augustin *et al.* [25], we could map local IXPs in our dataset, thanks to our adopted IXP mapping methodologies and the visibility given by our VPs. Finally, we exhibited in several case studies the impacts of an IXP launch on both AS paths and end-to-end delays among African networks. Each of those findings represents a particular contribution to the related work.

An outcome of these results is the need for supporting the growth of existing or newly launched IXPs, to give them the chance to grow in the competing environment surrounding them (Item 5). In November 2014, we began by defining, with the ISOC, the technical specifications of a system that will play such a role. A year later, the ISOC then partnered with Universidad Carlos III de Madrid (UC3M) on a project whose culmination point is the ARDA, a compass to support peering growth (in this case applied to AfriNIC). We emphasize that it has been built with in mind its possible extension to other Internet regions. Altogether, we designed and developed ARDA, an open source system that collects and analyzes from various angles the data of BGP route-collectors located in Africa to release statistics useful for researchers, Internet developmental institutions, network operators, and the community [87].

Internet users often connect to the Internet to not only communicate but also access information (web content), which are served by *content providers* (CPs). By *CPs*, we refer to (web-based) service providers which provide content (text, videos, websites, etc.) to end-users. We include traditional Content Delivery Networks (CDNs) such as Akamai that serve third-party content, as well as content providers such as Google, Facebook, and Netflix, which build and operate their own extensive networks. CPs tend to deploy their cache servers within local ISPs, at IXP infrastructures, or within their own network [99] to be as close as possible to end-users.

Although global web infrastructures have been inspected recently, prominent works [37, 76, 117, 123, 157, 210, 276, 315] have not (*i*) focused on developing regions such as Africa or (*ii*) explored if worldwide results apply to these under-connected areas. We, therefore, investigated from 2015 to 2017 (*cf.* Figure 1.5) content delivery to African networks (Item 7) [89] in collaboration with researchers from Queen Mary University of London, (QMUL, United Kingdom, UK) and University of Cambridge (UK). After refining the geolocation methodology mentioned earlier, we employed several measurement methodologies aiming at analyzing traffic data, traceroutes/RTT measurement outputs, as well as DNS queries outputs and HTTP requests outputs. We found Google to have an expanded content infrastructure in Africa compared to other CPs. Nevertheless, much web content is still served from the US and Europe. We discovered that many of the

problems faced are due to the persistent lack of peering and poor DNS configurations. We then mapped the infrastructure of the top 10 global and the top 15 regional websites to show that large-scale web infrastructure deployments are a rarity in Africa and that top regional content sources mostly host their services outside Africa. We present our results in Chapter 4.

By and large, our findings highlight the need for more peering in the region, while new IXPs are being set up. But before inspecting ways to achieve this, we investigated the causes of congestion in the African IXP substrate, as the prevalence of congestion at local IXPs may prevent these facilities from growing. We found no previous research to explore the nature of congestion and its causes at local IXPs, despite the interest in quantifying performance at those facilities raised by the context of promoting IXPs [6, 81, 85, 156]. The TSLP technique has recently been developed and validated by Luckie *et al.* [167] to infer congestion cases from RTT measurements. We applied it on fine-grained measurements collected over a whole year by Ark probes deployed at six strategically selected local IXPs [292] (Item 8). Next, we verified the events and interviewed the IXP operators to identify their causes and check whether these corroborate our expectations based on the collected datasets. We then examined to which extent the existence of congestion negatively influences communications between a given AS and its neighbor. We detail this work and present our findings in Chapter 3.

The need for better infrastructures in the region is another side of the problem which has been discussed in [61]. As an option to enrich connectivity on the continent and incentivize CPs to deploy caches closer to end-users, we suggested a framework based on existing physical fiber and local IXPs, which notably takes into account socio-economical and geo-political factors to build a distributed IXP spanning the African continent (Item 6). Compared to previous authors who thought about IXP interconnection as a way to attaining these goals [71, 72, 203, 211, 268, 275], the particularity of this study is to have succeeded in providing a concrete proposal for achieving it and a quantitative estimation of potential benefits from doing so. We showed using measurement data, simulation, and analysis, how such an initiative will help reduce both RTTs among African networks and AS path lengths, thus lowering the costs paid to their respective transit operators (Chapter 5). Note, the latter two studies have been launched in collaboration with CAIDA from November 2015, at the beginning of my five-month internship at their premises and conducted simultaneously till early November 2017.

1.2.3. Thesis structure and writing style

This doctoral thesis reports on contributions achieved from January 2013 to December 2017, whose timelines are presented in Figure 1.5. It is structured as follows: we discuss work related to each addressed topic in Chapter 2. We then explain in Chapter 3 how we built our measurement infrastructure before detailing the analysis of data collected through both active and passive measurements of the African interdomain routing as well as our obtained results and their implications. Next, we present the measurements and analysis aiming at mapping the African web ecosystem in Chapter 4 before highlighting our findings. After that, we suggest in Chapter 5 ways

to tackle the infrastructure gaps in the region, including a framework for building a distributed IXP in Africa. Finally, we conclude and outline our perspectives for the future in Chapter 6.

As for our writing style, it is worth emphasizing that in the core of this thesis, we present the related work, our probe deployments or teaching efforts, measurements, experiments, data collection, and interviews in the past tense. Meanwhile, our analysis of the collected datasets is reported in the present tense to highlight that these methods are reproducible on similar datasets. Finally, the flow of the thesis, as well as obtained results are also reported in the present tense, since reproducing the same analysis on our previously collected datasets will give the same results. We acknowledge that our results are related to the period during which measurements were performed, since the Internet is in constant evolution. However, by making these choices, we aim at drawing the attention of the reader to the chronology of the closely intertwined events, which are presented and to how they lead to the conclusions of this thesis.

1.3. Summary of publications

The core of this thesis is based on two journal papers [79,85], four conference papers [78,81, 89,185], one poster [83], and several invited talks mainly at network operator meetings. Next, we contributed to the report from Dagstuhl Seminar 14471 [61]. In addition, two more journal papers are under submission. Further, other publications such as a RIPE Atlas labs article, our web technical reports containing released measurement datasets, and two applications arose from our studies as well. In this section, while elaborating on these publications (ordered by publication dates and classified by category), we report, for each of them, the number of citations as of this writing (September 2017) between square brackets if applicable and briefly describe how we contribute to each of them.

1.3.1. Published journal articles

Both our published journal papers are JCR journal of impact factor 3.34.

1. **Rodérick Fanou**, Francisco Valera, Pierre Francois, and Amogh Dhamdhere. Reshaping the African Internet: From Scattered Islands to a Connected Continent. *Computer Communications*. Elsevier. November 2017.

   This work [79], which constitutes the core of Chapter 5, was launched in collaboration with CAIDA in the context of my internship in their premises. It was mainly conducted under the supervision of Dr. Amogh Dhamdhere during my internship, and that of both Dr. Amogh Dhamdhere and Prof. Francisco Valera from the end of my internship till its publication. Considering publicly available information on fiber deployments, etc., and inspired by our previous results as well as discussions with local network operators and Dr. Pierre Francois during our study of the interdomain routing in Africa [81], I suggested a rough idea for a framework for building a distributed IXP in the region. The main lines of the

proposed methodology were retained and improved after extensive discussions among all authors. The resulting approach was then presented to the CAIDA team for their feedback and expertise on January 7, 2016. I was then in charge of surveying of local IXP operators, contacting cable operators or the ISOC to notably get their opinions on IXP interconnection and on the environment in the region as well as collecting the datasets essential for the study. Based on these inputs, I suggested and implemented the utilized algorithms, ran the computations and simulations for obtaining the results before editing the first drafts that all authors kept improving later. The final version of the paper was obtained after several edits and full reviews by each author, including the comments of the reviewers from the different venues to which we submitted.

2. **Rodérick Fanou**, Pierre Francois, Emile Aben, Michuki Mwangi, Nishal Goburdhan, and Francisco Valera. Four Years Tracking Unrevealed Topological Changes in the African Interdomain. *Computer Communications*. Elsevier. July 2017 [2 citations].

   This paper [85], which continues the works [77,81,83] is included in the Section 3.2.1 of the thesis. It has been co-authored with the principal actors with whom we have been working from 2013 to 2016 on this topic. Among them, Emile Aben and Dr. Pierre Francois were from the industry when the paper was published; In the meantime, Michuki Mwangi and Nishal Goburdhan were members of Internet developmental institutions and the remainder from the research community. I was responsible for updating the related work of [81], continuing the RIPE Atlas probe deployments in the region (along with other co-authors), keeping the various measurements (whose datasets were studied) running till 2016 (except traceroutes performed among US and Europeans networks), and analyzing all the collected datasets. I was also in charge of defining the IXP mapping methods and implementing them for IXP detection, identifying more case studies of the impacts of peering on communications performance and highlighting the emergence of recently established IXPs. Further, I suggested the evaluation of inter-ISP communications performance within the US as well as European, and African countries. Following that, I inspected how frequently an AS path, which originated from and was destined to African countries, European countries, or the US, traverses an IXP located on each continent. Several edits and full reviews of the first draft were performed by each author, taking into consideration the comments of the reviewers from the different venues to which we submitted. This joint-work allowed us to obtain a genuine symbiosis of both the research knowledge and the industry know-how.

1.3.2. Conference or workshop papers

Our conferences and workshop papers can be listed as follows:

1. **Rodérick Fanou**, Francisco Valera, and Amogh Dhamdhere. Investigating the Causes of Congestion on the African IXP Substrate. In *ACM Internet Measurement Confer-*

ence (IMC), London, UK, November 2017. Computing Research & Education (CORE) 2017 [57] rank A.

Our methodology, analysis, and findings in this work [78] are unveiled in Section 3.2.2. This paper results from the second study that we conducted in collaboration with CAIDA from the beginning of our internship (November 2015) until November 2017. Before its launch, the TSLP technique was already developed, tested, and validated [167]. I, therefore, suggested to apply it in the African interdomain routing, an idea which was approved by my internship tutor Dr. Amogh Dhamdhere. I helped deploy 15 Ark probes within 14 African networks in 11 countries by identifying the potential probe hosts, contacting them, and getting their approval. Six of those VPs were later retained for [78], given that they have been strategically deployed at local IXPs. I was not responsible for implementing/running the border mapping process[11] or performing the TSLP measurements. Instead, I conducted the validation of the border mapping process outputs with the IXP operators, implemented a set of scripts to parse the collected time series and to perform the loss measurements, and performed an in-depth analysis of many different datasets with my co-authors to obtain our results. Altogether, we then carried out, either through emails or online meetings, several IXP operator interviews aiming at pinpointing the causes of inferred congestions cases. Further questions to network operators at any side of a problematic interdomain link were transmitted through the IXP operator. The final version came out after several iterations, edits, and full reviews by each author.

2. **Rodérick Fanou**, Gareth Tyson, Pierre Francois, and Arjuna Sathiaseelan. Pushing the Frontier: Exploring the African Web Ecosystem. In *The 25th International World Wide Web Conference (WWW 2016)*, April 2016, Montreal, Canada [11 citations]. CORE 2017 [57] rank A$\star$.

   From our discussions with Dr. Gareth Tyson and Dr. Arjuna Sathiaseelan at the Dagstuhl Seminar 14471 [61] emerged a working plan, which later evolved with the guidance of Dr. Pierre Francois towards the inspection of the web ecosystem in Africa (*cf.* Chapter 4). At every step of this work, I set up several meetings with Dr. Gareth Tyson to align our views on the methodology. Using both my deployed RIPE Atlas probes and those existing in the region, I performed the DNS queries, HTTP queries, traceroutes, RTTs measurements whose outputs were used in this paper. I was not in charge of performing the EDNS0 queries. Instead, I was responsible for the implementation and execution of the computation scripts, the graphs edition, etc. Together with my co-authors, I then analyzed the obtained results, and after several edits and reviews performed by all of them, we obtained the final version of this paper. The results presented in Chapter 4 represent, however, an improved version of those released in this paper [89]. In the manuscript that continues this work, we

[11] The border mapping process aims at obtaining sufficient information about the links observed from the AS hosting the VP toward every other AS to constrain the border router inferences [43, 168, 262]

indeed refined the Internet Protocol (IP) geolocation methodology and studied traffic data collected at a large European IXP, which is currently under submission at a top-tier journal.

3. **Rodérick Fanou**, Pierre Francois, and Emile Aben. On the Interdomain Topology of Africa (Invited Poster). In *The 5th PhD School on Traffic Monitoring and Analysis (TMA) and The 7th International Workshop on TMA*, April 2015, Barcelona, Spain.

4. **Rodérick Fanou**, Pierre Francois, and Emile Aben. On the Diversity of Interdomain Routing in Africa. In *The 16th International Conference on Passive and Active Measurement (PAM 2015)*, March 2015, New York City, NY, USA [25 citations]. CORE 2017 [57] rank B.

   We started this work [81] by helping build the measurement infrastructure essential for our studies, as detailed in Section 3.1: I traveled around the African region, deploying both our raspberry pis (raspis) and RIPE Atlas probes. As the RIPE Atlas probe deployments were successful, I could then perform full-mesh measurements among those devices. With the guidance of Dr. Pierre Francois and Emile Aben, we set up a sound methodology aiming at treating carefully the traceroute data collected during the measurement campaigns. Altogether, we defined an IP geolocation based on the combination of public datasets and latency-based measurements. I then contributed to analyzing path dynamics, AS paths length distribution, techno-economic insights in the interdomain routing, and impacts of traffic localization on end-to-end delay, before looking for insights into the emergence of IXPs. This paper was extracted from my master thesis [77], a study that we later deepened over the Ph.D. program since the results were unique and promising.

5. Miriam Marciel, Foivos Michelinakis, **Rodérick Fanou**, and Pedro Jose Muñoz-Merino. Enhancements to Google Course Builder: Assessments Visualisation, YouTube Events Collector and Dummy Data Generator. In *XV Simposio Internacional de Tecnologías de la Información y las Comunicaciones en la Educación (SINTICE 2013)*, September 2013, Madrid, Spain [5 citations].

   Although being part of our publications, this work [185] is not included in the thesis. It aimed at extending the functionalities of Google Course Builder (GCB), an open source platform that provides online educational courses to a broad public, for improving how it supports learning analytics. As a matter of fact, in platforms like GCB, it is imperative to understand the learning process and try to improve it. Together with my co-authors, I defined the proposed architecture including the external elements of GCB (YouTube events collector, YouTube API, dummy students generator, and visualization module). I was then in charge of building the Visualisation module, testing it, and editing the section entitled "Assessments Extensions for learning Analytics on GCB." These required the collection of data on each student and the definition of some metrics whose processing (with the classes of GCB) allows the computation of each student's assessment statistics. Finally,

I implemented the data visualization and recommendations component, which helps the teacher know at a glance how his students are performing. The complete code is freely available at [186].

1.3.2.1. Technical reports

We enumerate, in this section, our contributions to the report [61] from Dagstuhl Seminar 14471 at which 27 participants were invited. Only personal contributions to this report, which are related to the infrastructure needed for a better Internet access in Africa and the challenges in that developing region, are included in Chapter 5.

1. **Rodérick Fanou**. Which Infrastructure for a better Internet in Africa? In *Report from Dagstuhl Seminar 14471*: Jon Crowcroft, Adam Wolisz, and Arjuna Sathiaseelan, Towards an Affordable Internet Access for Everyone: The Quest for Enabling Universal Service Commitment, November 2014, Dagstuhl, Germany [4 citations].

2. Weverton Cordeiro and **Rodérick Fanou**. Challenges in developing regions. In *Report from Dagstuhl Seminar 14471*: Jon Crowcroft, Adam Wolisz and Arjuna Sathiaseelan. Towards an Affordable Internet Access for Everyone: The Quest for Enabling Universal Service Commitment, November 2014, Dagstuhl, Germany [4 citations].

3. **Rodérick Fanou**, Michael Fourman, Thomas Huhn, Renato Lo Cigno, Leonardo Maccari, Mahesh Marina, Henning Schulzrinne, and Marco Zennaro. Socio-Economic Models and Role of Community Networks. In *Report from Dagstuhl Seminar 14471*: Jon Crowcroft, Adam Wolisz and Arjuna Sathiaseelan, Towards an Affordable Internet Access for Everyone: The Quest for Enabling Universal Service Commitment, November 2014, Dagstuhl, Germany [4 citations].

1.3.3. Journal articles under submission

We detail, in this section, our contributions that are currently under review at top-tier journals.

1. **Rodérick Fanou**, Gareth Tyson, Eder Leao Fernandes, Pierre Francois, Francisco Valera, and Arjuna Sathiaseelan. Exploring and Analysing the African Web Ecosystem. Under submission.[12]

   This submission, which enriches our previous work [89], constitutes the core of Chapter 4. It details a large-scale measurement study of the web infrastructure serving end-users in Africa, which we undertook over the last two years. Among others, we have employed several methodologies that have collected a broad range of relevant data. Via each methodology, we have been able to explore the deployment strategies of multiple websites for serving

[12] This submission has been accepted with minor revisions at the ACM Transactions on the Web on December 06, 2017. This JCR journal has an impact factor of 1.526.

the region and identify core issues and bottlenecks. We have made substantial extensions and improvements to [89], including a significantly improved geolocation technique, which addresses a number of limitations in our previous work. The key enhancements were the inclusion of the speed of light checks to determine IP geolocation discrepancies and my addition of the multilateration technique to build a rigorous four-step geolocation approach. Towards this end, I repeatedly run new latency measurements from RIPE Atlas probes randomly selected worldwide towards the geolocated Google caches (GGC) and DNS resolver IPs. A large IXP dataset was then analyzed to quantify and understand the total traffic that fails to be localized in Africa: I conducted this specific task with Eder Leao Fernandes (Ph.D. Student at QMUL, UK), under the guidance of Dr. Gareth Tyson. After that, I was in charge of re-running all the computations scripts, necessary for updating our results with the outputs of the refined geolocation methodology. Several recently published references have then been added for the completeness of the manuscript. Finally, we performed a substantial textual and structural edits altogether, to improve the readability of the paper, combined with several reviews before its submission.

2. **Rodérick Fanou**, Víctor Sánchez-Agüero, Francisco Valera, Michuki Mwangi, and Jane Coffin. The ISOC Compass to Support Peering Growth in the African Region: a Route-collectors Data Analyzer. Under submission.

   This journal article under review presents the design and implementation of a route-collectors data analyzer. The rough idea behind this work, which has improved over time, is that of the ISOC that was in need of a system able to help the Internet community witness the evolution of the IXPs of an Internet region while supporting the growth of those infrastructures. As there is a great push for IXPs setup in the African region [81, 85, 156, 292], the prototype of this tool has been implemented for the specific case of route-collectors located in AfriNIC and is thus entitled African Route-collectors Data Analyzer (ARDA). To achieve this, I have worked with Víctor Sánchez-Agüero (Ph.D. Student at IMDEA Networks Institute and UC3M) under the supervision of Prof. Francisco Valera and in collaboration with the ISOC. The architecture of this open-source system can be split into three modules: the data collection, the metrics computations, and the visualizations modules. First, I contributed to the design of the system according to the technical specifications, which were jointly defined with the ISOC. Together, we then implemented the data collection module. Next, I implemented the computation module and the frequent transmission of the computed statistics to the visualizations module. Further, I contributed to the collection of BGP routing data and guided the design of the visualization module. Altogether, we analyzed the obtained results. The final manuscript was then edited and reviewed several times by each author before its submission.

1.3.4. Other contributions

The following contributions without printed publications have mostly aimed at triggering the interest of network operators, of other researchers, and of the Internet community to our research.

1.3.4.1. Invited talks at network operators meetings and seminars

I gave several talks, mostly at network operators meetings, which can be listed as follows:

1. **Rodérick Fanou**, Víctor Sánchez-Agüero, Francisco Valera, Michuki Mwangi, and Jane Coffin. African Route Collectors Data Analyzer: a compass to support peering growth in the region (Presentation). In *The 7th African Peering and Interconnection Forum (AfPIF 2016)*, August – September, 2016, Dar es Salaam, Tanzania.

2. **Rodérick Fanou**, Gareth Tyson, Pierre Francois, and Arjuna Sathiaseelan. Pushing the Frontier: Exploring the African Web Ecosystem (Presentation). In *The 7th African Peering and Interconnection Forum (AfPIF 2016)*, August – September, 2016, Dar es Salaam, Tanzania.

3. Cristina Márquez, **Rodérick Fanou**, Pierre Francois, and Michuki Mwangi. Assessing peering evolution in Africa (Remote presentation). In *The 6th African Peering and Interconnection Forum (AfPIF 2015)*, August 2015, Maputo, Mozambique.

4. **Rodérick Fanou**, Pierre Francois, Emile Aben, Michuki Mwangi, Nishal Goburdhan, and Víctor Sánchez. Tracking the evolution of intra-African traffic localization (Remote presentation). In *The 6th African Peering and Interconnection Forum (AfPIF 2015)*, August 2015, Maputo, Mozambique.

5. **Rodérick Fanou**, Pierre Francois, and Emile Aben. On the Diversity of Interdomain Routing in Africa (Presentation). In *RIPE 70 Meeting*, May 2015, Amsterdam, the Netherlands.

6. **Rodérick Fanou**. Which Infrastructure for a better Internet in Africa? In *Dagstuhl Seminar 14471*: Towards an Affordable Internet Access for Everyone: The Quest for Enabling Universal Service Commitment, November 2014, Dagstuhl, Germany.

7. **Rodérick Fanou**, Pierre Francois, and Emile Aben. From Africa to Africa: AS-level topology snapshot. In *The 5th African Peering and Interconnection Forum (AfPIF 2014)*, August 2014, Dakar, Senegal.

1.3.4.2. Web technical reports

We enumerate in this section our RIPE Labs article and our web technical reports that mostly contain released measurement datasets used in our journal articles and conference papers. In addition to improving the reproducibility of our different works, publishing our datasets aims at encouraging further experimentation and proposals from other researchers in this area.

1. **Rodérick Fanou**, Francisco Valera, Pierre Francois, and Amogh Dhamdhere. Reshaping the African Internet: From Scattered Islands to a Connected Continent (Technical Report). https://fourier.networks.imdea.org/external/techrep_reshaping/index, released in September 2017.

2. **Rodérick Fanou**, Pierre Francois, Emile Aben, Michuki Mwangi, Nishal Goburdhan, and Francisco Valera. Four Years Tracking Unrevealed Topological Changes in the African Interdomain: Technical Report. https://fourier.networks.imdea.org/external/techrep_amc_journal/index/, released in June 2016, updated in June 2017.

3. **Rodérick Fanou**, Gareth Tyson, Pierre Francois, and Arjuna Sathiaseelan. Technical Report: African Content Measurement Campaign. https://fourier.networks.imdea.org/external/techrep_cdma/index/, released in June 2015, updated in June 2017.

4. **Rodérick Fanou**, Pierre Francois, and Emile Aben. On the Diversity of Interdomain Routing in Africa (RIPE Labs). https://labs.ripe.net/Members/fanou_roderick/on-the-diversity-of-interdomain-routing-in-africa, May 2015 [2 citations].

5. **Rodérick Fanou**, Pierre Francois, and Emile Aben. African Measurement Campaigns: Technical Report. https://fourier.networks.imdea.org/external/techrep_amc/index/, released in September 2014, updated in June 2017.

1.3.4.3. Applications

Two applications targeting decision-makers, network/IXP operators, researchers, and the Internet community resulted from this work:

1. Víctor Sánchez-Agüero, **Rodérick Fanou**, Pierre Francois, and Francisco Valera. African Measurement Campaigns (AMC), http://amc.netcom.it.uc3m.es/. October 2017.

2. **Rodérick Fanou**, Víctor Sánchez-Agüero, Francisco Valera, Michuki Mwangi, and Jane Coffin. African Route-collector Data Analyzer (ARDA), https://arda.af-ix.net/, April 2017.

1.3.4.4. Press releases

1. Michuki Mwangi and **Rodérick Fanou**. ARDA 1.0: A pulse meter for Africa's peering and interconnection landscape. https://www.internetsociety.org/blog/2017/04/arda-1-0-a-pulse-meter-for-africas-peering-and-interconnection-landscape/, April 2017.

2. **Rodérick Fanou** and Pierre Francois. Drawing the Map of the West African Internet. http://netcom.it.uc3m.es/whats-new/news/2014/drawing-map-west-african-internet, February 2014.

Chapter 2

Related Work

In this chapter, we discuss prominent work on Internet topology discovery, measuring performance on communications among local IP networks, congestion, routing data analysis, content delivery, as well as IXP interconnection. While doing so, we shed light on *(i)* the technical state-of-the-art in general and on *(ii)* the works related solely to the African Internet. We contrast both of them with the methodologies adopted in this thesis. Finally, we present a taxonomy of the studies related to the African Internet, separating those published before, from those published during this work.

2.1. Interdomain routing

2.1.1. Internet topology discovery and end-to-end performance measurements

Over the last decades, several platforms have been launched for constantly measuring the Internet, observing its evolution, and understanding its main characteristics. One of the pioneers, CAIDA, has a long history of running Internet measurement platforms. Its latest active measurement platform, Archipelago (Ark), aims at reducing the efforts needed to develop and deploy sophisticated large-scale measurements [40]. Other measurement networks have then been set up for similar purposes, sometimes with different scopes or extended capabilities. As an example, the PingER project [223] aims at measuring Internet end-to-end performance and was notably used to quantify the digital divide. Its infrastructure contained 89 monitors and 1,090 remote monitored nodes at 956 sites in 169 countries. We can also list as measurement platforms SamKnows [255], BISmark [32], Dasu [22], RIPE Atlas [248, 250], M-Lab [191], Planet Lab [224], etc. Bajpai *et al.* [26] provided a taxonomy of these measurement networks. They explored in detail their coverage, scale, lifetime, deployed metrics, as well as their measurement tools, architecture, and overall research impact.

A key use of those measurement networks is to help discover the Internet topology. In fact, Internet topology discovery, both at the router level and the AS level, is a topic that has been investigated extensively [102, 119, 126, 161, 175, 182, 183, 281]. In particular, Spring *et al.* [176]

used Rocketfuel to analyze Routeviews BGP table dumps combined with traceroutes performed by 750 VPs targeting 10 ISPs in the US. As for the IXP substrate, Augustin *et al.* [25] used in 2009 various techniques to map IXPs on the Internet. They detected 223 of the 278 IXPs with 393 known prefixes located all around the world. Note, they obtained their full list of (278) IXPs and IXP prefixes by merging IXP information collected from PCH and the Peering DB databases, IXP websites, private communications with IXP operators followed by extensive checks on the validity of this information.

Despite previous studies on Internet topology discovery and the existence of the aforementioned measurement platforms, little was known about the topology of the African Internet at the beginning of this project. In [106] for instance, Gilmore *et al.* mapped both the router and the AS-level graph of intra-African Internet paths. To achieve this, they performed traceroutes from a source in South Africa (ZA) towards several randomly selected IP addresses in all AfriNIC IP ranges for a week. Their results were then enhanced by AS adjacency data extracted from BGP-speaking routers hosted of the ZA Tertiary Education Network. It resulted in one-way paths from which a tree was inferred, with ZA at the root. They, however, acknowledged that the link density might look different if the traceroute probes were sent out from other countries in Africa. Similarly, the attempts of Augustin *et al.* [25] to infer IXPs in Africa were unfortunately often unsuccessful, which can be explained by the existence of only four looking glasses on the continent. Besides, African IXPs sometimes utilize RFC1918 address space, which may have prevented the use of various detection techniques. The authors acknowledged that they lack sufficient information to infer the presence of these IXPs that are known to exist and be active.

A key aspect of [176] is the targeted analysis of a restricted set of ISPs instead of an attempt to map the whole Internet. We follow the same focused approach in Section 3.2.1, targeting African ISPs: we undertook to fill the lack of knowledge of the African Internet by studying extensively the interdomain topology in Africa and measuring performance on communications among local networks. Bearing the above in mind and given the actual geographical dimensions of the continent [110, 159], we quickly understood that for achieving these goals, a larger deployed base of Vantage Points (VP) was needed. Of the 94 Archipelago monitors deployed, only five were in Africa, of which two were hosted in West Africa: Archipelago was therefore not used at this stage. Although the infrastructure of the PingER project [58, 223, 316] involves 46 African countries, only two (Burkina Faso, BF and South Africa, ZA) host a monitoring site, which prevented us from doing large scale end-to-end measurements. PingER was thus not retained, either. Other measurement platforms had similar characteristics, *i.e.*, very few probes in the African region [32, 54, 191, 224, 255]. The RIPE Atlas network, however, contained a dozen of devices in Africa in February 2013, and about 83, six months later [85, 248]. Consequently, we started by helping build the measurements network in the region: as detailed in Section 3.1, our actions were two-fold and consisted of deploying our own raspis-based measurements platform, while extending the RIPE Atlas network in the region. To study the African Internet with BGP data, similarly to Spring *et al.* [176] for a restricted set of US ISPs, obtaining relevant local BGP feeds

is also essential, and is part of this work (*cf.* Sections 2.1.3 and 3.3.1).

Early September 2013, Gupta *et al.* [117] performed, using BISmark nodes, traceroutes from access networks in Tunisia (TN), Kenya (KE), and ZA to sites hosting popular content to investigate Internet connectivity in Africa. They did not specify the period of these measurements in [117]. Their results, published in March 2014, revealed that 66.8 % of the paths going from their VPs towards Google cache servers located in Africa leave the continent. Since broadband access networks in those three countries are more developed [54] than in most of the 51 remaining African countries (Section 1.1.2), the results of this study may not reflect connectivity in other countries, as acknowledged by the authors. Next, Chavula *et al.* [49] examined communications among African research networks: they launched traceroutes from five Ark monitors located in residential/university networks to 95 university locations in 29 countries. These measurements were performed for 14 days (April 6 – 20, 2014). They found that 75 % of the paths were routed via Europe (EU) and the US and observed that RTTs on those circuitous paths were therefore affected by an increase of 150 ms on average. The percentage of intercontinental paths from their VPs was evaluated to 95 % in *WAf*, 70 % in *CAf*, and 60 % in *SAf*. Hence, they suggested the use of Software Defined Networks (SDN) in IXPs, multi-path traffic engineering, and application-specific traffic engineering. Most recently, the authors of the 2017 white paper [94] performed a large-scale measurement of the African Internet covering 52 countries and 319 networks across Africa with the commercial measurements service Speedchecker [269]. They highlighted an excessive reliance on international transit providers as well as the existence of communities, in which countries have built up low delay interconnectivity.

Our studies [81, 85], which we detail in Section 3.2.1, contrast with the related work for the following reasons. First and foremost, their timelines covering 2013 – 2017 makes them an early and longitudinal work on the African Internet as never conducted before (*cf.* Figure 1.5). Given that one of the objectives of this work is to *help researchers undertake interconnection studies by making our measurement datasets available and freely accessible* (Section 1.2.1), we did not consider using any commercial measurement service, contrary to Formoso *et al.* [94] who have recently adopted Speedchecker [269] for measuring the African interdomain topology. We believe that the use of commercial measurement networks, apart from bringing a financial barrier to the typical researcher, does not allow the latter to have a full control of the measurements process. Outsourcing the data collection process indeed brings an opacity that may have prevented other researchers from trusting our results, from being interested in reproducing our measurement techniques in other underdeveloped regions, or even investigating other topics related to the African Internet. Most of all, it may have prevented us from sharing our measurement outputs or analysis with the Internet community, as we did in [84, 87, 88, 90, 256]. We thus considered instead existing open measurement platforms, notably the RIPE Atlas network [248, 250] and the Archipelago [40] platforms for diversely delving into the analysis of the African interdomain routing.

Formoso *et al.* [94] enumerated as reasons for preferring Speedchecker [269] the fact that

RIPE Atlas has a strong bias towards university networks and that around half of all probes in Africa are hosted in South Africa (ZA). We performed some checks in September 2017, based on which we show why this is not the case in the following paragraphs. While doing so, we present the results for only IPv4 probes, as it is the only protocol family that [94] used in their study.

As of September 25, 2017, the RIPE Atlas network contains 13,767 probes, of which 7,775 are active. These active devices are hosted in 2,963 ASes. Similarly to Xenofontas *et al.* [67], we used information in the data from RIRs to distinguish university networks from among those ASes. For every probe IP, we checked whether or not the RIRs records data contain the names of any of the academic institutions, universities, colleges, or research institutes worldwide [118] or any keyword hinting to an academic institution (*e.g.*, laboratory, school, campus, institute, research, academy, to only name a few). Of the 2,963 ASes, in total 288 ASes hosting 902 active probes are marked university networks. That is to say, university networks represent roughly 9.7 % of all ASes hosting an active RIPE Atlas probe. In Africa, roughly 15 % of the 120 ASes hosting a RIPE Atlas probe are university networks; these networks host only 21.2 % of online probes. Therefore, one can not conclude that the RIPE Atlas measurement infrastructure is biased towards university networks neither worldwide, nor in Africa.

Our checks also reveal that 548 (resp. 189 online) IPv4 probes are hosted in Africa, of which 125 (resp. 61 online) probes are located in South Africa (ZA). In a nutshell, 32.3 % online IPv4 RIPE Atlas probes in Africa (*i.e.*, not about a half) are located in ZA. Further, of all African networks hosting a RIPE Atlas probe, 31.2 % are based in ZA. We then looked into the AfriNIC allocations to find that 28 % ASNs and 31.1 % IPv4 blocks are allocated to ZA by September 25, 2017. In summary, RIPE Atlas probe deployments in the African region, far from being biased towards ZA, are consistent with the portion of Internet number resources allocated to ZA and representative of the Internet development in that country compared to other African countries.

Moreover, our studies present discoveries of the Internet infrastructure in the region based on measurements performed from access to access networks regardless of their type (residential, university networks, ISPs, etc.). In fact, we aim at studying how African networks are interconnected to one another from an end-user perspective (*i.e.*, seen from our VPs). Contrary to [49], we do not only focus on university networks. Instead, we perform our measurements from a wide variety of networks, and at random periods of time covering 2013 to 2016 for highlighting topological changes: by doing so, we make sure that our datasets are not biased towards an African country or sub-region.

Unlike [49, 117], we used a broader base of VPs. We ran full mesh paris-traceroutes among all (324) RIPE Atlas probes scattered throughout Africa to assess the interdomain routing. We also performed paris-traceroute among subsets of probes in countries where sustained traffic localization efforts are made by local networks to determine the effects of the launch of new IXPs. We showed in [81] varying ISP transit and peering habits. We underlined the reliance on ISPs based outside the continent for serving intra-continental traffic but found that in the meantime new IXPs were launched in the region. Further, we explored in the longitudinal study [85] the evolution of

the interconnectivity among local African networks over the last four years, highlighting the prevailing dominance of intercontinental ASes. By inspecting both existing and recently established IXPs located in Africa, we show that ISPs do peer locally. With five case-studies, we then evaluated the impact of IXP infrastructures on AS path length and end-to-end delays among peers, illustrating the benefits of initiatives to promote peering.

The computation scripts used in [81, 85] are all written in the Python programming language and query a local MySQL database containing the collected data parsed following a well-defined format. We release them under the format of an application accessible by everyone at [256], as part of our contributions in this thesis (Section 1.3). By contrast, the already released open python code base *IXP Country Jedi* by Aben *et al.* can be used by anyone to create a snapshot of a country and does not require a back-end database [2, 4]. This code produces visualizations [3] that show if paths with end-points within the same country stay in the country and if local IXPs are used.

2.1.2. Congestion in the IXP substrate

In the US or Europe, some studies have found that interdomain congestion occurs due to peering disputes [103, 167]. Genin *et al.* [103] studied patterns of congestion distribution in Asymmetric Digital Subscriber Line (ADSL) and cable ISPs networks. Luckie *et al.* [167] then inspected challenges in inferring Internet interdomain congestion. They developed the TSLP method, which consists of inferring from RTT measurements run from a VP to the near and far ends of an interdomain link, the occurrence or not of congestion. They also validated the TSLP method using traffic data from a research network. This technique has the advantage of allowing an outside player to monitor congestion without explicit cooperation from the network operators. In their study of the effects of routing changes and congestion on the latencies between servers in the core of the Internet, Chandrasekaran *et al.* [47] found a vast majority of the interconnection links with congestion to be private interconnects. In contrast, very few studied links established through IXPs were found to experience congestion: most IXPs provide Service Level Agreements (SLAs) [48] or automatically assign (and charge) ports when a given switch port is utilized by, *e.g.*, more than 60% for a period, *e.g.*, peak hours. Most recently, the usage of throughput measurements to infer congestion on points of interconnections between ISPs [283] as well as TCP congestion signatures were also inspected [282].

Still, much less is known about the nature of congestion and its causes at IXPs in developing regions in general, and at those located in Africa in particular. In fact, Chetty *et al.* [54] measured broadband performance in ZA using measurement software implemented on mobile phones and home routers. They found that users in ZA do not get advertised speeds, and the interconnection (or lack of it) between local ISPs mainly influences reliability and users performance. Despite these and the great push of stakeholders to setup IXPs in Africa [6, 81, 85, 156], we found no previous research to inspect congestion in the African IXP substrate. To fill this lack of congestion-related measurements, we investigated the causes of congestion in the African IXP substrate in [78]. Towards this end, we notably applied the TSLP method on fine-grained

measurements run over a year using Ark probes strategically deployed at six African Internet eXchanges [292]. Further, we extensively analyzed these collected datasets to understand the extent and the nature of the detected congestion cases. Following that, we interviewed the IXPs operators to delve into the root causes of the observed congestion events.

2.1.3. Routing data analysis

An extensive amount of research has been carried out on RouteViews data such as [50, 113, 195, 241, 263]. Recently, CAIDA provided to the Internet community its open source software framework BGPstream, which facilitates live and historical BGP data analysis [39, 209]. In fact, the University of Oregon RouteViews project [189] and RIPE RIS [253] are the most popular projects operating route-collectors and continuously updating their information. RouteViews manages a passive raw routing data collection system, which stores the BGP routes exchanged among the peers at the IXPs at which it is deployed. Its data have been daily collected since 2004 and are publicly accessible. IXP participants which peer with RouteViews may agree or not to exchange their full routing tables, thereby providing respectively either a *global viewpoint* or a *peering viewpoint*, seen from their respective IXPs. In this thesis, the term *peering viewpoint* refers to the set of AS paths received by a route-collector (deployed at an IXP) to which IXP members solely announce their networks and those of their customers (but neither those of their peers nor those of their transit providers). As of September 2017, there are in total 19 RouteViews collectors in the five Internet regions. In the meantime, 21 RIPE RIS route-collectors, all deployed at IXPs in Europe, aim at achieving the same purposes. Similarly, PCH [219] adopted an open peering policy thanks to which it peers with all IXP members that are willing to do so. Contrary to RouteViews collectors, PCH boxes always offer a peering viewpoint, since their peers only exchange the routes of their customers, rather than their full routing tables. Since 2003, PCH has been peering worldwide at 139 IXPs covering 68 countries. The collected data is also made public at [213].

Unlike in other Internet regions, only three RouteViews collectors are located in Africa (at Kenya Internet eXchange Point (KIXP) in KE, Johannesburg Internet eXchange (JINX) and recently NAPAfrica in ZA) as of September 2017. In contrast, PCH route-collectors are deployed at 23 (63 %) IXPs, including at KIXP and JINX. These are hosted in 18 (33.3 %) countries (Section 3.3.1.3.1). Furthermore, some local IXPs deployed their private route-collectors or route-servers with which each member is suggested to peer. These infrastructures enable the collection of exchanged routes locally and facilitate peering setup for newcomers, as noticed by [243]. Contrary to RouteViews and PCH datasets, these data are not publicly accessible.

Despite the existence of these facilities, there is a lack of studies on historical routing data collected at African IXPs. In fact, analyzing such data may give a glimpse of how ASes have been behaving at those IXPs, the evolution of those facilities over time, their richness regarding reachable ASes or prefixes, etc. In the context of overall efforts [6] to localize traffic, this study is critical for decision-making stakeholders, and the results can also incentivize new ISPs or content

providers to join the existing IXPs of their choice, given their interests.

To correct this lack, the ISOC[1] partnered with UC3M[2] to start the African Route-collectors Data Analyzer project (ARDA) project [87]: we built an open source web platform, which regularly computes and displays statistics based on routing data collected in the AfriNIC region and which is easily replicable in other ones. With this tool, we can evaluate in real time key statistics that could help IXPs market their features and reporting on routing inefficiencies, make everyone witness the interconnection growth and gaps, etc. (Section 3.3.1).

2.2. Content delivery

With respect to broadband services, Bischof *et al.* [31] explored the performance experienced by end-users in their analysis of data collected from end-hosts and residential gateways in 160 countries. They provided insight into the impact of broadband services market characteristics such as pricing, cost of increasing capacity, and connection capacity on network usage. There have also been various studies on content delivery infrastructures. Calder *et al.* [37] enumerated the IP addresses of the infrastructure of Google, finding their geographic locations, analyzing its growth, and matching users to clusters. Similarly, Farahbakhsh *et al.* [91] depicted and studied the global picture of the current Facebook network infrastructure, including native Facebook servers and Akamai nodes. Otto *et al.* [210] examined the role of Domain Name Server (DNS) in the redirection process, exploring the potential of the Extension mechanisms for DNS (EDNS0). We note similar studies have been expanded to other CDNs such as EdgeCast and CacheFly [274]. Prominent works have further analyzed redirection strategies to understand how CDNs map users to edge caches. For example, Su *et al.* found that Akamai primarily redirects clients based on active network conditions [276]. More recently, Fan *et al.* [76] have evaluated the dynamics of the mapping of network prefixes to Google Caches (GGCs). They underlined a high variance across the servers mapped to each location with nearby clients often being redirected to clusters, which are far apart.

Expanding Internet deployment in Africa has received a lot of attention recently, mainly from local organizations and Internet developmental institutions, such as the AU and ISOC [6, 157, 207]. There has also been an expanding push from companies like Google, Facebook, Cloudflare, Akamai, Microsoft, etc. who see the economic potential of Africa [111]. Of particular interest has been the use of IXPs [78, 81, 85, 87, 117], which are seeing an expanding uptake. It has been followed by a range of performance studies. For example, Pejovic *et al.* [222] concluded from their research on broadband services adoption in the rural areas of sub-Saharan Africa that restricting access to public terminals and workplaces severely hinders the type of applications used online. Chetty *et al.* [54] investigated mobile performance, finding that it can often be superior to wireline. Zaki *et al.* [315] then focused on web performance, highlighting that critical

[1] www.internetsociety.org
[2] www.uc3m.es

bottlenecks include slow DNS resolution and a lack of local caching. They found that DNS caching, redirection caching, and the use of SPDY [74] can all yield substantial improvements to user-perceived latency. We take this as clear evidence of the limitations of solely provisioning better connectivity and not considering the upper layers.

After the extensive inspection of the African interdomain routing using active and passive measurements in Chapter 3, another major theme of this doctoral thesis is thus exploring the web ecosystem serving Africa (Chapter 4). Our work [89] is orthogonal to that of Bischof *et al.* [31], focusing on web infrastructure, rather than end-user choices. Our focus also differs from previous work aiming at analyzing redirection strategies in that we target web deployments in the African region. In fact, we inspected this topic to understand the current state of content infrastructure in the region, while improving existing methodologies through the combination of several measurement approaches. Following that, we take a broad perspective, looking at several different websites, CPs, and network operators.

2.3. Topology and infrastructure

The primary business model of an IXP consists of operating and managing a physical infrastructure to support public and private Internet interconnections [13]. Striking examples are those of AMS-IX, NetNod, 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 is a result of their governance structure [48].

Interconnecting IXPs is a contentious issue since there are as many arguments for it as there are against it [48, 93, 200]. 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 [93, 98]. In 2012, Nipper [200] argued that an IXP should not go beyond its diameter[3] 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. As for the particular case of smaller IXPs, Nipper, however, acknowledged that interconnecting such IXPs could contribute to gain more critical mass. In fact, the point at which the value of participation at a given small IXP is equal to the cost of participation (*i.e.*, the critical mass point) becomes higher, as there are more potential members after the interconnection. In this respect, Fenioux [93] argued in 2015 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 (FR) for example, Rezopole operates two IXPs that are interconnected (LyonIX and GrenoblIX). Moreover, these IXPs are interconnected to other IXPs in FR or abroad such as FranceIX, Equinix, NetIX, SFINX, fr-IX (all in Paris, FR), EuroGIX (in

[3] We note that the term *diameter of an IXP* can be defined as the direct reach of the IXP with its own equipments, which depends on the geographic location of the infrastructure [200].

Strasbourg, FR), TouIX (Toulouse, FR), CIXP (Geneva, Switzerland (CH)), and Top-IX (Turin, Italy) [35, 48, 93, 114, 171, 200]. Further, FranceIX deployed interconnections with not only the above-listed IXPs but also LU-CIX in Luxemburg, enabling its members to peer with theirs [98]). Another example of IXP interconnection is that of InterLAN (Bucharest, Romania (RO)) and BalkanIX (Sofia, Bulgaria (BG)) via a 10 Gbps link [48, 200].

Sprague *et al.* [270] listed the infrastructure among the barriers to Internet adoption in developing regions, thereby raising the need for looking more closely at it in those areas. In this regard, Galperin *et al.* [100] reported on the connectivity in Latin America and the Caribbean, advocating for the development of IXPs as an essential step to improve the quality and the coverage of access in those regions. In the meantime, several projects have been investigating the infrastructure of the African Internet, underlining the bottlenecks or encouraging points of its evolution. To begin with, the importance of reliable energy as a corner-stone for Internet Technology (IT) development in developing regions was underlined by Crowcroft *et al.* [61] and Ncube *et al.* [196]. Projects [177, 198, 267, 277–279] can also be listed as constant efforts to map submarine and terrestrial cables serving the continent. Using these maps, Nyirenda-Jere *et al.* [207] highlighted, for instance, that from 2009 to 2014, the international bandwidth of Africa increased twenty-fold and its terrestrial network more than doubled. This drastic increase was detailed in [258] along with the corresponding investments. Still, as of September 2017, the terrestrial fiber network remains fragmented [85, 198].

As a solution for improved QoS for local traffic and decreased transmissions costs for ISPs, Kende *et al.* [156] reported on the benefits that IXPs have had in KE and Nigeria (NG), two emerging markets in sub-Saharan Africa. They concentrated on these countries, as their respective IXPs (KIXP and Internet eXchange Point of Nigeria (IXPN)) appear as leading examples of growing IXPs in the region. This ISOC study reveals that in 2012, KIXP localized approximately 1 Gbps of peak traffic and reduced latency from 200 –600 ms to 2 –10 ms, allowing ISPs to save almost US$1.5 million on international transit costs per year. In the meantime, IXPN localized 300 Mbps of peak traffic and reduced latency from 200 –400 ms to 2 –10 ms, leading to savings of up to US$1 million per year on international transit costs [156]. The authors further highlighted how these facilities paved the way for a significant increase in performance for end-users, a boom in content usage and corresponding revenues for mobile traffic, as well as social benefits from e-governments access to IXPs.

From 2012 up to now, the number of IXPs set up and active in Africa has drastically increased from 26 to 38 [85, 292]. Accordingly, the findings of our longitudinal studies [77, 78, 81, 85, 87, 89] on the African Internet ecosystem all suggest that local stakeholders intensify peering in the region. Intensifying peering in the region [85] could be achieved by enabling ISPs present at any two isolated local IXPs to peer. A possible way to realize this is to establish a link between the IXP infrastructures. In this respect, we propose and evaluate in [79] a framework to build a distributed African IXP, interconnecting existing IXPs using available fiber networks, given external factors that influence the Internet in the region (Chapter 5). A direct consequence of the implementation

of this framework would be that paths from any African country to another, rather than traversing a different continent, are routed within Africa through a hierarchical IXP substrate.

Along these lines, Noordally *et al.* [201] later suggested, in their 2017 study, the setup of a regional IXP in the Indian Ocean Area, demonstrating the interest about this topic in the region. However, unlike our study, they did not include socio-economic and political constraints, which we found to give our work a solid grounding in reality and make our framework realizable under the present-day constraints. Note that our constrained solution covers not only the African continent but also nearby islands, including those of the Indian Ocean Area such as Mayotte (YT), Seychelles (SC), Reunion (RE), Mauritius (MU), and Madagascar (MG).

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 Center (IDRC) and the ITU [134] 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 [275]. In the same year, Pehrson *et al.* proposed [221] 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 [198, 265, 279]. East African IXP operators proposed [72] 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 the equal promotion of all IXPs and the absence of competition between IXPs and their members. In contrast, other sub-regional communications organizations [71, 203, 268] 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.

2.4. Taxonomy of the studies on the African region

As a takeaway from this chapter, we provide in Table 2.1 a taxonomy of all studies related to the African Internet ecosystem. We categorize them according to the main topics addressed in this work. To highlight how our research has contributed to creating a renewed interest in the African Internet, we order all studies of each category by publication date and precise their respective year of publication: academic publications (peer-reviewed) are preceded by the symbol ‡, while the symbol † precedes white papers and other articles. Moreover, we separate those released before from those published during our studies. Furthermore, we label, in the latter category, the studies in which we participated and that we co-authored by the symbol ⋆ to separate them from those carried out by other authors.

Table 2.1: Taxonomy of articles, white papers, academic/scientific papers, and other research studies related to the African Internet ecosystem.

Addressed topics	Year of publication	Interdomain routing				Content delivery	Topology and infrastructure	
		Topology discovery	Communications performance	Interdomain Congestion	Routing data analysis		Infrastructures	IXP interconnection
Published before this work (before 2013)	N/A		PingER [223]		Smith [263]		NSRC [198] Song [267]	
	2006						‡Pehrson *et al.* [221] ‡Stucke [275]	†Pehrson *et al.* [221] ‡Stucke [275]
	2007	‡Gilmore *et al.* [106]						
	2012		†Kende *et al.* [156] ‡Pejovic *et al.* [222]			‡Pejovic *et al.* [222]	†Kende *et al.* [156] STF [280]	
Published during this work (2013 – 2017)	2013	Les Cottrells [58]	‡Chetty *et al.* [54]				STF [277] †Manyika *et al.* [181] Mahlknecht [177]	
	2014	‡Gupta *et al.* [117] ⋆Fanou [77] ‡Chavula *et al.* [49]	‡Gupta *et al.* [117] ⋆Fanou [77] ‡Chavula *et al.* [49] ‡Zaki *et al.* [315] †Sprague *et al.* [270]			‡Gupta *et al.* [117] ‡Zaki *et al.* [315]	‡Gupta *et al.* [117] ⋆Fanou [77] STF [278] ‡Chavula *et al.* [49] †Sprague *et al.* [270]	
	2015	⋆‡Fanou *et al.* [81] Aben [3]	⋆‡Fanou *et al.* [81] Aben [3]			†Kende *et al.* [157]	⋆Crowcroft *et al.* [61] [a] ⋆‡Fanou *et al.* [81] Aben [3] †Nyirenda-Jere *et al.* [207]	
	2016	⋆‡Fanou *et al.* [89] ‡Noordally *et al.* [202]	⋆‡Fanou *et al.* [89] ‡Noordally *et al.* [202]			⋆‡Fanou *et al.* [89]	⋆‡Fanou *et al.* [89] STF [279]	
	2017	⋆Fanou *et al.* [85] †Formoso *et al.* [94] [b] ⋆‡Fanou *et al.* [79] ⋆‡Fanou *et al.* [78]	‡Fanou *et al.* [85] †Formoso *et al.* [94] ⋆‡Fanou *et al.* [79] ⋆‡Fanou *et al.* [78] †Noordally *et al.* [201] ITU [136] ITU [138]	⋆‡Fanou *et al.* [78]	⋆Fanou *et al.* [87]	⋆‡Fanou *et al.* [79]	⋆‡Fanou *et al.* [85] †Ncube *et al.* [196] ⋆‡Fanou *et al.* [85] ⋆‡Fanou *et al.* [79] ⋆‡Fanou *et al.* [78] ⋆‡Fanou *et al.* [78]	⋆‡Fanou *et al.* [79] †Noordally *et al.* [201]
	2013 – 2017		PingER [223]		Smith [263]		NSRC [198] Song [267]	

[a] We only consider here the sections “Which infrastructure for a better Internet in Africa?” and “Challenges in developing regions” of this Dagstuhl report that are respectively copyrighted “Rodérick Fanou” and “Weverton Cordeiro and Rodérick Fanou” under Creative Commons BY 3.0 Unported license.

[b] As of this writing (September 25, 2017), this work is made public, although being non-peer reviewed and has thus been considered as a white paper.

Chapter 3

African Interdomain Routing

In this chapter, we first present our deployment efforts to help build in the region the measurement infrastructure based on which we perform our studies. Next, we describe our methodology to inspect the interdomain routing using data collected during our active measurements as well as those collected through passive measurements by existing route-collectors located in Africa. We then detail our findings and suggest ways to improve communications performance. The key topics covered include measurement infrastructure deployment, diversity of the African Interdomain routing, IXPs mapping, impacts of existing and recently launched IXPs, causes of congestion in the African IXP substrate, and route-collectors data analysis.

3.1. Building the Internet measurement infrastructure in Africa

Confronted with a near non-existence of measurements devices and the lack of data on IP networks in Africa [54], we started by helping build its measurement infrastructure. We had the option of either building our own infrastructure or extending existing measurements platforms. On the one hand, the former option has the drawback of preventing other researchers from easily accessing or sharing our platform/data for their studies and the disadvantage of preventing us from using the few existing devices. However, it gives us more leeway to construct a platform responding to all our needs regarding data collection. On the other hand, the latter option prevents us from defining and controlling the parameters that we measure but gives us the possibility to use probes or data collected by other researchers or even share our measurements datasets.

After weighing the pros and cons, we chose to start both options, *(i)* selecting West Africa as our first focus, since most[1] existing measurement devices were in Southern, East, and North Africa, and *(ii)* planning to enlarge our deployment efforts to other sub-regions later. We then adopted as materials raspberry pis (raspis) for our measurement infrastructure as well as the RIPE Atlas network [248,250] and the CAIDA's Archipelago (Ark) measurement infrastructure [40] as existing measurement platforms to extend. We further detail the reasons for these choices in the

[1] A dozen of RIPE Atlas probes [248] and five Ark probes [40] as of February 2013

Figure 3.1: Deployed probes: from the left to the right, a raspi used in our own raspi-based measurement infrastructure, a RIPE Atlas probe, and an Ark probe.

subsequent sections. The selected devices are presented in Figure 3.1. To each of them is added a UTP cable and a power cable, essentials for its installation.

3.1.1. Deploying our own raspi-based measurement infrastructure

3.1.1.1. Technical specifications, tools selection, and platform implementation

We chose raspis [96] as devices for our own measurement infrastructure because, since they are cheap devices (US$45 each), we can populate them in many hosting sites with low expenses. Most importantly, these small computers developed on a simple card are compatible with all types of primary input devices. Next, their hardware (700 Mhz processor, Ethernet 10/100 Base-T, 512 MB memory) can support our measurement scripts that do not request a high workload. Moreover, their OS (Linux) supports not only the Python programming language (with which our scripts are edited) but also MySQL for the databases hosting the measurement results.

We then chose to setup a client/server architecture in which the raspis play the role of clients and must often exchange information with a server located in the UC3M premises (Madrid, Spain (ES)). These information include the latest versions of the measurement scripts to be run by the raspis, the updated list of IPs to probe, the type of measurements to be transmitted by the server to the raspis. They also include the status of each probe and the measurement outputs, which are transmitted from the raspis to the server.

The ideal situation for achieving successful measurements would be that each deployed raspi is connected to the Internet 24 hours/7 days. However, the targeted countries are characterized by frequent power cuts due to the lack of energy or the inability of energy suppliers to support the increasing demand all the time [61, 196]. Moreover, our devices can suffer from random Internet outages. We, therefore, chose to configure each raspi so that it: *(i)* restarts and reruns all the scripts automatically after any blackout or Internet outage period, *(ii)* always keeps a backup copy of its latest measurement results that have not yet been transmitted to the server for a period *P1*. We also found necessary that each of them *(iii)* sends all remaining measurement results to the server after a period *P1* and gets the corresponding acknowledgments before deleting them from its database, *(iv)* retransmits all the backed up results that could not be transmitted at the end

of the next period *P1*.

After that, we opted for using the open source software Puppet [228, 230] to ensure a centralized management of the system. In fact, Puppet has been largely adopted and is, for instance, used for running vast server farms [229, 230]. Moreover, it is well tested and can help support the remote management of an arbitrary amount of nodes securely, while ensuring scalability. Furthermore, Puppet allows us to configure cron jobs centrally. Using it, we set up the system as explained below starting from early July 2013.

The raspis were configured with the capabilities of pushing and pulling files to/from the server, which plays the role of a puppet master. Since some hosts might not have public IP addresses and since we needed a mechanism to secure the files transfer between the components of the system, we built a Virtual Private Network (VPN) between the raspis and the puppet master. That way, the raspis could fetch their configuration from the puppet master under the format of syslog messages. They could also ping one another, a capability that is essential for running full mesh measurements. Their availability and proper functioning were monitored continuously by a Munin instance [192] installed on the server. After that, we used the Secured SHell (SSH) protocol for delivering the measurement results.

Additionally, a driver (set of python scripts) running on the puppet master was responsible for coordinating measurements and sending to each raspi, the list of IP addresses to probe. A prober (set of python scripts) running on each raspi was in charge of executing the measurements and frequently sending back the results to the server. Our first tests only involved 8 raspis and included paris-traceroutes [24] (implemented using scamper [165]) and MTR [162] measurements for respectively IP paths discovery, end-to-end delay, and packets loss measurements. They took place in the servers room of UAC (BJ) at end July 2013 before the deployment tour (Figure 3.2a).

Finally, each raspi was storing its latest measurement results in its MySQL database. Meanwhile, the server hosted a similar MySQL database per raspi in which it gathers the measurement results transmitted by the corresponding raspi over time. Those parsed outputs were sent (every $P1 = 5$ days) under the format of a MySQL database backup. After their analysis on the server side, the list of IPs to probe by each raspi was frequently updated. We built this system with the goal of later extracting high-level information from these data for generating periodical reports accessible by the probe hosts.

3.1.1.2. Raspis deployment efforts

Early February 2013, we planned the deployment of 16 raspis split into two sets of 8 raspis. The targeted networks were those operating in the West African countries Benin (BJ), Togo (TG), Niger (NE), Ghana (GH), Ivory Coast (CI), Burkina Faso (BF), Mali (ML), and Nigeria (NG). We then started by building a list of potential prospects to host our probes using for instance [197], defining a probe deployment strategy, and preparing its implementation. From May to July 2013, we contacted our potential hosts (in universities, residential networks, local ISPs headquarters, hotels, cybercafes, members of ISOC local chapters, to only name a few) through online calls

(a) *Raspis configuration in the servers room of Université d'Abomey Calavi (BJ) before the deployment tour, July 2013.*

(b) *Raspi and RIPE Atlas probes deployment at Hotel le Châtelet (Togo, TG) in the router of the host network, August 2013.*

(c) *RIPE Atlas probes deployment at UAC (BJ), August 2013.*

Figure 3.2: Setting up and deploying the raspi-based measurements platform while extending the RIPE Atlas network in the African region.

and emails, highlighting the interests of our project for network operators, CPs, or researchers, and requesting their adherence. Prospects contacted in 12.5 % of targeted countries notified their adhesion to the project. As they accepted the terms of our pre-defined Non-Disclosure Agreement (NDA), we undertook to co-sign it. The remaining either requested face-to-face meetings or promised to further discuss with their hierarchy before taking their final decisions. We were thus obliged to travel throughout *WAf* *(i)* to meet those who were still hesitant for clearly explaining the objectives of this research project and triggering their interests or *(ii)* to deploy the probes within the networks of those who adhered (sometimes in network installations with makeshift compromises as shown in Figure 3.2b). Finally, we did so *(iii)* to meet as much as new prospects in the region as we could. This step lasted from August to December 2013, at the end of which our raspis were hosted in 10 ASes operating in six countries (Table 3.1).

Table 3.1: ASes and countries hosting our deployed raspis

CC	Countries	ASes	#Deployed raspis
BJ	Benin	AS28683, AS37090	2
BF	Burkina Faso	AS25543, AS37073	2
TG	Togo	AS30982, AS24691	2
CI	Ivory Coast	AS36946, AS29571	2
GH	Ghana	AS29614	1
NG	Nigeria	AS37480	1

Nevertheless, for many raspi hosts, the NDA signature process was then stalled by the hierarchy before December 2013. Before we could perform the full-mesh measurements, collect a significant amount of data for analysis, we noticed that most raspis got down one after the other. While some of the corresponding hosts did not respond to our emails, others explained to us that their hierarchy considers the raspis as intrusive devices and threats to the security of their network. They specified that by hosting them, they allow us to run at will scripts that they are not able to always control. Hence, most deployed raspis were unplugged. In the meantime, Gupta *et al.* [117] released their study on the first look at ISP Interconnectivity in Africa using data col-

lected with BISmark nodes in networks operating in Kenya (KE), South Africa (ZA), and Tunisia (TN). Fortunately, most RIPE Atlas probes concurrently deployed (Figure 3.2c) remained online: their hosts were confident in keeping them plugged for reasons detailed in Section 3.1.2.2.

We learned from our experience that deploying a measurements network is a challenging task of the data collection process: it requires to construct and look after a sustainable human network across stakeholders, while human relationships need time and trust to be built. Moreover, we learned from this initiative that ensuring trust and partnership with the industry is a sine qua none condition for successfully setting up a sustainable measurement infrastructure in the region. More specifically, we understood that we needed to setup trustful relationships (so that our host always make sure the devices are online) and to use VPs that are practically maintenance-free (deployed in record time within networks whose wiring respect recognized standards) for achieving our goals. Considering the broad adoption of RIPE Atlas probes on other continents, and the possibility to complement our deployed RIPE Atlas probes with those existing in the region [34, 36, 248, 250], we decided to enlarge right away our focus to Africa while relying exclusively on those VPs for improving the related work at this stage of our research.

3.1.2. Extending existing measurement platforms

3.1.2.1. Building trust and partnership with local operators and stakeholders

We present in this section our teaching efforts destined for decision-makers, regulators, network engineers, and IT stakeholders, playing a pivotal role in the Internet furniture chain. We also explain the reasons underlying our attendance to several network operators meetings worldwide.

In August 2012, the ISOC was chosen by the AU to conduct technical aspects workshops for supporting the establishment of IXPs in its member states as part of the AXIS project [6, 141]. The AXIS capacity building workshops on "Technical Aspects of Setting up, Operating, and Administering IXPs" have thus been organized in each African country to help raise the awareness of the stakeholders on the necessity to build a local IXP, join it as a member, peer with existing members to localize traffic, and help address the issues mentioned in Section 1.1. They aimed at pointing out the importance of those facilities and giving the participants the technical capabilities to set up those infrastructures themselves.

As illustrated in Figures 1.5 and 3.3, I was then selected by the ISOC to lead the AXIS workshops in Burkina Faso (March 2013), Niger (July 2013), Benin (August 2013), Mauritania (October 2013), Congo-Brazzaville (March 2014) [131, 132]. Moreover, I co-led the AXIS workshop in Liberia (May 2014) and the ICT workshop in Ethiopia (December 2013). The number of participants registered at those events is on average 21. The courses lasted 40 hours and spanned five days each (Appendix A). Additionally, I trained the Network Operating Center (NOC) of ISOCEL Telecom, a local ISP in BJ, on routing protocols and IXP setup and gave a networking and protocols course at "Institut de Formation et de Recherche en Informatique" (IFRI/UAC, BJ) in July – August 2014. These teaching activities contributed to building trust and partnership with

local operators, essential for the deployment of a considerable amount of RIPE Atlas probes in several networks operating in the country host, close to local ISPs headquarters.

(a) *AXIS Workshop in Burkina Faso, BF (March 2013)*

(b) *AXIS Workshop in Benin, BJ (August 2013)*

(c) *AXIS Workshop in Congo-Brazzaville, CG (March 2014)*

(d) *AXIS Workshop in Niger, NE (July 2013)*

(e) *AXIS Workshop in Mauritania, MR (October 2013)*

(f) *AXIS Workshop in Liberia, LR (May 2014)*

Figure 3.3: Contributing to African Union and Internet Society's initiatives for promoting IXPs [6, 141] by leading AXIS workshops, while building trust and partnership with local operators.

In the meantime, I also attended several operators meetings, mostly in the region (Figure 1.5). During those events, I constantly reported on our measurement results to the community and discussed with the network operators to get their feedback for a more impactful research. These include AfPIF (2014, 2015, 2016) [142–144, 147], RIPE (70, 73) meetings [246, 247], the Workshop on Active Internet Measurements (AIMS) (AIMS-8) [55], BGP Hackathon 2016 [62], etc. Notably, at each AfPIF conference [147], operators meeting that promotes national and cross-border interconnection, I could spread the probes all over the continent (along with other RIPE Atlas ambassadors) by giving them out to the ISPs engineers or the IXP operators of each country.

3.1.2.2. RIPE Atlas probe deployment efforts

As already mentioned, measuring African networks involves many challenges, which influenced our choice of the measurement infrastructure. First, operators are hesitant to deploy foreign devices into their networks, for security and privacy reasons. Meanwhile, we had to find a relevant number of hosting locations for the measurement devices so that our study covers the whole continent. Second, any device deployed for this purpose has to be robust, as power outages and surges frequently occur in the countries under study. Third, the devices cannot be expensive, since we have no guarantees that all our collaborators will keep them online. Finally, we preferred an open measurement infrastructure (contrarily to the recent work [94]), as it provides the means for other network operators and researchers to also utilize the infrastructure and its publicly available

data to study the African Internet.

To best deal with these challenges, we chose, amongst other options, to extend the RIPE Atlas measurement platform: RIPE Atlas consists of over 10,000 online devices deployed worldwide in various locations, as of September 2017 [248, 250], which traduce a massive adoption. Any individual wanting to host a RIPE Atlas probe can do so. For individual users, the probes are free to obtain and to deploy; they are secure, robust against power outages, and require no maintenance. They can perform multiple types of measurements on IPv4 and IPv6, including the ping and the paris-traceroute, which we use in Sections 3.2.1.2 and 4.1.1, as well as Hypertext Transfer Protocol (HTTP) requests and DNS queries, which we adopt in Section 4.1.1. A RIPE Atlas probe host can perform measurements from any probe around the world. Also, the measurement source code is publicly available [34, 36, 251]. Another non-negligible reason for which the hosts were confident in keeping them within their networks is that they belong to RIPE NCC, an RIR that is well-known to African networks for being in constant collaboration with AfriNIC.

(a) *RIPE Atlas network on June 2013: about 83 devices were deployed in African networks. Green triangles correspond to active probes while red triangles, to disconnected ones.*

(b) *RIPE Atlas network on October 2014: in total 227 devices are deployed in Africa. Green anchors correspond to active RIPE Atlas probes while red anchors, to disconnected ones.*

(c) *On July 2015, in total 326 devices were deployed in Africa. The green portion of each doughnut represent active probes in the geographical area while red portion are disconnected ones.*

(d) *RIPE Atlas network on August 2017: 526 devices are deployed in Africa. Green dots correspond to active RIPE Atlas probes while orange dots represent disconnected ones.*

Figure 3.4: RIPE Atlas network evolution from May 2013 to August 2017 [248].

Africa only hosted a few (about 83) active RIPE Atlas devices in June 2013, with almost none in the West (Figure 3.4a). Till then, the RIPE Atlas network coverage in the region was low and therefore considered by researchers to be a source of limited data [54]. To improve the situation, we actively contributed to the deployment of 148 RIPE Atlas probes in 69 networks covering 31 African countries over four years (2013 – 2016) [85]. We list further the ASes and countries hosts (Section 3.2.1). The overall increase in the number of probes hosted by the Internet community

in the African region over time is perceptible in Figure 3.4.

We highlight in Figure 3.5 the number of new devices connected per African sub-region in the period November 2010 – February 2017. It shows how from May 2013 to February 2017, there is a higher new probe connections rate compared to the period before May 2013. Probe deployment efforts are more intensive in the Southern part of the continent than in other sub-regions. Note, I deployed in total 53 RIPE Atlas probes in 28 local ASes covering 25 African countries, with a focus on West Africa (*WAf*). RIPE Atlas volunteers and collaborating institutions concurrently deployed a considerable amount of probes in Southern and East Africa, which we also used in this study. We detail in Section 2.1.1 the reasons why the RIPE Atlas network is not biased contrary to the claims of Formoso *et al.* [94].

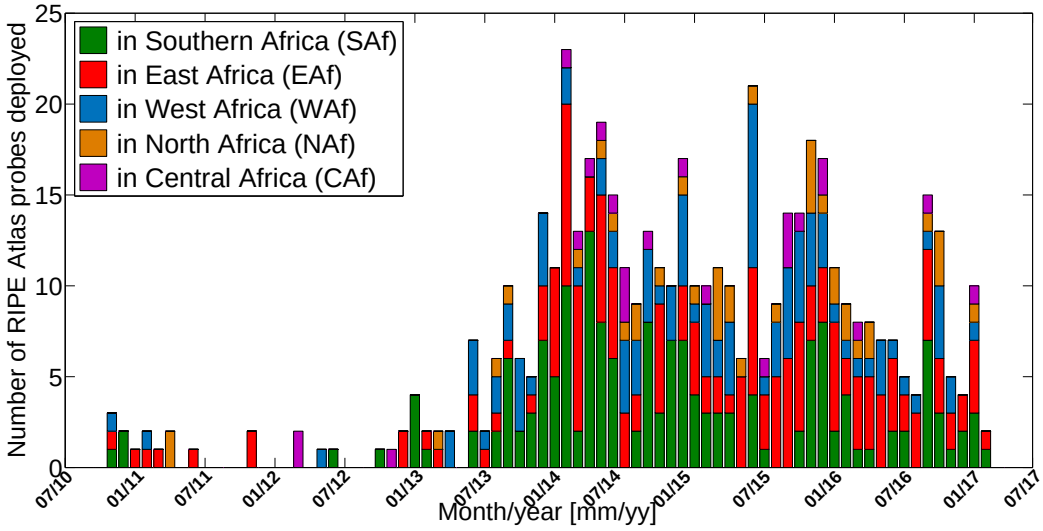

Figure 3.5: Number of new RIPE Atlas probes connected per month in each African sub-region from November 2010 to February 2017.

Our deployed RIPE Atlas probes are hosted either by ISPs, universities, or residential networks. None of them are behind a wireless access link, which reduces the impact of last mile latency on our results. Despite all efforts, it is still challenging to get probes in North and Central Africa, where resistance to hosting external devices in the network is highest. As a consequence, very few ASes/countries from these regions are covered in our study.

3.1.2.3. Ark probe deployment efforts

Ark probe deployment efforts have been conducted from end August 2015 to early September 2016. The goals were not only to extend the reach of the Archipelago (Ark) infrastructure in Africa but also to give ourselves the means to investigate interdomain congestion in its IXP substrate, as detailed in Section 3.2.2. We adopted Ark [40] because of its capacity to perform fine-grained measurements. Ark offers us the ability to run scamper [165] on its monitors for limiting the TTL value of the Internet Control Message Protocol (ICMP) packets, sending a burst

of packets through the congested link, and so one. By using Ark, we also gain more visibility on the events occurring on the IP layer, while performing a longitudinal TSLP based study.

(a) *Ark probes in Africa on May 2013: 5 devices deployed in 5 ASes in 5 African countries.*

(b) *Ark probes in Africa on August 2017: 15 devices deployed in 14 ASes in 11 African countries.*

Figure 3.6: Increase in the number of Ark probes in African ASes from August 2015 to August 2017, highlighting deployment efforts done in that period [40].

The map showing the spread and the locations of the 180 Ark monitors (deployed worldwide as of September 2017) is available at [40]. To the five VPs hosted within African networks at the beginning of our deployment efforts, we actively contributed to adding 10 monitors deployed within 10 ASes including IXP infrastructures, ISPs, and universities networks (Figure 3.6). In the congestion study presented in this thesis, we only consider probes that support TSLP measurements and are deployed at six IXPs [292] strategically selected. These are Ghana Internet eXchange Association (GIXA) [104] in GH, JINX [140] in ZA, KIXP [287] in KE, Serekunda IXP (SIXP) [261] in GM, and Tanzania IXP (TIX) [284] in Tanzania (TZ). Details on the reasons why these IXPs were chosen and on how we set up the measurement devices are provided in Section 3.2.2.

3.2. Active measurements

3.2.1. Four years tracking unrevealed topological changes in the African interdomain routing

The interdomain routing mainly gathers the domains of today's highly commercial Internet, *i.e.*, ASes as well as the economic relationships between them namely *p2p*, *c2p*, or *s2s* relationships [102,107,130]. BGP is the single interdomain routing protocol used on the Internet [33,238] (*cf.* Section 1.1.1). As a first step to reveal the African Internet, we chose to examine in depth its interdomain routing. Recent work targeting such a goal [54,106,117] relied on a very limited set of well-connected VPs and had different focuses (Section 2.1.2). In contrast, the key contribution

of this section is to obtain an interdomain map that covers the entire continent and is not biased towards any particular country or sub-region. Towards this end, I met and convinced 28 local ISPs in 25 African countries (out of 54), to deploy 53 RIPE Atlas probes [248] within their networks. Added to the efforts of our co-authors on this work, we could reach a total of 148 probes hosted in 69 African ASes and located in 31 countries, giving a 278.3 % rise in the number of deployed VPs. We complemented this set of deployed probes with those already present in the region. For obtaining relevant topological data on access-to-access interconnection and tracking the evolution of traffic localization, our measurement campaigns monitored both IPv4 and IPv6 end-to-end paths between RIPE Atlas probes scattered throughout Africa at random periods over the last four years. We propose different techniques, which can be used to treat any set of similar measurements, for analyzing the collected datasets and inferring results that depict ISP behavior.

This section gathers results obtained from our studies [77, 81, 85], highlighting the evolution of the IXP substrate in Africa from 2013 to 2016, analyzing and reporting on more measurements performed among local networks. It then focuses on the detection in our dataset of the usage of/launch of IXPs and compares performance experienced within African countries to those of European countries and the US. Our results illustrate that, except for ISPs based in ZA, the provision of intra-continental paths is dominated by ISPs based outside Africa, while ZA is being adopted as a hub for East-West African communications (Section 3.2.1.4.4). We discover a large variety of ISP transit habits, notably correlated with the location, the official language, and the monetary union of the country in which the ISPs operate (Section 3.2.1.4.4). We further study the impact of those routing trends on the AS path lengths (Section 3.2.1.4.3) and end-to-end delays between ISPs (Section 3.2.1.4.5), notably among networks based in the same country.

Using two methodologies based either on the detection of known IXP prefixes in the traceroute data or on the evolution of AS path lengths and RTTs among local ASes over time (Section 3.2.1.3), we map 23 of the 37 African IXPs and improve previous studies that were not able to infer existing IXPs in the region [25]. As opposed to Gupta *et al.* [117] which indicates that, by and large, local ISPs are not present and do not peer at local IXPs, we highlight how many local ISPs are found to peer at African IXPs in our dataset (Section 3.2.1.5). We also expose the benefits of the setup of new IXPs concerning end-to-end delay (Section 3.2.1.5.2). Next, we evaluate how frequent it is for IXPs from other regions to be traversed by intra-African communications (Section 3.2.1.6) *to reveal that further efforts need to focus on increasing the number of local members at African IXPs so that peering is intensified in the region.*

The remainder of this section is organized as follows. In Section 3.2.1.1, we describe the African interconnection landscape and present our motivations for this work. In Section 3.2.1.2, we give an overview of the data collection and sanity check. In Section 3.2.1.4, we present and analyze our results, which we further discuss and compare to previous work in Section 3.2.1.8.

3.2.1.1. African interconnection landscape

We provide in the subsequent paragraphs an overview of the evolution of the African telecom infrastructure and briefly describe the current state of the African Internet, before concluding with the motivation of our study.

In the early 60's, the incumbent national operators were the sole licensees of the international gateways and phone networks. Since the late 90's, however, there has then been a gradual shift towards the creation of more liberalized telecommunications market environments in Africa. As a result, many competing operators have emerged across the entire range of telecommunications services, such as mobile, fixed, wireless phone, and data services. This has contributed to the partial or full privatization of some of the incumbent operators [172].

Similarly, telecoms operators have invested in both domestic long haul and intercontinental optical fiber deployments to reduce their reliance on satellites links [163,180,259,277–279,289]. As a consequence, Africa is September 2017 linked through 32 submarine cables of various lengths and bandwidth capacities, but the terrestrial optical fiber deployment is still fragmented. Central Africa and the Sahel are the main gaps on the map that segregate other areas of connectivity [177,198,267,277–279,288,289].

The low penetration rate of Africa [137–139,149] contrasts with the boom in mobile networks infrastructures and mobile users. For instance, the percentage of online inhabitants in Africa has increased from 2.4 % in 2005 to 20.7 % in 2015 as shown by [137]. Meanwhile, the rate of mobile users has risen from 12.4 % to 73.5 % in the same period, with a percentage of active mobile-broadband subscriptions of only 17.4 % in 2015. These, however, highlight the substantial potential in Internet users that may be reached and positively affected in the region by the network and the web, especially when QoS increases and prices are lowered [61,137,181]

A challenge in attaining this goal is to ensure that local networks can easily and cheaply exchange traffic within the region instead of exchanging traffic via remote locations [5]. We shed light on this phenomenon and its drawbacks in Sections 3.2.1.4.3 and 3.2.1.4.5. The ability to localize traffic will have significant performance and eventually monetary benefits since local networks will save those costs. We present next our methodology aiming at measuring and better understanding the interdomain topology for identifying where this situation can be addressed.

3.2.1.2. Methodology

We begin by describing the approach followed to identify ISPs playing a pivotal role in transiting Internet traffic between any pair of ASes hosting a RIPE Atlas probe. We then detail the sanity check performed on the collected dataset. Next, we explain how we dealt with unresponsive IPs addresses in the traceroutes outputs, *unknown* ASes in the results of the IP to AS mapping process, or loops in the inferred AS paths. Further, we describe our geolocation methodology based on 10 DSes cross-correlated with ping measurements towards the considered IPs. Finally, we explain the methods used to detect peering links, or IXPs and their members in the dataset.

3.2.1.2.1. Data Collection While actively contributing to the deployment of 148 RIPE Atlas probes in 69 different networks covering 31 African countries (Section 3.1.2.2 and Figure 3.4), we used both deployed and existing probes in the region to conduct 7 measurement campaigns from November 2013 to June 2016 (Table 3.2). This data collection aimed at investigating the interdomain routing in the region. Instead of running our measurement periodically, on the full timeline, or among the same set of probes, we launched them over random periods while collecting the data destined to assess the behavior of the involved networks. There are three other reasons for this choice: first, given the low quality of service experienced by end-users in the region, continually performing measurements from the hosts' devices may have a negative impact on their Internet access. Second, the RIPE Atlas platform sets, for each user, a maximum number of measurements that we chose not to exceed[2] too often, unlike cases where we run full-mesh measurements such as *Meas1A*, *Meas2B* (Table 3.2). By doing so, we also avoid overloading the probed networks with our measurement packets. Third, massive loads of measurements consume RIPE Atlas credits at a faster rate than our probes gain them: therefore keeping them running for four years is impossible.

Our measurement campaigns consisted of full mesh paris-traceroutes between the sets of probes listed in the column "involved probes" of Table 3.2. We used paris-traceroute [24] for all our measurements to discover path diversity and to reduce the number of inconsistencies caused by load balancing when using classic traceroute [298]. The probes performed traceroutes with 16 different paris_id to prevent the outputs from leading to the discovery of inaccurate IP paths due to routers, which employ load balancing on the packet header fields. We used User Datagram Protocol (UDP) traceroute to reduce the potential bias caused by differentiated traffic handling of ICMP packets [59]. The outputs of our measurements are publicly available in a Technical Report [84].

Careful sanity-checking and cleaning of the collected raw data is an essential step in our analysis. Before filtering, our raw data involved 324 probes hosted in 169 ASes operating in 40 African countries. It also contained data collected from 626 probes hosted in 380 ASes in 8 EU countries (Belgium (BE), FR, Finland (FI), Ireland (IE), Germany (DE), Netherlands (NL), Sweden (SE), and Switzerland (CH)) as well as 329 probes hosted in 195 ASes operating in the US. The geographical spread of all those devices used during our measurement campaigns is depicted in Figure 3.7.

More specifically, Table 3.3 summarizes the geographical and networking spread of the probes in Africa used in our study. ASes in italics host probes that participated only in IPv6 measurements, and those in bold, probes used to perform in both IPv4 and IPv6 measurements. Moreover, we put Southern African countries in bold, while the ones in West Africa are in italics. We also add the symbol $\star$ to the names of countries in which operate ASes hosting our deployed probes.

The percentages of ASes and IPv4 prefixes covered per country are computed based on

[2] For being able to exceed the maximum, we requested the RIPE Atlas team to increase the number of concurrent measurements we were allowed to run daily and our maximum daily spending limit regarding measurement credits.

Table 3.2: Datasets collected as parts of this work during our measurements covering 2013 to 2016

Name (IP Type)	Involved probes	Period	Frequency	# Traceroutes (Valid Traceroutes) outputs	Coverage of valid Traceroutes outputs	Goal	Used in Sections
Meas1A (IPv4)	All IPv4 probes in Africa (AF)	30/11/2013 to 06/04/2014	$\approx 3h$	675,421 (593,087) IPv4		Investigate IPv4 inter-domain routing	
Meas1B (IPv4 & IPv6)	All probes in AF countries hosting IPv6-enabled probes	01/06/2013 to 01/08/2014	$\approx 3h$	408,383 (397,234) IPv4 21,744 (19,593) IPv6	238 IPv4 probes hosted in 136 ASes in 35 AF countries & 30 IPv6 probes hosted in 20 ASes in 11 AF countries	Compare IPv6 to IPv4 interdomain topology	
Meas1C (IPv4)	All IPv4 probes in Gambia	04/08/2014 to 10/08/2014	$\approx 1h$	3,161 (2,747) IPv4		Highlight the launch of SIXP in Gambia	
Meas2A (IPv4 & IPv6)	All probes in AF	07/11/2014 to 18/02/2015	every week	361,267 (313,268) IPv4 1,584 (970) IPv6		Update our data and track evolution	§3.2.1.2 §3.2.1.4
Meas2B (IPv4)	Only IPv4 probe in Liberia to online IPs in local ASes	04/08/2015 to 10/08/2015	$\approx 200s$	50,960 (45,978) IPv4		Highlight the launch of LIBERIA-IX (LR)	
Meas2C (IPv4)	Randomly selected probes in same EU countries (resp. US)	08/12/2014 to 23/02/2015	every week	257,508 (227,021) IPv4	599 (319) IPv4 probes in 373 (190) ASes in 8 EU countries (US)	Compare results within AF countries to those within EU ones and the US	
Meas2D (IPv4)	All probes in Madagascar (MG)	04/04/2016 to 04/08/2016	$\approx 200s$	361,344 (318,597) IPv4	11 IPv4 probes in 04 ASes in MG	Highlight the launch of MGIX (MG)	

Figure 3.7: Geographical spread of the RIPE Atlas probes used in all our 7 measurement campaigns [84].

AfriNIC allocations [11] to offer the reader a glimpse of the granularity of our results. On average 23.8 % of allocated ASes (and 47.6 % of allocated IPv4 blocks) are covered per country. While computing these percentages, we include IPv4 spaces of local operators whose ASes have been allocated by other RIRs.

Using the techniques described below, we first map IP addresses into CCs to infer the set of countries traversed by the packets on the forward path during each traceroute. Second, we map IP addresses into ASes to infer the AS sequences.

3.2.1.2.2. IP to CC Mapping Geolocation of Internet infrastructure is known to be of poor quality [105, 124, 225], especially for IP addresses located in Africa [153]. To geographically locate the 42,412 public IPv4 and 1,425 public IPv6 addresses found in the traceroute data as accurately as possible, we analyzed 10 public DSes that we cross-correlated with delay measurements, as explained in this section. We used the following DSes:

1. OpenIPMap (*OIM*) [245], which aims at obtaining city-level accuracy of Internet infrastructure by crowdsourcing this information from network operators and other interested parties. 25 contributors, mostly operators, currently participate in this effort.

2. Reverse DNS lookups (*RDNS*): we deduced geolocation from location information embedded in hostnames by network operators such as CCs, airport codes, or abbreviated city names. For instance, "xe-3-2-1.**was14**.ipv4.gtt.net." corresponds to a TINET (DE) router located in Washington (US), "if-4-1-2.core2.**COV-Cochin**.as6453.net", to a TATA (US) router located in Cochin (IN), while "be2321.ccr22.**ams03**.atlas.cogentco.com.", to a Co-

Table 3.3: Before filtering, ASes and involved probes per African country

CC	African Country (#Probes used)	ASes	%ASes	%IPv4 blocks
AO	**Angola**⋆ (3)	36907, 17400, 3741	5.1%	7.1%
BJ	*Benin*⋆ (18)	37090, 28683, 37292	33.3%	67.2%
BF	*Burkina Faso*⋆ (7)	25543, 37577, 37073	42.9%	84.7%
BW	**Botswana**⋆ (5)	**14988**, 37678, 37537	14.2%	73.2%
CI	*Ivory Coast*⋆ (6)	29571, 36974	15.4%	73.4%
CG	Congo⋆ (1)	37451, 37281	15.4%	9.9%
CM	Cameroon⋆ (6)	16637, 15964, 37475, 36905	23.5%	43.3%
DZ	Algeria (1)	36947	6.7%	88.1%
ET	Ethiopia⋆ (4)	24757	50%	33.3%
GA	Gabon (2)	16058	10%	74.2%
GH	*Ghana*⋆ (7)	30988, 29614, 37012, 37623, 37140, 37047	11.1%	61.6%
GM	*Gambia*⋆ (10)	37309, 37524, 327719, 37323, 25250, 37503	75%	96.2%
GQ	Equ. Guinea (1)	22351	N/A	N/A
KE	Kenya⋆ (15)	**37061**, 36914, 36866, 37406, 30844, 327748, 12556, 9129, 37662, 15808, 15399	18.3%	20.5%
LR	*Liberia*⋆ (1)	37557	16.67%	9.5%
LS	**Lesotho**⋆ (2)	37057, 3741	10%	43.6%
LY	Libya (1)	21003	12.5%	89.8%
MA	Morocco⋆ (3)	**30983**, 36925, 6713, 36884	50%	92%
MG	**Madagascar**⋆ (6)	37054, 21042, 37608	20%	62.9%
MR	Mauritania⋆ (1)	8657	25%	39%
MU	**Mauritius**⋆ (15)	**37708**, **37100**, 37662, 23889	13.3%	72.3%
MW	**Malawi**⋆ (4)	37098, 37187, 3741	16.7%	12.4%
MZ	**Mozambique**⋆ (6)	42235, 31960, 30619, *6939*	13.8%	10.4%
NA	**Namibia**⋆ (5)	36996, 33763, 36877	12.5%	44.2%
NE	*Niger*⋆ (4)	37205, 37385	28.6%	33.1%
NG	*Nigeria*⋆ (3)	30988, 36932, 30988	2%	0.6%
RE	**Reunion** (3)	**37002**, 3215, 49902	66.7%	43.7%
RW	Rwanda⋆ (8)	21174, 30844, 37006, 37228, 36934, 16637	30%	69.5%
SC	**Seychelles**⋆ (19)	**36958**, 36867, 37343, 36930	28.6%	1.7%
SD	Sudan⋆ (4)	**37197**, 33788	28.6%	6.8%
SN	*Senegal*⋆ (6)	8346, 37196	50%	76.2%
SS	South Sudan⋆ (1)	14938, 37406	14.3%	30.8%
SZ	**Swaziland**⋆ (1)	**3741**	11.1%	98.7%
TG	*Togo*⋆ (5)	30982, 24691	66.7%	95.2%
TN	Tunisia (11)	*2609*, 37492, 37705	21.4%	83.1%
TZ	**Tanzania**∗ (7)	37045, 36909, **37084**, 37126, 12143, 36930, 37182, 33765	13.3%	18.3%
UG	Uganda⋆ (6)	327687, 37063, 36997	8.8%	13.8%
ZA	**South-Africa**⋆ (109)	37542, 22355, 36874, 36877, 37519, 37457, 37199, 37199, 37315, 12258, **33762**, 29975, 37618, 327813, 37596, 32653, 327805, **3741**, 16637, 36982, **10474**, 37253, 37251, 36937, 327750, **37105**, 22351, **6083**, **2018**, *6939*, 3491, 37519, 37312, 37266, *18931*, 22690, 5713, 37403, 11845, 37497, 37100, 37358, 37403, 37172, **37520**, **327817**	19.7%	46.5%
ZM	**Zambia**⋆ (3)	**37043**, 37154, 30844	16.7%	60.1%
ZW	**Zimbabwe**⋆ (5)	37204, 30969, 36986, **30844**	11.1%	43.1%

gent (US) router located in Amsterdam (NL).

3. MaxMind GeoIP2City (*MM*) [187] is a well-known geolocation database often used in applications for end-user geolocation (*e.g.*, credit card fraud detection). Therefore, it is

most accurate for geolocating end-user IP addresses and far less accurate for router IP addresses that we see in traceroutes.

4. Team Cymru (*TC*) [286], whose data is obtained directly from the RIRs.

5. RIR delegated files: RIRs report their allocations and assignments in so-called *delegated* files that are publicly available [11, 20, 23, 160, 252]. We collected these delegated files up to July 03, 2016.

6. RIR Databases (widely known as *WHOIS*).

3.2.1.3. Data analysis

Our mechanism to map an IP address to its CC can be described as follows: when all Data Sources (DSes) providing an entry for an IP return the same CC, we retain it for that IP. Next, we use a latency-based method to resolve instances of inconsistency among the DS entries. We launch three sets of ping measurements towards each IP from up to 10 random RIPE Atlas probes located in each country returned by the DSes.[3] For each group of probes per country, we then compute the minimum delay measured and used the CC for which the minimum delay is the lowest. To evaluate the accuracy of public and commercial geolocation databases, the checks of the consistency of country-level resolution by a given database against the majority answers and the calibration of the IP geolocation against measured RTTs have been adopted, among other techniques, by Huffaker *et al.* [124] and recently Gharaibeh *et al.* [105]. However, our set of retained databases differs from theirs in that we have only used publicly available DSes. In Table 3.4, we compare those selected DSes. The column "Coverage" corresponds to the percentage of IP addresses in our dataset for which the DS provides a valid country field.[4] The column "Trust" corresponds to the percentage of IP addresses in our measurement outputs for which the DS entry is equal to the country that is finally selected for that IP address.

Table 3.4: Comparison of geolocation data sources.

DSes	**IPv4 entries**		**IPv6 entries**	
	Coverage	Trust	Coverage	Trust
OIM	27 %	98.2 %	36.2 %	96.1 %
RDNS	42.7 %	94.7 %	49.4 %	90.7 %
MM	89.7 %	85.8 %	92.9 %	59.2 %
TC	90.5 %	83.7 %	100 %	52.3 %
AF	16.4 %	92.1 %	38.5 %	75.8 %
RI	28.6 %	79.3 %	22.2 %	87.1 %
AR	35.8 %	87.4 %	26.5 %	29.8 %
AP	0.84 %	86.9 %	0.1 %	0 %
LAC	0.002 %	100 %	0 %	0 %
WHOIS	94.6 %	46.5 %	33.7 %	24.1 %

[3] The raw data for these delay measurements can be found in [84].
[4] Any entry of these DSes, which is not a valid CC is ignored ("EU", "AP", "ZZ", "A1", "A2", etc.)

15,412 IPv4 (resp. 472 IPv6) addresses out of the 42,412 IPv4 (resp. 1,425 IPv6) addresses had a consistent CC among all DSes for which a valid entry was available. Our delay-based method to resolve inconsistent answers was then applied to the rest of the IP addresses: it allowed us to geolocate all IP addresses that respond to our pings. That is for 18,603 IPv4 (resp. 766 IPv6) addresses, we could deduce the country by using the delay based technique. At the end of this process, 80.2 % IPv4 (resp. 86.9 % IPv6) addresses in our dataset are associated with a location. The rest of the IP addresses corresponds to either offline IP addresses, *i.e.*, IP addresses which did not reply to our pings, or cases in which there was no RIPE Atlas probe in one or more possible countries given by the DSes: hence, they are not geolocated.

With the obtained geolocation data, we can compute the country path corresponding to the IP path of each traceroute output, as defined by [153].

3.2.1.3.1. IP to AS Lookup and Raw Data Sanity Check We first map, using *Team Cymru* (*TC*), public IP addresses of our traceroute data into ASes. We then apply the following filtering procedure: we keep traceroutes for which the obtained AS Sequence contains source and destination ASes corresponding to the ASes which are known to host the probes. Next, we try and complete remaining path ends based on learned AS adjacencies from this first check: for each non-valid AS sequence, we check if the first AS on the path is a known direct upstream of the source, or the last AS on the path, a known direct downstream of the destination, as observed in the previous set of traceroutes. If these checks succeed, we keep the traceroute as well. However, we only use this second set of inferred AS sequences for AS path analysis and exclude them from our RTT analysis.

To give ourselves the means to later assess the accuracy of the inferred AS paths, we keep track of intermediate traceroute hops for which the IP address has no entry in *TC*, or for which we did not receive a reply [119]. We respectively refer to them as *unresolved* and *unknown* ASes. We then compress AS paths into AS sequences. *Unresolved* or *unknown* hops found between two resolved hops of the same given AS are considered as belonging to that AS. Consecutive equal AS numbers are compressed into a single AS hop. We only infer an edge between two ASes if there are no *unresolved* or *unknown* hops in the IP path, and if both ASes are consecutive in the AS sequence. We identify 4,648 traceroutes with inferred AS path loops in the valid outputs of *Meas1A*, 1,419 traceroutes with inferred AS paths loops in those of *Meas1B*, and 1,195 inferred AS paths with loops among the valid IPv4 traceroutes collected during *Meas2A*. Since these paths are a small fraction of the total dataset, we filter them out. Note that we find no AS path with loops within the valid AS paths of *Meas1C*, *Meas2C*, and the valid IPv6 paths of *Meas2A*.

By the end of this raw data cleaning method, 87.8 % of IPv4 traceroutes are retained for *Meas1A*, while 97.3 % of IPv4 traceroutes and 90.1% of IPv6 traceroutes are retained for *Meas1B*. In the meantime, 86.9 % of IPv4 traceroutes outputs are selected for *Meas1C*. We keep 90.2 % and 88.2 % of IPv4 traceroutes outputs for *Meas2B* and *Meas2C* respectively. The corresponding total numbers for all the sub-campaigns are listed in Table 3.2.

The dataset resulting from this filtering process comprises paris-traceroutes outputs from 243 probes located in 35 African countries (covering over 60 % of Africa) hosted in 138 ASes, 599 probes hosted in 373 ASes in 8 European countries, and 319 probes hosted in 190 ASes operating in the US. Moreover, the filtered dataset involves 10,689 IPv4 AS pairs and 224 IPv6 AS pairs in Africa, 33,886 IPv4 AS pairs in Europe (EU), and 31,687 IPv4 AS pairs in the US. Furthermore, we could collect among them in total 27,481 unique IPv4 and 433 IPv6 AS paths within Africa, 38,326 IPv4 within EU, and 36,978 IPv4 AS paths in the US.

Finally, we estimate the RTT between each source and destination AS (denoted *RTT between ASes*) as the difference between the RTT from the source probe to the ingress point of the destination AS, and the RTT from the source probe to the egress point of the source AS. We also estimate the corresponding *RTT between probe IPs* as the end-to-end delay between the probes source and destination of the considered paris-traceroute measurement.

3.2.1.3.2. IXP detection We explain in this section the process followed to detect IXPs in the dataset. To begin with, we built a complete list of IXPs by collecting IXP information (ASes, prefixes, peers, IP addressing of the IXP) available in African IXP websites, Euro-IX, PeeringDB, PCH, IXP toolkit, Telegeography Internet Exchange Map, CAIDA AS relationships dataset [41, 218, 220]. After that, we ran the "*WHOIS*" command for the subnets in AfriNIC IXP blocks (196.49.0.0/16, 196.216.0.0/16, and 196.223.0.0/16) and extracted the corresponding prefixes, organizations names, and CCs from the outputs. In the rest of this thesis, we term the information obtained above *IXPs public datasets*. It involves IXPs of all regions (Africa, EU, North America (NAm), South America, Middle East, Australia, Asia Pacific) contained in those datasets. By *IXP-AS*, we refer to the ASN allocated by an RIR to an IXP platform.

Since mapping an IXP highly depends on the location of the probes used for measurements [25], the next step consists of checking whether our VPs were present in the networks of some IXP members. It appears that, in the AfriNIC region, 13 of the 37 IXPs (Table 3.5) under study had no member hosting a probe. We then apply the following techniques to detect IXPs in our traceroutes:

***Method1 (M1)*—IXP prefix search in IP paths** We consider the IP paths collected for any given pair of ASes in all our measurement campaigns. If any of those paths are via IP prefix in the same subnet as those assigned to an IXP (of Africa, EU, North America, or Asia), we deduce the IP Path is traversing the considered IXP.

***Method2 (M2)*—Tracking the launch of an IXP** With this method, we confirm the existence of an IXP whose prefix is not known (*e.g.*, IXPs using RFC1918 address space), by proving its launch based on the collected data. It consists of showing that during a certain period, the length of the AS path among its peers is higher than 3 and suddenly becomes 2 till the end of the measurement period, with delays considerably reduced. Towards this end, we track AS paths

length and substantial delay drops between networks operating within the same African country, the same sub-region, or the region by observing the evolution of these metrics over the measurement campaigns. In practice, we compute for all pairs of African ASes in our dataset the ratio of the average RTT between ASes collected from the first 25 % of traceroutes outputs to that of the last 25 %. If this ratio is greater than or equal to 2, we check if simultaneously to the drop of the measured RTTs, the most common AS path length drops to 2 as well. Note, we do not deduce the detection of an IXP with *M2* unless we find 3 or more peers and the RFC1918 address space traversed in all cases is the same.

3.2.1.3.3. Are IXPs prefixes routed on Internet? To investigate whether IXP prefixes are routed on the Internet, we pinged all the IP addresses in the ranges of the IXP prefixes from machines whose addresses belong to routed prefixes on Internet. These measurements were launched (*i*) three times from July 16, 2015, to July 25, 2015, from a single location in Spain and (*ii*) three times from December 28, 2015, to January 03, 2016, from a unique location in the US. Next, we performed DNS lookups of the online IP addresses, based on which we deduce (if possible) the IXP members from ISP names embedded into the corresponding hostnames.

3.2.1.3.4. Technical description The collected datasets are parsed from our measurement campaigns (from a JSON format) into a MySQL database structured according to a well-defined format. Among others, this database stores per measurement campaign all information related to each traceroute, the IPs geolocation, the results of the IP to CC mapping, those of the IP to AS mapping, or the IXP detection. Our computation scripts are all written in the Python programming language and run queries over the MySQL database. Their outputs are stored either in the database or text files. They are then used as inputs of our Matlab (.fig) plots, Pyplots, or R scripts for plotting the graphs included in this section. We release our measurement results under the format of an online application freely accessible with interfaces showing statistics on the African interdomain in [256].

3.2.1.4. Results

In this section, we first examine the limitations of our dataset, before comparing it to the view of the African topology that can be made from public BGP data. We then highlight the remaining dominance of ISPs based outside Africa to provide interdomain connectivity between studied ASes, except those in ZA. Socio-economic patterns are also discovered. After that, we illustrate the impact of the intercontinental aspect of paths on the RTTs among African ISPs. Next, we evaluate inter-ISPs communications performance within African countries, European countries, and the US. Following that, we map African IXPs in our traceroute dataset and successfully infer 62.2 % of the existing IXPs. We detail the inference of Seychelles-IX (SC) and SIXP (GM), and exhibit, as case studies, the launch of BENIN-IX (BJ), LIBERIA-IX (LR), and Madagascar-IX (MG). Finally, we inspect and compare how frequently an AS path among ISPs of our dataset

operating on each continent traverses a local IXP.

3.2.1.4.1. Dataset limitations We first acknowledge that the RIPE Atlas infrastructure continuously evolves since probe deployment increases steadily. This evolution led us to add new probes to the set of probes that we use on a daily basis. Moreover, not all the probes are online and usable all the time, due to downtime. For diverse reasons detailed in Section 3.2.1.2, we adopted both full mesh measurements in Africa and measurements among subsets of probes in the same countries in Africa, Europe, and the US (Table 3.2).

Although the probes used in the African region are deployed in 60 % of countries, our dataset covers in total 13.3 % of the ASes, and 43 % of the IPv4 ranges allocated by AfriNIC. The coverage per delegated IP range is summarized in Table 3.3. At last, we acknowledge the shortcomings of IP to AS mapping. As an example, 36.4 % of the unique IPv4 AS paths between probes in Africa, 39.7 % of those between probes in the same EU countries, and 58.5 % of those between probes in the US contain at least one either *unknown* or *unresolved* AS, as defined in Section 3.2.1.2. One of the implications for this work is that we thus excluded AS paths which contain any *unknown* or *unresolved* AS to accurately evaluate the various distributions of AS paths length (Section 3.2.1.4.3). In contrast, the evaluation of our metrics related to end-to-end delays, such as RTTs between probe IPs or ASes were not affected. Further, we are aware that off-path IP addresses can cause false AS path inferences [166] and we acknowledge that including this so-called "3rd party address" problem exemplified in [184, 271] remains an open challenge.

3.2.1.4.2. Dataset completeness We validate our dataset (notably the AS paths inferred from all measurement campaigns except *Meas2D*) against data extracted from Routeviews, RIPE RIS, and PCH [189, 213, 253]. Although the results of this comparison are the only ones, which were not updated to include paths from *Meas2D* and were removed from the final version of our journal paper [85] notably for concision, we present them in the three subsequent paragraphs for the sake of completeness. We also note their similarity with those shown in our conference paper [81].

We extracted from the data collected from January 2013 to July 2015 by the route-collectors available in the African region, JINX and KIXP [189], all AS paths containing any of the African ASes hosting a probe used in this study. We then split those 62,312 distinct AS paths into 191,257 AS path fragments. The term *AS path fragments* refers to subgroups of the total AS path with ordering preserved. They are of minimum length 2 and maximum length, the length of the AS path. We later break the AS paths inferred from our traceroute dataset into 42,263 AS paths fragments, excluding all those containing an *unresolved* or *unknown* AS.

Our dataset is more precise when it comes to end-to-end African paths: of the 37,776 AS adjacencies that we inferred from the discovered paths, 99 % are not visible in these public datasets. Note that most of the AS adjacencies found in both datasets are between ASes based outside the continent. Quite intuitively, entire African AS paths fragments, *i.e.*, 190,726 are not visible in RouteViews.

In addition, we extracted from the data collected by existing PCH route-collectors for the same period [213, 214] 2,425 distinct paths, which we split into 3,492 AS path fragments. We then compare these paths with our set of AS path fragments previously extracted from all discovered AS paths, which contain neither *unknown* nor *unresolved* ASes. Unsurprisingly, we only find 8.9 % AS paths fragments in both datasets (Figure 3.8).

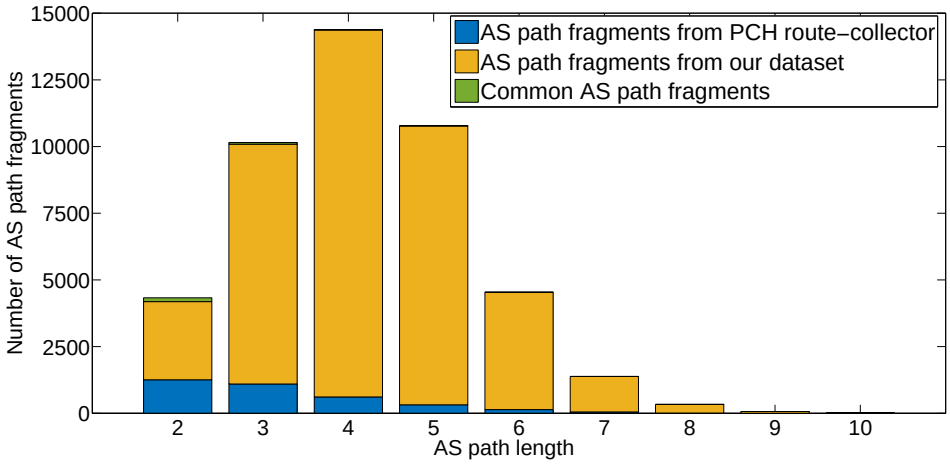

Figure 3.8: Comparison of AS paths of various lengths extracted from our dataset with those extracted (in the period 2013 - 2015) from PCH route-collectors deployed at African IXPs.

3.2.1.4.3. AS path length distribution We study the distribution of the length of AS sequences among pairs of ASes operating in Africa. We notably take a perspective focused on the sub-regions *WAf*, *SAf*, *EAf*, and on ZA. We separate IPv4 from IPv6 paths to highlight differentiated trends. We also carry out a specific analysis for pairs of ASes located within the same country. Moreover, we compute AS path distributions within EU countries and the US for comparison. As already mentioned, we only consider the set of paths containing neither *unknown* nor *unresolved* ASes for plotting the graphs of Figures 3.9 and 3.10. Thus, the AS paths in those cases could be even longer than what is presented.

In Figure 3.9e, we show the AS path length distribution for all the intra-African paths of the dataset. Since ASes in *WAf* are based in geographically collocated countries, one could presume that paths would be shorter. However, given the specific view provided in Figure 3.9a, we discover unusually long AS paths of five ASes on average in West African communications. It is worth noting that we find a higher proportion of national paths going through only three intermediate ASes than in [81] (Figure 3.9c). This could be explained by the discovery of new AS paths during *Meas2A*. They connect, for example, Connecteo with Onatel in BF, Sonitel with Atlantique Telecom in NE, AFRICELL-GM with Unique-Solutions in GM, and GHANATEL-AS with InternetSolutions in GH.

Figures 3.9b, 3.9d, and 3.9f highlight that short paths tend to be found in *SAf*, and precisely in ZA for which the set of AS path lengths has a mode of 3. Further, paths between ASes operating

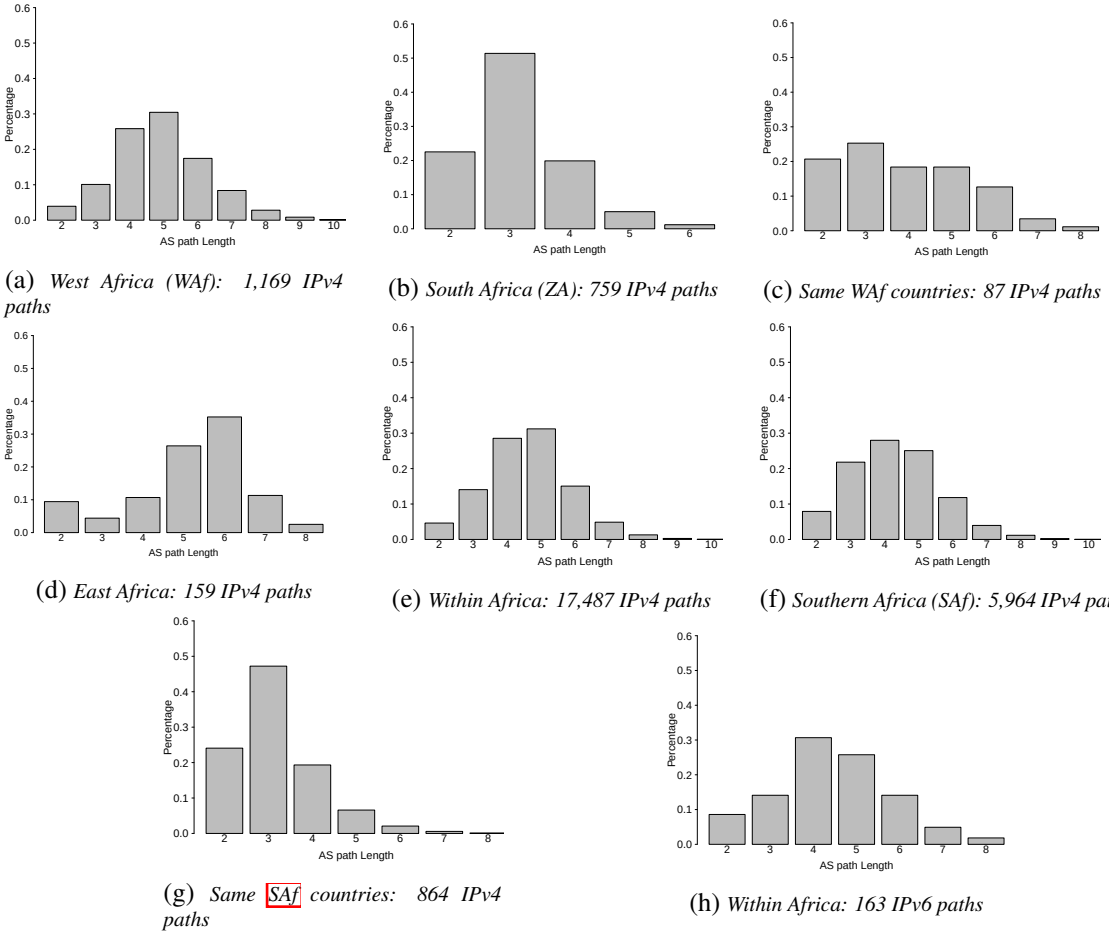

(a) *West Africa (WAf): 1,169 IPv4 paths*

(b) *South Africa (ZA): 759 IPv4 paths*

(c) *Same WAf countries: 87 IPv4 paths*

(d) *East Africa: 159 IPv4 paths*

(e) *Within Africa: 17,487 IPv4 paths*

(f) *Southern Africa (SAf): 5,964 IPv4 paths*

(g) *Same SAf countries: 864 IPv4 paths*

(h) *Within Africa: 163 IPv6 paths*

Figure 3.9: Path length distributions for all (IPv4 & IPv6) AS paths within Africa and for some African sub-regions

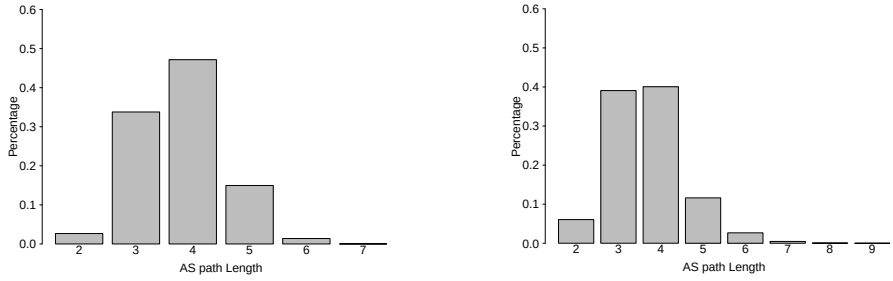

(a) *Within the US: 22,285 IPv4 paths*

(b) *Within same EU countries: 15,917 IPv4 paths*

Figure 3.10: Path length distributions for IPv4 paths within involved European countries and US

in the same country (Figures 3.9b, 3.9c, and 3.9g) are much shorter in ZA than in *WAf*. IPv6 AS paths, of which 77 % are observed in *SAf*, tend to be short, reflecting similar peering and localized transit habits as for IPv4 in the region (Figure 3.9h). These observations confirm that focusing solely on measurements from ZA like Gilmore *et al.* [106] does not provide a representative sample of Internet path characteristics for the rest of Africa.

Figures 3.10a and 3.10b present the distribution of the length of AS paths between ASes operating within the US and involved EU countries. AS paths inferred from our measurements between probes in the US never exceed a length of 7, while those inferred from data collected in EU countries attain a maximum length of 9. Similarly to same *SAf* countries, and contrary to same *WAf* and *EAf* countries (Figure 3.9), both have a mode of 4. These results highlight two key points. To provide end-users in the African region a better connectivity and an improved QoS for intra-African communications, it is essential to (*i*) shift the average AS path length within *WAf* and *EAf* sub-regions and thus, within Africa (Figure 3.9e) to 4; in other words, about 70 % of the AS paths within Africa should have a length below 5. Further, (*ii*) local ISPs should be encouraged to never exceed an AS path length of 7, in the worse case, for the communications between them.

3.2.1.4.4. Trends in African Interdomain Routing

AS-Centrality We now study the role of transit played by each ISP found within the AS paths extracted from our dataset. To this end, we define the *"AS-centrality of an AS"* as the percentage of observed paths containing that AS, but for which the said AS is neither the source nor the destination. We only account for presence within AS paths among pairs of ASes, radically diverging from betweenness centrality [199] in the AS graph. We then define the concept of *"joint AS-centrality"*, which captures the centrality of tuples of ASes present together on AS paths.

To provide insights into the African sub-regions, we classify the 255 ASes of our dataset into five categories, depending on their sub-region of operation. The category *WAf* ASes gathers ASes based in West Africa; *SAf* ASes are based in Southern Africa, while *EAf* ASes are those based in East Africa. *RAf* ASes are ASes operating in Africa but in none of the previous regions, while the category *Intercontinental* (*Int*) ASes gathers all ASes based outside the continent. An AS belongs to the sub-region in which are geolocated most of the IP addresses allocated by its RIR. Any AS having a significant amount of IP addresses located on more than one continent is classified in the *Int* category: 87 *Int* ASes are found in the dataset.

Figure 3.11 depicts the AS-centrality of each AS, in the whole set of paths (blue curve), among *WAf* networks (orange curve), and among *SAf* networks (black curve). We sort the ASes according to their centrality on the whole set of paths and represent them with different markers given the category to which they belong. In total, 168 ASes have an AS-centrality value greater than 0. We only plot those that play a non-negligible role of transit in Africa, *i.e.*, their AS-centrality is greater than the threshold 0.7 %, leaving 98 ASes out.

As indicated by the blue curve, the four most central ASes in our view of the IPv4 African interdomain topology are all intercontinental ones, namely TATA (US) with 23.5 % of the 27,481 AS paths, Level3 with 20.6 %, Cogent (US), 16.6 %, and France Telecom-Orange (FR)[5], 10.6 %. 54.9 % of the AS pairs are served using at least one of these four ASes. The

[5] Note that the ASN of France Telecom-Orange is AS5511. The name of that AS, which was formerly **France**

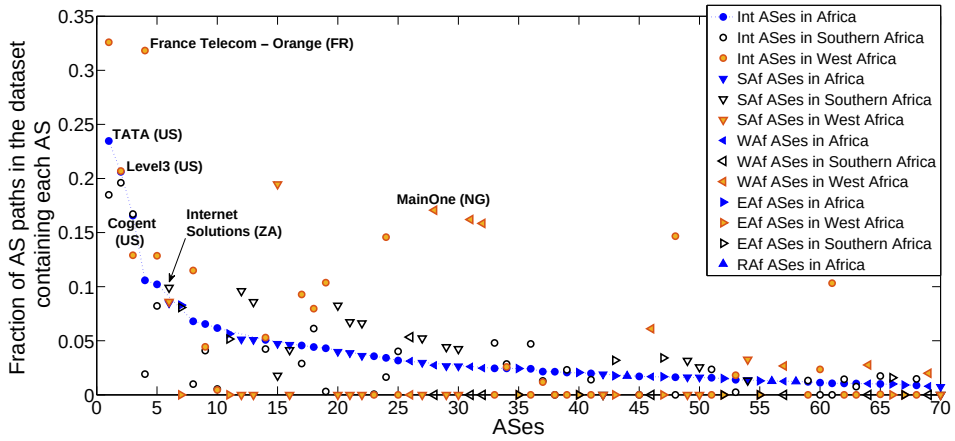

Figure 3.11: AS-centrality. ASes are sorted according to their AS centrality within the African interdomain topology (blue curve).

most central African AS, InternetSolutions, has an AS-centrality of 8.5 %. By contrast, France Telecom-Orange becomes the second dominating ISP with an AS-centrality of 31.8 % when it comes to paths between ASes in the *WAf* category (orange curve). Meanwhile, TATA and Level3 play the role of transit on respectively 32.6 % and 20.7 % of the AS paths. We also notice that a relevant percentage of paths (19.5 %) connecting *WAf* ASes transit via MTN (ZA). The most central local AS is MainOne, found in 17.1 % of the paths.

By contrast, the reliance on *Int* transit providers is lower within the *SAf* sub-region. In fact, the top three ASes remain Level3 with 19.6 %, TATA (18.5 %), and Cogent (16.7 %), but InternetSolutions, SAIX-NET (a private IXP owned by Telkom SA) and MWEB follow with 9.9 %, 9.6 %, and 8.2 % respectively. *SAf* ASes appear to benefit from diversity in their transit offerings and resort a lot to peering. Note that the reliance of *SAf* ASes on ISPs based in other African regions is insignificant.

Some ASes, which are not relevant for IPv4 routing, show a high AS-centrality for IPv6 routing. The top two ASes in IPv6 are Hurricane Electric (US) with 28.8 % and TENET (ZA) with 22.9 %. They are followed by TATA (18.7 %), Cogent (16.6 %), and Liquid Telecom (AS30844, UK) with 16.6 %.

Techno-Economic Insights on Routing Trends We have also appreciated in our measurements how some techno-economical factors affect transit trends. To give a glimpse of such facts, we present in Figure 3.12 the AS-centrality of TATA, Level3, and France Telecom-Orange, discussing whether these ASes jointly serve a path or are lying on a path on their own. From the left, the first three triplets of barplots are based on all the paths of the dataset, while the last triplet focuses on the *WAf* category. We use for that graph the color circle: TATA, Level3, and France Telecom-Orange correspond to the primary colors blue, red, and yellow respectively. When they

Telecom, is currently **Orange**. In this doctoral thesis, we term this AS **France Telecom-Orange** to better highlight the techno-economic insights discovered in our analysis.

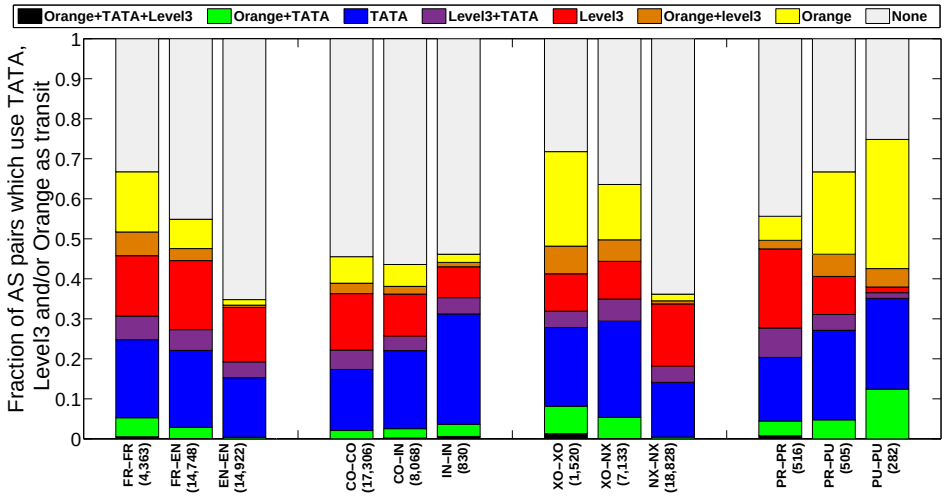

Figure 3.12: Joint AS-centrality of AS3356 (Level3), AS6453 (TATA), and AS5511 (France Telecom-Orange) for paths among various categories of ASes.

all appear on a path, that path is classified in the category for which the three colors are mixed (black). If none of them is found, the path is classified in the category colored in grey. All the other colors are obtained by mixing the two primary colors of the corresponding ASes.

In Figure 3.12, FR means French-speaking countries, EN English-speaking countries, CO Coastal countries, and IN Inland countries; besides, XO stands for Countries in the XAF-XOF region, NX Countries not in the XAF-XOF region, PR Privately owned ASes, and PU stands for Publicly owned ASes. Figure 3.12 shows that ISPs in French-speaking countries mostly rely on France Telecom-Orange, which serves 15 % of the West African (*WAf*) AS pairs, without TATA or Level3. Another 11.2 % of AS pairs are served by France Telecom-Orange, but jointly with TATA or Level3. Note that these results, found in [81] and confirmed by our longitudinal study [85], have recently been validated by the findings of [94]: in fact, Formoso *et al.* [94] similarly noticed "a significant presence of the French operator, Orange" and precised that "17% of networks in French speaking countries utilise Orange, which add up to almost 40% of Orange's downstreams." By contrast, we find that when it comes to communications among English-speaking countries, France Telecom-Orange disappears from our African internetworking map.

Such diverse transit habits are also observed when classifying ASes according to the monetary region to which they belong. Within the XAF-XOF (CFA Franc) monetary union, France Telecom-Orange has alone a centrality of 23.6 % but is barely present (1.6 %) in the market of communications among ISPs operating in countries that do not belong to this union. From the same figure, we learn that France Telecom-Orange and TATA are together on 12.4 % of the paths among the publicly owned *WAf* ASes.[6] Further, France Telecom-Orange is alone on another 32.3 % of these paths. Meanwhile, few publicly owned operators (1.4 %) seem to get transit from only Level3. In the same sub-region, however, a relevant proportion of pairs of ASes (19.8 %) involving a privately owned AS are served via Level3. Finally, the second triplet of

[6] We categorized the *WAf* ASes as owned by a public or private company, based on gathered private information

barplots shows that pairs of ASes operating in African inland countries rely much more on TATA (35.8 %) than on Level3 (13.5 %), dominating France Telecom-Orange. Such differences can be explained by the scarcity of Internet transit offerings in inland countries, which mostly rely on Satellite transport companies that peer with Level3 and TATA. It is important to emphasize that as of this writing, only the studies [81, 85] have attempted and reported on the exercises mentioned above regarding techno-economic insights on routing trends in Africa.

3.2.1.4.5. Impact of transit localization on end-to-end delay Our objective in this section is to characterize, based on the observed RTT, the QoS experienced by communications following the diverse categories of IPv4 and IPv6 paths registered during our measurements. To achieve this, we first identify, per AS path among ISPs operating in Africa, the IP path over which the minimum RTT is observed as well as its corresponding country path. After that, we group AS paths into two categories. Continental AS paths (20.5 % of the AS paths in our dataset) are those for which the corresponding country paths only traverse African countries, and which thus stay within Africa. By contrast, intercontinental AS paths (79.5 % of the AS paths) traverse at least one node geolocated outside the continent.

Figure 3.13a shows the CDF of the minimum RTTs among our probes in Africa, comparing continental (IPv4/IPv6) AS paths to intercontinental ones. We notice, for instance, that continental IPv4 AS paths in our dataset have a median of 32.5 ms and an Interquartile Range (IQR)[7] of 97.9 ms, whereas intercontinental AS paths have a much higher median of 238.1 ms and an interquartile range of 168.5 ms. Also, we observe that approximately 75 % of continental IPv4 AS paths have a delay below 100 ms, while this is only 16 % for intercontinental AS paths. The results are similar for IPv6 AS paths. They highlight the severe consequences on performance among local ISPs induced by the adoption of intercontinental tromboning of local traffic.

Let us now examine Figure 3.13b, a boxplot of the minimum RTTs among our probes, on which boxes are ordered based on their median. The values in parenthesis on its x-axis correspond to the number of AS paths classified in the corresponding category. Continental AS paths with very low RTTs mostly correspond to paths between pairs of ASes based in the same country, or those traversing collocated regional ISPs. 82.3 % of such paths are through ZA, acting as a regional hub. The IPv4 paths not passing through ZA have a median of 27.5 ms with an IQR of 182.7 ms, while those traversing ZA have a higher median (33.1 ms) with a lower IQR (78.9 ms). Note, all the continental IPv6 paths traverse ZA; their median is 41.6 ms.

Slightly longer RTTs (50 ms – 150 ms) are seen among AS pairs from geographically distant countries. As an example, a path from a KE ISP to a ZA ISP, only served by African transit ISPs, shows a minimum RTT of 80 ms. A striking result comes from the presence of very long RTTs in paths that are categorized as continental ones. These IPv4 AS paths are typically those between *EAf* and *WAf* ISPs, which are served by ZA transit ISPs. The following long RTTs (> 2 s) are

[7] The interquartile range is a measure of statistical dispersion, which is resistant to the presence of outliers. It is computed as the difference between the first and the third quartile. In this case, it highlights how the RTT values are spread out around their median.

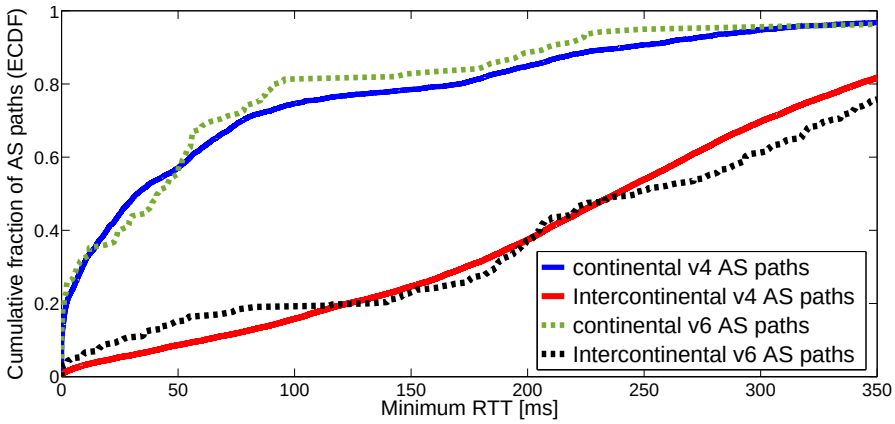

(a) *CDF of the minimum RTT between probes in Africa.*

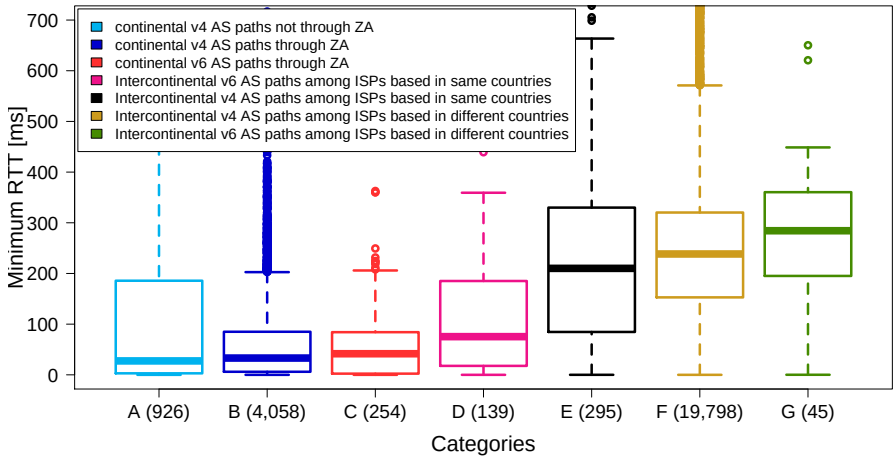

(b) *Boxplot of the minimum RTT between probes in Africa.*

Figure 3.13: Minimum RTT distribution over the AS paths between ISPs operating in Africa.

recorded on paths from TZ to ZA via SEACOM (MU), from InternetSolutions to Simbanet (TZ) via KE, or from SAIX-NET to TENET in ZA. They are having a mis-categorization issue, as per our manual checks, since their IP level traceroutes contain many non-answering hops. But we have no data allowing us to certify that they leave the continent.

Intercontinental paths with a low RTT (*i.e.*, < 100 ms) also reveal the weakness of geolocation. These AS paths contain *Int* ASes, as per *TC*, and have been consistently geolocated in either the UK, NL, FR, or the US by the DSes. These correspond to cases where all DSes are returning the same CC, located outside Africa, although latency measurements indicate that the IP address is located on the continent.

Nevertheless, most of the measured RTTs in this category reflect intercontinental transit of continental traffic, with an RTT of around 238.4 ms on average. 95.4 % of the paths with an RTT between 100 ms and 400 ms are through EU. AS paths with RTTs scattered around 750 ms are mostly from and towards ISPs, which are served by Satellite providers, routing traffic through another continent. A path in this group is, for example, from Connecteo in BF to AFNET in

CI, passing through SkyVision, Level3 (in New York), Level3/Global-Crossing (in London), and MTN (ZA). The paths measured with an RTT above 1 s are mostly those served via 2 satellite links. For instance, one is from Connecteo in BF to Sonitel in NE, going through the US and EU but arriving in NE via another satellite, provided by IntelSat. Finally, we highlight the RTTs between ISPs operating in the same African countries, exchanging packets over intercontinental AS paths. These are notably observed in BJ, CM, MU, MA, and MZ.

3.2.1.5. Mapping African IXPs in our collected dataset

We then conduct a detailed analysis of paths revealing the use of IXPs to exchange traffic. We began by building a complete list of IXPs and their information, as explained in Section 3.2.1.3.2. In this data termed *IXPs public datasets*, we found for the AfriNIC region 29 IXPs to which an ASN has been allocated as of August 2016. Their respective allocations dates are specified in Table 3.5 and vary between 2005 and 2015. Table 3.5 summarizes the information related to IXPs of the AfriNIC region. In this table, the symbol $\star$ follows the name of the IXPs for which no members hosted our probes during our measurements. Table 3.5 shows that PeeringDB and PCH public datasets are not up-to-date with regards to the number of peers at each IXP in Africa. Moreover, some IXP members do not register in those datasets or add their prefixes, while some IXPs (*e.g.*, Madagascar Internet eXchange (MGIX), DJIBOUTI-DC-IXP, ZINX, LIXP) do not have a website. Note that this is one of the issues tackled by the African Route-collectors Analyzer (ARDA), whose design and implementation are presented in Section 3.3.1: this freely accessible platform automatically profiles local IXPs and monitors in real-time their growth.

3.2.1.5.1. Mapped African IXPs By crossing the collected IXPs information with our traceroute outputs, we detect IPs used to address interfaces to these IXPs in our traceroute data. We map in our dataset a total of 23 African IXPs located in 16 countries (Table 3.5) thanks to method *M1*. These IXPs are CINX, JINX, KIXP, NAPAfrica (Johannesburg and Cape Town), SIXP, UIXP, TIX-ASN, MGIX, etc. Among them, five IXPs are recently established. With method *M2*, we can prove, for instance, the launch of BENIN-IX and MGIX (Section 3.2.1.5.2). Internet Exchanges BENIN-IX and SIXP are detected when using both *M1* and *M2*, since those IXPs first adopted an RFC1918 address space before acquiring their prefixes from AfriNIC.

We discover that 11 IXPs have their prefixes routable on the Internet. These correspond to prefixes allocated as either the IXP peering or IXP administration block. By the time we performed our checks, the IP addresses could not be resolved and nor did the collected Reverse DNS outputs contain the names of the IXP members using the corresponding interfaces. Instead, some DNS lookups outputs contained the name of the IXP (*e.g.*, SIXP and BENIN-IX). We informed those IXPs through ISOC so that this is corrected by the peers to avoid security attacks.

Table 3.6 presents a partial view of the IPv4 peering matrix of KIXP. It highlights the positive impacts of having each IXP member peering with all the others. We put in green minimum delays between two ASes (in ms) when they are present and peer at the IXP. In those cases, the AS path

Table 3.5: List of African IXPs [292] collected in public datasets as of December 31, 2016. N/A means "Non Available" and ?, "Unknown".

IXP-AS	IXP Name	CC	AS Allocation date	#IPs up (ping from Internet)	IXP mapped?	Mapping Methods	# Peers IXP Websites	#Peers Peering DB	#Peers PCH	#members (local) members found
N/A	DINX	ZA	N/A	0	Yes	M1	N/A	10	5	2 (2)
N/A	WHK-IX/IXP-NAMIBIA	NA	N/A	18	Yes	M1	N/A	4	5	4 (2)
N/A	REUNIX	RE	N/A	0	Yes	M1	N/A	5	10	3 (1)
N/A	KINIX /RDC-IX Kinshasa⋆	CD	N/A	0	No	N/A	N/A	1	6	N/A
N/A	SEYCHELLES IX	SC	N/A	0	Yes	M1	N/A	N/A	?	4 (4)
N/A	LIBERIA-IX	LR	N/A	0	Yes	M1	N/A	N/A	N/A	4 (4)
4558	KIXP	KE	2010-09-22	0	Yes	M1	27	11	29	18 (8)
24736	CAIX⋆	EG	2007-09-20	0	No	N/A	7	N/A	8	N/A
30997	GIXA-AS⋆	GH	2005-03-02	89	No	N/A	12	1	24	N/A
33791	TIX-ASN	TZ	2005-08-02	20	Yes	M1	27	11	25	9 (8)
36932	IXPN⋆	NG	2007-01-16	58	Yes	M1	33	7	30	4 (4)
36946	CIVIX	CI	2007-04-24	10	No	N/A	5	0	5	N/A
37143	ARUSHA-AS/AIXP	TZ	2009-09-02	0	Yes	M1	N/A	N/A	6	8 (8)
37186	NAPAFRICA	ZA	2010-03-15	0	Yes	M1	108	172	41	55 (30)
37195	NAPAFRICA	ZA	2010-04-01	0	Yes	M1	108	172	41	55 (30)
37221	LUSAKA-IXP/ZAMBIAIXP⋆	ZM	2010-06-14	0	Yes	M1	13	2	13	4 (3)
37224	RINEX c/o RICTA	RW	2010-06-18	0	Yes	M1	9	5	5	14 (4)
37228	RINEX c/o RICTA	RW	2010-06-18	0	Yes	M1	9	9	5	14 (4)
37299	LIXP⋆	LS	2011-03-08	0	No	N/A	2	0	?	N/A
37355	ZINX	ZW	2011-07-07	0	No	N/A	N/A	N/A	5	N/A
37383	ANG-IXP/ANGOLA-IXP	AO	2011-10-12	0	Yes	M1	12	6	10	2 (1)
37386	UIXP	UG	2011-10-18	23	Yes	M1	17	3	8	5 (4)
37481	CGIX⋆	CG	2012-07-16	0	No	N/A	6	0	6	N/A
37551	ATI-TUNIXP⋆	TN	2013-02-05	52	No	N/A	12	2	9	N/A
37635	MOZIX⋆	MZ	2013-10-04	0	No	N/A	17	3	17	N/A
37651	MIX-AS	MW	2013-11-06	0	Yes	M1	N/A	2	21	2 (2)
37695	BURUNDIX⋆	BI	2014-03-20	5	No	N/A	2	0	1	N/A
37699	JINX	ZA	2014-05-23	36	Yes	M1	57	53	57	48 (23)
37701	CINX	ZA	2014-05-23	21	Yes	M1	23	21	23	26 (15)
327719	SIXP	GM	2014-01-06	0	Yes	M1, M2	15	0	0	4 (4)
327740	AMSIX- East-Africa	KE	2009-07-24	0	Yes	M1	N/A	N/A	3	2 (2)
327775	AO-IXP⋆	AO	2014-07-07	6	No	N/A	N/A	N/A	?	N/A
327779	DJIBOUTI-DC-IXP (DjIX)⋆	DJ	2014-07-31	0	No	N/A	N/A	2	N/A	N/A
327788	ANGONIX⋆	AO	2014-09-01	0	No	N/A	4	3	3	N/A
327818	BENIN-IX	BJ	2014-11-12	53	Yes	M1, M2	5	2	5	3 (3)
327821	MIXP	MU	2014-11-24	0	No	N/A	N/A	N/A	6	N/A
327834	MGIX	MG	2015-01-29	0	Yes	M1	N/A	N/A	N/A	N/A

lengths are either 2 or 3. The latter all contain an *unknown* AS (corresponding to IP addresses belonging to KIXP prefix) surrounded by the peers ASes. Minimum RTTs in red correspond to cases in which both ASes (although present) do not exchange traffic via KIXP. In those cases, we add the AS path length between parenthesis. Further, low RTTs with a path length of 2 correspond to cases in which ASes have a private interconnection (*e.g.*, from KENET-AS to WANANCHI-KE or from JTL to WANANCHI-KE), or peer at another Internet Exchange (*e.g.*, from KENET-AS to Liquid Telecom at NAPAfrica). High RTTs correspond to cases in which both ASes transit via others to communicate (*e.g.*, from JTL to Liquid Telecom). Finally, N/A corresponds to cases in which we could not have any RTT value due to the absence of probe in one of the AS or to non-valid (and filtered) measurements between the corresponding ASes.

AS30844 (Liquid Telecom) is used by Gupta *et al.* as an example of a network that connects at JINX, and is present but does not peer at KIXP [117]. Nevertheless, our measurement campaigns *Meas1A*, *Meas1B*, and *Meas2A* running from 2013 to 2015 show that Liquid Telecom is present and peers at both IXPs (see Table 3.6 for details on KIXP peering) as well as at other ones (NAPAfrica, Lusaka-IXP, RINEX, and UIXP). At KIXP, however, Liquid Telecom has also been peering using the ASes of the networks they acquired: it is common for large networks to use a BGP confederations feature [69, 296] during network mergers and acquisitions before the implementation of the network strategy of new organizations.

Table 3.6: Partial KIXP IPv4 peering matrix extracted from our dataset. The minimum RTTs between ASes presented are in ms. Minimum RTTs are in red, followed by the AS path length in parentheses, when both ASes (although present at KIXP) do not exchange traffic via the IXP.

ASes	**36914**	**36866**	**15399**	**21280**	**12556**	**15808**	**30844**
36914		0.6	0.01 (2)	0.1	1.8	0.4	0.3
36866	0.7		1.8 (2)	0.4	0.3	0.001	165 (4)
15399	0.1 (2)	N/A		0.1 (3)	0.4	1.2 (2)	0.8 (3)
21280	0.1	0.9	0.1		0.2	0.7	0.8 (3)
12556	1.7	1.1	0.9	1.3		0.9 (2)	1.8
15808	0.8	N/A	0.1	0.6	0.1 (2)		0.8 (3)
30844	0.1	0.6	0.1 (2)	0.1 (3)	0.1	0.01	
9129	0.4	0.6	0.4	0.8	0.5	N/A	1.5 (3)

Moreover, we notice that over the same period only 6.8 % IPv4 (respectively 6.3 % IPv6) AS pairs in Africa have their RTTs dropped to a half of the initial values or more. For 0.4 % (resp. 59.8 %) IPv4 AS pairs, the AS path length has dropped to 2 (resp. 3), while this is only 0.4% (resp. 0.9%) for IPv6 AS pairs. After cross-checking with the IXP prefixes, we remark, for instance, that RTTs between SAIX-NET and InternetSolutions changed from 22 ms to 6.8 ms on average, since they peered at JINX. We also observe a drop of the RTTs from ISOCEL Telecom to Benin Telecom from 229.3 ms to 35.9 ms, which correspond to the period both ASes started peering at BENIN-IX (Section 3.2.1.5.2). Furthermore, we notice the drop of RTTs between SEACOM-AS and HABARI-CO-TZ-AS from 31.1 ms to 0.6 ms, from the period they peered at CINX. Finally, we detect the drop of RTTs between SIXP platform and GAMTEL in both directions (from 93.7 ms to 0.4 ms in one and 45.9 ms to 22.6 ms in another).

3.2.1.5.2. Emergence of recently established IXPs

Detection of Seychelles-IX and Serekunda Internet eXchange Point (SIXP, GM) Apart from the information available on public datasets and IXP websites, we were advised in August 2014 [190] that new IXPs were being deployed in BJ, SC, and GM [261]. We thus looked for and found those IXPs in our traceroutes. As a matter of fact, at the beginning of the 2nd campaign (*Meas1B*), four members of the IXP were hosting RIPE Atlas probes in SC. We could observe in the dataset a delay around 1 ms among each pair of this clique, formed by CWS-AS, ASIntelvision, Telecom Seychelles Ltd, and Kokonet-BGP.

In the data collected during *Meas1C*, the probes hosted in networks QCell, NetPage, and GAMTEL are connected to SIXP, with RTTs around 1.5 ms among QCell, NetPage, and the SIXP platform: there is a direct link in both directions between QCell and SIXP. Moreover, GAMTEL and SIXP appear within the path from SIXP to QCell. These hint at the fact that GAMTEL is also a peering partner. Indeed, to peer at the IXP, both GAMTEL and QCell use the IXP address space, while the BGP peering is set up either directly or via a route-server. The IXP typically holds the AS that announces the IXP address space, which causes its appearance in between the peers at that IXP when using IP to AS mapping with Team Cymru (*TC*) [286] (or data from RIPE RIS [253]). In the meantime, our measurements show a direct link from GAMTEL to SIXP. Besides, paths from GAMTEL to QCell all contain the IXP-AS. These prove the success of SIXP, the IXP of Gambia (GM), recently launched by the time of these findings.

However, not only RTTs between GAMTEL and QCell but also those between GAMTEL and SIXP fluctuate between low (0.9 ms) and high values (460 ms) with a median (and mean) of 14.4 ms (56 ms) and 8.9 ms (40.1 ms) respectively. After comparison with measurements performed between NetPage and QCell (0.04 ms – 18.9 ms), we deduce that the link from GAMTEL to the IXP platform is unstable and responsible for such delays.

We then learn that during *Meas2A*, which lasted from November 2014 to February 2015, RTTs among GAMTEL and QCell dropped to a set of values with a median of 0.9 ms (1.1 ms on average); likewise, the corresponding AS sequences have a length of 3, and the IP sequences traverse the IXP. These significant improvements in the RTTs highlight the correction by the peers of the previously mentioned shortcomings. In the outputs of *Meas2A*, we also find two others SIXP members (AFRICELL-GM and Unique-Solutions). Table 3.7, drawn following the same rules as Table 3.6, presents the SIXP IPv4 peering matrix extracted from our measurements. The number of discovered IXP members depends on the number of local ASes hosting our probes. Table 3.7 shows that most SIXP members were peering with one another during the measurements campaign: the AS paths often have a length of 3 (with an *unknown AS*), and the IP paths pass via the IXP platform. This has a positive impact on the minimum RTTs among any two of them, as those delays are low (0.001 ms – 6.6 ms).

Nonetheless, AS paths from AFRICELL-GM and Unique-Solutions to the IXP platform often traverse their respective transit ASes, leading to high minimum RTT values (44.9 ms – 55.7 ms).

In other words, if the peers are sharing resources hosted at the IXP platform, they will still have to pay transit fees when accessing them, although they are peering locally. Moreover, a properly configured IXP should not have its AS number in the AS path attribute. The fact that it is visible means the IXP is using a route-server, which does not support transparent AS feature. This situation, which makes the AS path appear longer than it is, should be corrected by the peers.

Table 3.7: SIXP IPv4 peering matrix extracted from our dataset.

ASes	37309	37524	37323	25250	37503	327719
37309		2.6 (3)	N/A	1.2 (3)	0.3 (3)	0.001 (2)
37524	1.1 (3)		N/A	6.6 (3)	1.5 (3)	55.7 (7)
37323	0.003 (3)	N/A		N/A	N/A	N/A
25250	0.8 (3)	0.03 (3)	N/A		1.1 (3)	0.3 (2)
37503	0.3 (3)	2.4 (3)	N/A	1.5 (2)		44.9 (4)
327719	0.01 (2)	N/A	N/A	0.3 (2)	45.8 (2)	

On the launch of BENIN-IX (BJ) The launch of BENIN-IX [28] in the period of our measurements gave us the opportunity to measure its impact on communications among its different members (Benin Telecom, ISOCEL Telecom, and OTI Telecom). While examining the dataset of *Meas1A*, we find that RTTs measured between the above-listed ASes considerably drop from a median of 326.5 ms (314 ms on average) between November 30, 2013, and December 19, 2013, to a median of 22.1 ms (42 ms on average) from December 20 to April 6, 2014. According to the traceroute data, those two ASes started peering on December 20, 2014. Figure 3.14 illustrates the benefit brought by this IXP for end-users and ISPs, depicting RTTs among two of its members and the length of the measured AS sequences. The figure also shows that our probes lost Internet connectivity during the establishment of the IXP, as very few traceroutes succeeded during that period (December 20 – 30, 2014).

A positive outcome of our longitudinal study of the African interdomain routing is that we could observe the dynamics of the AS paths over time. As an example, we later notice that during *Meas2A*, the AS path length was fluctuating from time to time between 4 (when the AS sequence traverses Cogent or Tinet SpA and France Telecom-Orange) and 3 (when the two ASes peer via the IXP platform) in both directions. According to our checks, this instability of the delay does not depend on the IP addresses (source or destination) of the probes and hence, is not due to misconfigurations while advertising the networks of the peers on the BGP sessions. We deduce that a possible cause is the instability of the interdomain link between the peers and that such a situation needs to be corrected by a careful check of the routers configurations and the introduction of redundant connections between peers.

On the launch of LIBERIA-IX LIBERIA-IX was expected to be launched by local networks in August 2015. To measure the impacts, we planned to perform paris-traceroute and ping measurement campaigns between hosts in local ASes. Despite our attempts to previously deploy RIPE Atlas probes in that country, only one RIPE Atlas probe was online. We, therefore, scanned

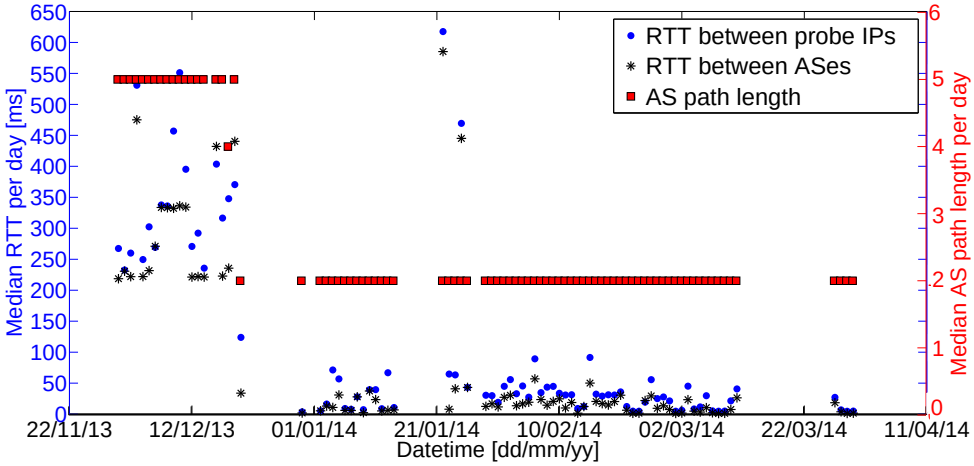

Figure 3.14: RTTs between probes in AS28683 (Benin Telecom) and AS37090 (ISOCEL Telecom) during BENIN-IX (BJ) establishment.

all the IP ranges assigned to Liberia (LR) by AfriNIC and randomly selected online IP addresses in each local AS. We then launched our measurements from the only available probe towards those IP addresses roughly every 200 s (*Meas2C*) from December 2014 to February 2015.

Figure 3.15 highlights the impact of the launch of LIBERIA-IX by depicting RTTs among our probe hosted in NOVAFONE (LR) and carefully selected online IP addresses in local networks LONESTAR, CELLCOM, and LIBTELCO. Our analysis of the collected dataset reveals the following.

At the beginning of our measurements, NOVAFONE had only one upstream: France Telecom-Orange. The upstream of LIBTELCO was Cogent, while that of LONESTAR was MTN. In contrast, CELLCOM was multihomed and served by Cogent, Belgacom, and DiViNetworks LTD. All networks were transiting for exchanging communications among local networks. Consequently, the set of AS paths collected for communications from NOVAFONE to LONESTAR had a median of 5 (via France Telecom-Orange, Cogent, and MTN). The median of the AS paths from NOVAFONE to LIBTELCO was 4 (via France Telecom-Orange and Cogent), while that of AS paths from NOVAFONE to CELLCOM was 5 (via France Telecom-Orange, NTT, and Cogent). While such routing policies were applied, the corresponding medians of the measured RTTs values between those ASes (respectively probe IP addresses) were 244.1 ms (resp. 248.1 ms), 238.4 ms (resp. 240.4 ms), and 131.9 ms (resp. 133.9 ms). In the meantime, the average RTT between ASes (resp. IP addresses) was 248.1 ms (resp. 254.9 ms) to LONESTAR, 248.5 ms (resp. 250.7 ms) to LIBTELCO, and 157.3 ms (resp. 165.2 ms) to CELLCOM (see Figure 3.15).

The peering session between NOVAFONE and LIBTELCO was then established earlier on August 04, 2015, as shown in Figure 3.15c. The AS path thus dropped to a length of 3 till the end of the measurements: it contained an *unknown* AS corresponding to an IP address which belongs to the IXP LAN (196.223.44.0/24). Meanwhile, RTTs between ASes (resp. probe IP addresses) dropped to a set of values with a median of 0.9 ms (resp. 2.6 ms) and a mean of 3.9 ms (resp. 6.3 ms).

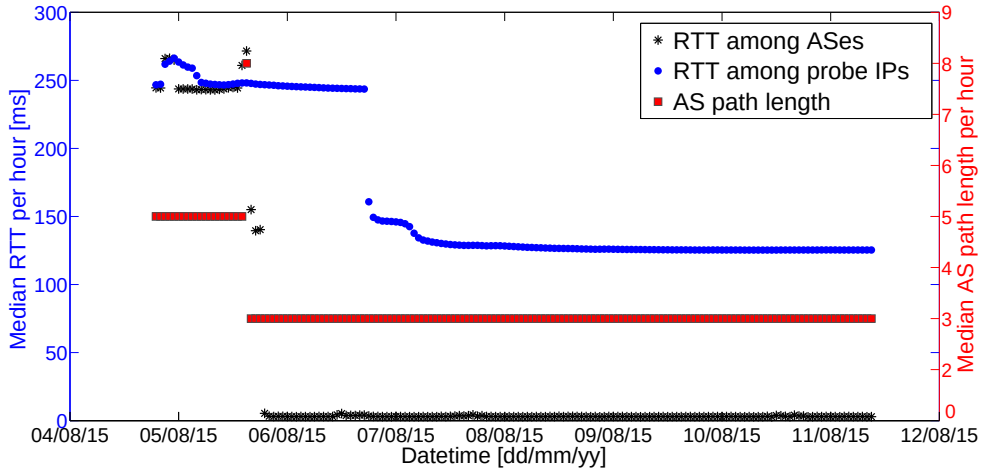

(a) *from AS37557 (NOVAFONE) to AS37410 (LONESTAR)*

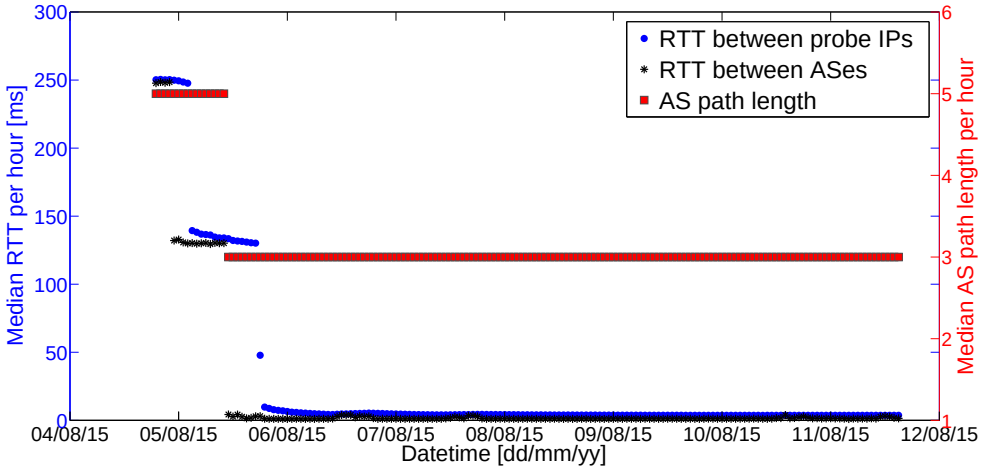

(b) *from AS37557 (NOVAFONE) to AS37094 (CELLCOM)*

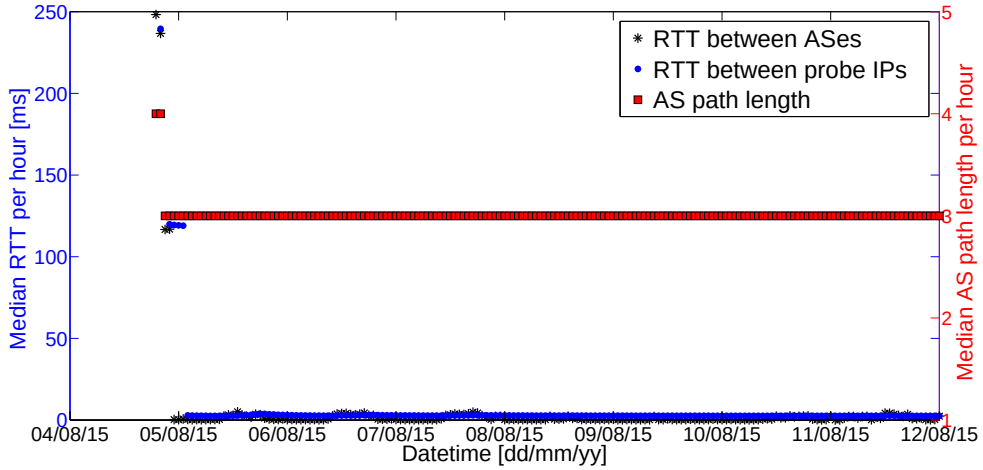

(c) *from AS37557 (NOVAFONE) to AS37203 (LIBTELCO)*

Figure 3.15: RTTs between ASes/probe IPs in NOVAFONE (Liberia, LR) and other LIBERIA-IX (LR) members during the IXP establishment in August 2015.

According to our measurement outputs, the peering session between NOVAFONE and CELLCOM was then established the day after, August 05, 2015. The AS path had a length of 3 until the end of the measurements due to the same reason as above. The BGP peering induced the drop of RTTs between ASes (respectively probe IP addresses) to values with a median of 1.2 ms (resp. 3.6 ms) and an average of 2.1 ms (4.9 ms).

The median length of AS paths between NOVAFONE and LONESTAR also dropped to 3 on August 05, 2015. On the one hand, RTTs between both ASes first decreased to values with a median of 140.9 ms, before declining to 2.9 ms. On the other hand, RTTs between probe IPs stayed at a median of 245.9 ms and then decreased to values with a median of 125.9 ms. The latter RTTs are so high because of either the mediums within the LONESTAR network or its intradomain routing, highlighting the need for the operator to work on reducing them.

In a nutshell, a given NOVAFONE customer communicates with a better QoS with a LIBTELCO, LONESTAR, or CELLCOM customer thanks to the setup of LIBERIA-IX. The IXP also appears as a platform where content or shared resources can be hosted for the benefit of end-users. Meanwhile, all local ISPs save on their transit costs previously paid for local traffic.

MGIX, the Madagascar Internet eXchange We summarize in this section the key findings from the paris-traceroutes measurements carried out between probes hosted by ASes operating in Madagascar (MG).

To assess peering among local networks, we carried out full-mesh paris-traceroutes measurements every 200 s among all active probes located in Madagascar from April to August 2016, as shown in Table 3.2. Local ASes hosting RIPE Atlas probes during this campaign were Orange Madagascar, TELMA, iRENALA, and GULFSAT-AS. Although AIRTELMADA host no probe, we randomly selected online IPs from its allocated prefixes towards which we also launched paris-traceroutes measurements from all the retained probes. As already mentioned, the traceroutes outputs are made publicly available in [84].

The measurement outputs confirm that ASes, which actually peer at the MGIX experience the smallest RTTs among their networks. In fact, AS paths between two pairs of ASes are found to traverse MGIX: Orange Madagascar – AIRTELMADA and GULFSAT-AS – AIRTELMADA. The set of AS path lengths corresponding to the AS pair Orange Madagascar – AIRTELMADA has a median of 3 (passing via an *unknown* AS) over the measurements period, while the set of RTTs between ASes has a median of 0.9 ms with an IQR of 0.03 ms (Figure 3.16). For the AS pair GULFSAT-AS – AIRTELMADA, the set of AS path lengths has a median of 3, while RTTs between ASes have a median of 9.5 ms with an IQR of 8.6 ms. Note, MGIX looking glass [108] lists all three ASes among the members of the IXP, which confirms our findings.

To direct links between local ASes (private peering or not) also correspond low (and relatively better) RTTs between networks on both sides of the interconnect. For instance, the link GULFSAT-AS – Orange Madagascar, which is never found to traverse MGIX, has a median of RTT values between ASes of 5.7 ms and an IQR of 18.4 ms.

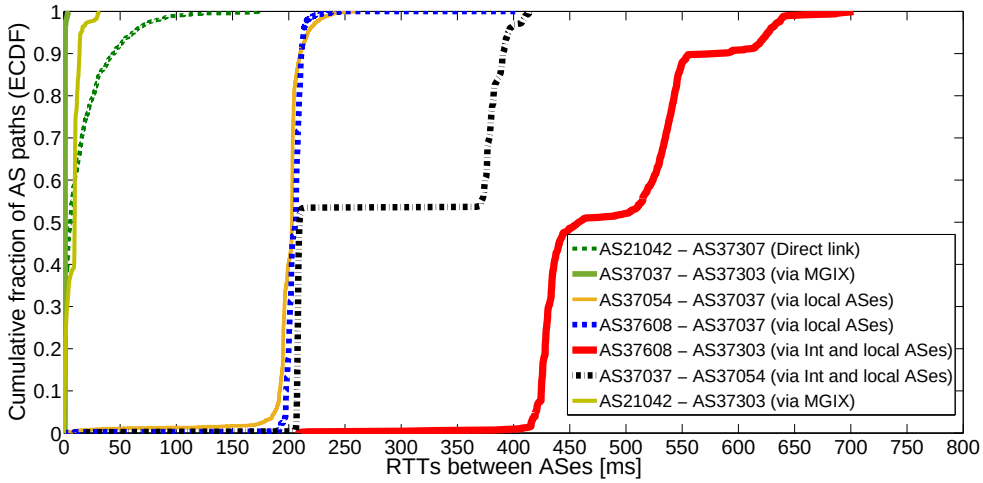

Figure 3.16: RTTs between ASes, which operate in Madagascar (MG), *i.e.*, AS37037 (Orange Madagascar), AS37054 (TELMA), AS37608 (iRENALA), AS21042 (GULFSAT-AS), and AS37303 (AIRTELMADA), showing the effects of being or not a member of MGIX (MG).

AS pairs for which the corresponding AS paths do not traverse the IXP, but traverse instead at least one local AS, have a slightly higher RTT. As an example, the AS path lengths of the pair iRENALA – Orange Madagascar have a median of 3 (often traversing TELMA): the median computed from the recorded RTT values is 206.1 ms, and the IQR is 8.2 ms; unsurprisingly, iRENALA is not listed as a member of MGIX [108]. Another AS pair classified in this category is TELMA – Orange Madagascar, for which AS paths do not traverse the IXP, although those ASes are members of the IXP. In this case, we registered a median AS path length of 4 and a median RTT value of 202.8 ms with an IQR of 8.8 ms.

Finally, AS pairs whose AS paths do not traverse the IXP and transit at least one *Int* AS experience the highest RTTs among their networks. As an example, the lengths of AS paths from Orange Madagascar to TELMA have a median of 6 (via France Telecom-Orange, Cogent, etc.), while RTTs between ASes have a median of 209.5 ms with an IQR of 172 ms. Another example is that of the AS pair iRENALA – AIRTELMADA, for which the median AS path length is 5 (via TELMA, BBIL-AP BHARTI Airtel) and the median RTT between ASes is 458.4 ms with an IQR of 109.8 ms.

3.2.1.6. A look into other IXPs in the dataset

We now examine how frequently an AS path originated from and destined to African countries, EU countries, or the US, traverses an IXP located on each continent. To achieve this, we only consider per pair of ASes the latest discovered AS paths of our 7 measurement campaigns run from 2013 to 2016 (Table 3.2), which contain no *unknown* ASes. We classify the measurement outputs per category, listed in Table 3.8, depending on the region of operation of the AS source and that of the AS destination. We then compute for each category the percentage of AS

paths that do or do not traverse a local IXP, as well as the percentage of paths going through an IXP located on each continent. Table 3.8 presents the results.

Table 3.8: Percentage of AS paths passing via an IXP or not in each continent per category of measurements.

	#AS			%AS paths via an IXP in		
	paths	via IXP	no IXP	Africa	EU	NAm
Among African countries	27,056	32.4 %	67.6 %	16.6 %	16.6 %	0.1 %
Within *SAf* countries	2,663	55.7 %	44.3 %	52 %	3.7 %	0 %
Within EU countries	37,192	67 %	32.9 %	0 %	67 %	0.01 %
Within the US	29,473	22.6 %	77.4 %	0 %	0.006 %	22.6 %

Table 3.8 shows that *no AS path* used for communications among the randomly selected probes in the same EU countries, or within the US, traverses an IXP located in Africa. It also highlights the extent to which communications between devices located in EU countries traverse an IXP in North America (0.01 %) or vice versa (0.006 %). These exceptions are indeed paths traversing Equinix (Dallas, San Jose, New York, Ashburn), in the first case, or Equinix Paris and AMS-IX, in the second. Such patterns are quite similar to that exhibited by communications within *SAf* countries (52 % via an IXP in Africa vs. 3.6 % via an IXP in EU): it is worth noting that for this sub-region and contrary to the US and EU countries, all (136) available RIPE Atlas probes are involved in the measurements as both sources and destinations.

When all African countries are considered, 16 % of AS paths are, however, found to traverse IXPs in Europe: these can be broken down into 67.7 % of paths traversing LINX (Juniper/Extreme), 16.7 % going through AMS-IX, and 12.7 % via DE-CIX. Meanwhile, only 16 % of the AS paths pass through IXPs in Africa. We identify the top three African IXPs as JINX (27.1 % of those paths), NAPAfrica Johannesburg (21.7 %), and CINX (11.4 %). We then evaluate the mean of the set of RTTs between the ingress points of any two African ISPs peering at an IXP located in Africa to 27.4 ms, while it is 70.4 ms for any two African ISPs peering at an IXP located in Europe. In addition to our previous results, these differences prove that *it is often better in terms of QoS for an ISP operating in Africa to peer at its closest IXP in Africa than at an IXP located on another continent*.

3.2.1.7. Evaluating inter-ISP communications performance within the US, EU countries, and African countries

To measure how geographic distances between our probes impact communications performance, we introduce the concept of *normalized RTT*, which refers to the ratio of the minimum measured RTT to the best possible RTT. We compute this metric based on traceroutes outputs collected within the US, EU countries, and countries in African sub-regions before comparing the results.

For each AS path (illustrated by Figure 3.13), we first compute, per corresponding probe pair, the *RTT between probe IPs* (defined in Section 3.2.1.3). We then identify the minimum value

$Min_{RTT}(s,d)$ and the corresponding probe pair (s,d). Next, we compute, using great-circle distances [127], the geographic distance $C_h(s,d)$ between the two probes composing the probe pair identified above. After that, we estimate the best possible RTT between each such probe pair as the RTT that would have been recorded if the two considered probes s and d were directly communicating via an optical fiber of length $C_h(s,d)$. We refer to this value as the *theoretical RTT per probe pair*, denoted $Th_{RTT}(s,d)$. Bearing in mind that light travels about 1/3 slower through optical fiber cables than it does through a vacuum [226, 235], we estimate $Th_{RTT}(s,d)$, as shown in Equation 3.1.

$$Th_{RTT}(s,d) = \frac{2 * C_h(s,d)}{2/3\mathbf{c}} = \frac{3 * C_h(s,d)}{c} \tag{3.1}$$

with $C_h(s,d)$ the great-circle distance (km) between probes s and d, and **c** the speed of light in vacuum (km/ms).

Figure 3.17a presents the distribution of the minimum RTT values $Min_{RTT}(s,d)$ among pairs of probes (in different ASes) operating within countries on each continent. The median of the empirical RTT values collected within EU countries is 13.5 ms, while in the US it is 61.3 ms. Interestingly, in *EAf*, *SAf*, and *WAf* countries, we record a median (and average) of 14 ms (102.9 ms), 22.6 ms (71 ms), and 227.8 ms (655.2 ms) respectively. Comparing Figure 3.17a with Figure 3.17b helps point out that these median values correspond to respectively 5, 4, and 1,891 times the median of theoretical RTTs $Th_{RTT}(s,d)$. In contrast, the medians of the minimum RTTs measured between probes within EU countries and the US correspond to 7 and 3 times those of the theoretical RTTs, respectively.

Next, we compute the ratio $R = \frac{Min_{RTT}(s,d)}{Th_{RTT}(s,d)}$ for the set of measurements targeting any pair of probes located in the US, EU countries, and African countries. We obtain a median of 6.9 for the ratio of values corresponding to AS pairs based within EU countries and 3.1 in the US, vs. 15.7 in *EAf*, 5.1 in *SAf* countries and, unsurprisingly, 1,940.5 in *WAf* countries. Combining the above, it goes without saying that *WAf operators need to deploy more terrestrial fiber within/across countries*: precise suggestions guiding fiber deployments in the whole region are made in Chapter 5. We also encourage them *to implement traffic engineering techniques* and routing policies, which aim at shifting the percentage of AS paths having a ratio below 10 to at least 70 % (*i.e.*, the case of EU countries) and at most 95 % (*i.e.*, the case of the US).

3.2.1.8. Discussions

The purpose of this section is to understand the global African interdomain routing topology without bias towards any sub-region or country and reveal hidden topological changes that have occurred over the last four years. To achieve this goal, we enhanced the RIPE Atlas infrastructure by around 278% by adding new probes. Overall, we collected traceroutes data at random periods from 2013 to 2016, using all (or subsets of) the 324 probes hosted in 169 ASes operating in 40 African countries, the randomly selected 626 probes hosted in 380 ASes in 8 European countries

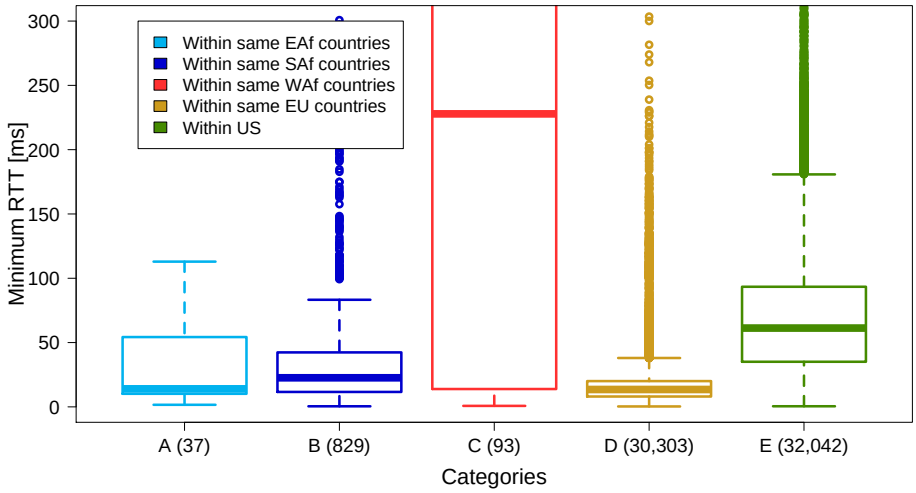

(a) *Minimum RTT measured per probe pair*

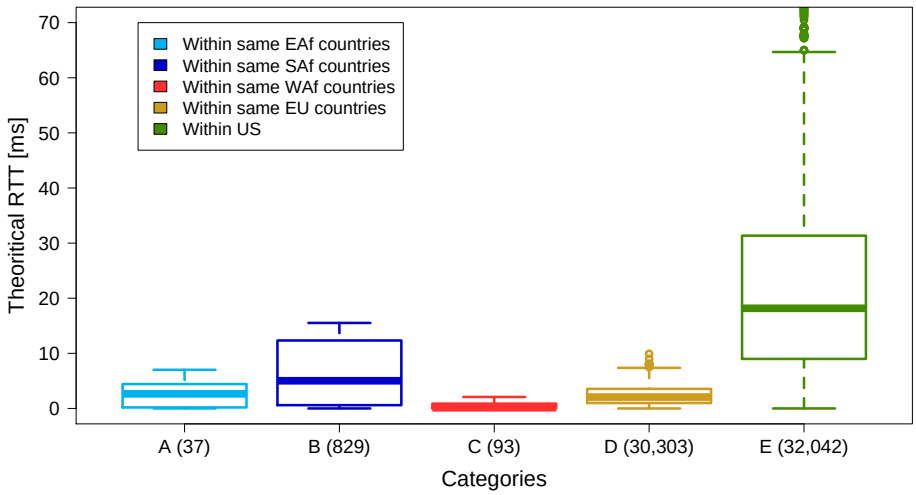

(b) *Theoritical RTT per probe pair*

Figure 3.17: Distributions of the minimum RTT measured ($Min_{RTT}(s,d)$) and theoretical RTT per probe pair ($Th_{RTT}(s,d)$) in same *SAf*, *EAf*, *WAf*, EU countries, and in the US.

and the randomly selected 329 probes in 195 ASes operating in the US. We then adopted, as a best effort, a comprehensive method based on 10 data sources combined with ping measurements to geolocate the IP addresses of routers with high accuracy. While the IP to AS mapping with *TC* allowed us to obtain the corresponding AS paths, we deduced the corresponding country paths with the geolocated IP addresses.

That said, we have first highlighted the caveats in our dataset and their implications on this work before showing that most of the AS adjacencies and AS paths discovered in Africa are not visible in public route-collectors dataset. Our in-depth analysis has revealed a diversity of transit operators playing a role in the provision of both IPv4 and IPv6 African interdomain paths. It has also highlighted the dominant reliance on intercontinental ISPs for the establishment of

continental connectivity. This leads to long AS paths and RTTs, sometimes among ISPs in the same country. We have shown a prevailing lack of interconnection among African ISPs in IPv4 (ZA being an exception) confirming the interest of initiatives to promote peering on the continent. We have also noticed striking differences in ISPs transit habits, notably depending on the official language of the country and the monetary region, specifically in West Africa.

We used RIPE Atlas to inspect the African interdomain topology, as it is open and trustful, it fits our needs regarding measurements and is widely adopted by both operators and researchers. In contrast, Formoso *et al.* [94] have recently adopted Speedchecker [269], a commercial measurement network in their 2017 experiments to achieve the same purposes. We have discussed in detail their reasons for this choice in Section 2.1.1. In fact, by enhancing the RIPE Atlas network and publishing both our findings and raw datasets [84], we have made African network operators realize the importance of network measurements. We have also made both researchers and networks engineers eager to use these VPs for monitoring African networks; we believe these would have been hardly achieved with a commercial network measurement infrastructure.

It is worth underlining that our study is based on longitudinal analysis of datasets collected from full-mesh measurements among 169 ASes operating in 40 (74 %) African countries from 2013 to 2016. Using this data, we could shed light on the dynamics of the African interdomain routing notably by mapping the launch of diverse local IXPs, highlighting their impacts, and drawing attention on performance improvements or decline among local AS pairs, etc. We are pleased to have noticed that a 2017 study (that of Formoso *et al.* [94]) covering 319 ASes in 52 countries across the continent has validated some of our findings. Our remaining results are yet unique insights, however.

As far as our methods are concerned, we adopted a rigorous IP geolocation methodology, which combines the use latency measurements with that of 10 public DSes, contrary to Formoso *et al.* that have used only one of them: *Maxmind (MM)*. As shown by [105, 225] using individual public DSes may lead to wrong inferences in IP geolocation or introduce discrepancies in the geolocation results. To map IP addresses to ASes, Formoso *et al.* [94] have used RIPE RIS, which according to [300] may induce several unresolved mapping. To avoid using single or incomplete data sources that can undermine integrity of research and analysis results [125], we chose instead the *TC* mapping service [286], which is based on data obtained directly from all RIRs [11, 20, 23, 160, 252].

From their 2014 experiments, Gupta *et al.* [117] revealed that, by and large, African ISPs are often either *(i)* not present at the local Internet exchanges or *(ii)* do not peer with one another at those IXPs. As we show in this section, the trends are different when considering a better view of the African Internet from an African point of view involving ASes operating in all its sub-regions. We have indeed detected in our traceroutes outputs, 23 of the 37 existing African IXPs and identified local networks as their members. The number of local networks (or not) peering at each discovered African IXP is also specified in our results. Five of the mapped Internet exchanges (SEYCHELLES-IX, BENIN-IX, SIXP, LIBERIA-IX, and MGIX) are recently established. In

case studies dedicated to each of them, we have highlighted the way they reduce RTTs among peers for a better QoS experienced by customers and produce a drop in the AS path length, which leads to savings on costs paid for transiting local traffic. All this serves to illustrate how critical it is to have quantity and diversity in the VPs used in our measurement campaigns to better assess interdomain routing on the continent. Thanks to them, we shed light on the success of projects aimed at fostering IXP establishments in the African region. Most of all, we can encourage local operators to continue making sustained efforts in this direction for improving QoS as experienced by local users, while reducing their transit costs.

Jensen proved in [152] that at the end of 1996 only 11 countries had Internet access. He added that, by September 2000, all 54 African nations except Liberia (LR) achieved permanent connectivity. Indeed, Liberia was connected in 1999 but lost its link when the local ISP failed to meet commercial viability. We have shown in this section how Liberia has succeeded in deploying its IXP before other African countries (which are currently targeting the same goal) and has four local ISPs connected to it.

The authors of [58] compared different graphs depicting the way communications from South Africa (ZA) (in September 2005 and August 2009) and BF (in August 2009) to other African countries were provisioned. They reported the absence of direct connections between ZA and the following countries: DR Congo (CD), Malawi (ML), Namibia (NA), Tanzania (TZ), Rwanda (RW), Uganda (UG), Mozambique (MZ), Kenya (KE), and Zambia (ZM). Meanwhile, Botswana, Swaziland, ML, MZ, NA were found to have direct routes to ZA (*i.e.*, not via Europe). It also appeared that the more northern countries in particular (CD, KE, TZ, RW, UG, and ZM) were not directly connected. Our results reveal that direct connections have been adopted for communications from ZA to those countries, except CD. However, the situation remains the same for BF, which is still connected to most African countries via Europe with satellites links.

As far as public peering is concerned, JINX was, for instance, listed by Winther [307] in 2006 among the largest Network Access Points in the world. The number of IXPs in Africa then rose, from 8 in 2008 [193] to 18 in 2014 [215, 216, 220]. Actually, in 2013 about a third of African countries hosted an IXP [100], while a half (29) hosted at least one in 2016. The rate of increase was at its highest from July 2014 to July 2015, during which time the number of African IXPs doubled from 18 to 36. As of this writing (September 2017), 38 IXPs are active and functional in the region [218, 292].

Despite this positive evolution, operators and stakeholders still have to devote much more effort. First, the number of IXPs in Africa increased from 5% of the 435 IXPs in the world in February 2014 to 7.5% of the 491 IXPs globally established by July 2016 [215]. Thus, there are still few African IXPs. In comparison, 46 IXPs are operating in Latin America and the Caribbean (LAC), 93 in North America (the US and Canada) and over 130 in Europe [100]. Second, the African interdomain routing is still characterized by the dominance of ISPs based outside Africa for the provision of intra-African communications. Third, Augustin *et al.* concluded in [25], while mapping the IXP substrate, that most IXPs in Africa are small and isolated. Based on information

available on IXPs websites and public databases [218, 220], we found that the average number of African IXP members is 16. As of September 2017, the largest IXPs in Africa are in ZA (with a maximum of about 160 members). Besides, the average number of members at African IXPs is, for instance, lower than the 1/6 of the members of PTT Metro Sao Paulo, where over 300 ASes exchange traffic. It is also insignificant when compared to those of large IXPs in Europe, which are peering over 500 ASes [100]. These confirm that African IXPs are relatively small, compared to those on other continents.

Therefore, local operators need to intensify peering and intra-African traffic localization, while increasing fiber deployment within nations and across sub-regions. IXP members should also update their information on public datasets (PeeringDB, PCH, IXPs websites). Moreover, IXPs should make sure their peering Local Area Networks (LANs) are not routable on the Internet. Finally, all stakeholders need a public tool, which triggers suitable routing policy changes and monitors the African interdomain routing based on real-time measurement data. We release at [256] this tool built based on the dataset collected and the scripts implemented for obtaining the results presented in Section 3.2.1.4. Analyzing peering evolution in the region using publicly available BGP feeds collected since 2005 is also needed for supporting the growth of local IXPs: the design of this system [87] as well as its functionalities and the results obtained from its implementation, are detailed in Section 3.3.1. But before moving to the analysis of those routing data resulting from passive measurements, we present in Section 3.2.2 our investigation of the prevalence, the nature of interdomain congestion at local IXPs as well as our inspection of its causes and impacts.

3.2.2. Investigating the causes of congestion in the African IXP substrate

The growing popularity of bandwidth-hungry applications such as streaming video has generated renewed interest in understanding the nature, location, and causes of performance degradations in the Internet infrastructure. In the US or Europe, some studies have found that interdomain congestion often occurs between networks boundaries, due to peering disputes [103, 167]. However, much less is known about such congestion and its causes at IXPs, particularly those located in developing regions such as Africa. While the possibility of performance problems due to congestion is not unique to IXPs and could also occur within ISPs operating in the region, IXPs are of particular interest due to their position as hubs, which facilitate traffic exchange between hundreds of connected networks. Since there is a significant push to promote peering at IXPs in the African region [6, 81, 85, 156, 292], it is of interest to quantify the performance at those infrastructures. The absence of congestion may contribute to motivating ISPs which are still reticent to join those IXPs. In cases where there is evidence of poor performance, it is also essential to be aware of the causes (peering disputes or other reasons).

To fill the lack of congestion-related measurements at IXPs in Africa, we selected six IXPs located in three of the five African sub-regions [9, 309] for reasons discussed in Section 3.2.2.1. Notably, the only sub-regions involved are *WAf*, *EAf*, and *SAf*, as we were not able to find hosts to

deploy our probes in the other African sub-regions.

We use techniques allowing continuous, fine-grained, and longitudinal measurements. From the outputs of TSLP [167] measurements ran by Ark probes deployed at those selected IXPs over a year (from February 2016 to April 2017), we infer whether or not each of the discovered AS links were congested. We then evaluate the extent to which this phenomenon influenced RTTs to the near and far ends of those links and the characteristics of the observed patterns. After that, we investigate the causes by interviewing the IXP operators. We also evaluate the impacts on the AS links regarding packet loss.

We detect cases of congestion at four IXPs. We show how RTTs and loss rates to the far end increase drastically during the congestion events, and delve into the root causes of the observed congestion. Although we do not find any evidence of widespread congestion, our findings suggest the need for ISPs to monitor the provisioning of their peering links for avoiding or quickly mitigating the occurrence of congestion. Regulators may also define the maximum permissible level of packet loss in those links with the goal of improving performance at local IXPs, thereby making those infrastructures, attractive hubs for local interconnection.

The rest of the section details how we conducted this part of the research. We describe our measurement infrastructure in Section 3.2.2.1. Next, we present, in Section 3.2.2.2, the data collection process as well as the AS relationship inference and validation. Then, we detail the analysis performed on the dataset in Section 3.2.2.3. In Section 3.2.2.4, we present the most interesting case studies of congestion discovered on AS links probed from our VPs, discuss their causes, and examine their consequences.

3.2.2.1. Measurement infrastructure

As mentioned in Section 2.1.2, the time-sequence latency probes (TSLP) method [167] consists of frequently performing RTT measurements from a VP to the near and far routers of an interdomain link. It uses TTL-limited probes set to expire at the near and far ends of that link. When the queue lengths of the routers increase, measured RTTs also increase. We can thus infer from a pattern showing an increase of only RTTs to the far end of an AS link, that a queue between the two routers on both sides of the link instigated the observed delay.

For this study, we adopted the Archipelago (Ark) measurement infrastructure [40] for reasons that are detailed in Section 3.1.2.3. Further, we only considered Ark monitors that support TSLP measurements and are deployed at six African IXPs [87, 292] located in three sub-regions out of five. These are JINX [140] in ZA (launched in 1996), KIXP [287] in Kenya (2002), TIX [284] in Tanzania (2004), Rwanda Internet eXchange (RINEX) [244] in Rwanda (2004), GIXA [104] in Ghana (launched in 2005), and SIXP [261] in GM (2014).

These are interesting IXPs, as *(i)* they are mature and large Internet markets or *(iii)* they have the potential to become regional IXP hubs in the near future, as they are susceptible to attract more members [87]. The term *regional IXP hub* can be defined as a local IXP at which peer most networks operating in the sub-region of the IXP country host, and which thus help localize

traffic among countries located in that sub-region (see Chapter 5). We deployed the VPs in two different settings: some (VP1 – 3) are plugged into the content network of the IXP. By *content network of the IXP*, we refer to the network, usually connected to the IXP switches, which hosts all resources destined to offer common services to the members, namely cache instances, Internet portals, search engines, NTP servers, routing registry, looking glass. Commonly, the content network is not separated from the peering network. As we will see later (Section 3.2.2.4.2), this is not the case of all IXPs. From VPs deployed on the content network of the IXP, we expect to discover all the networks accessing to content available at the IXP. Others (VP4 – 6) are hosted by ASes that peer at the IXPs. From these VPs we will discover, among others, the peers of the network at the IXP.

3.2.2.2. Data collection

We detail in the subsequent paragraphs how the studied dataset is collected and how AS relationships are inferred and validated.

We automatically infer the boundaries of the host network and discover their respective border links using CAIDA's border mapping tool *bdrmap* [43, 168]. The border mapping process [168] consists of gathering routing and addressing data used for data collection and analysis. The input datasets are prefix-AS mappings constructed from RouteViews [189] and RIPE RIS [253], CAIDA's AS-rank algorithm [42] used to infer AS relationships, the RIRs delegated files [11, 20, 23, 160, 252], a list of IXP prefixes from snapshots provided by PeeringDB [220] and PCH [218], a list of sibling ASes of the AS hosting the VP, etc. The creation of the sibling list is a manual process seeded with CAIDA's AS-to-organization mapping: missing siblings are manually added, and spurious siblings are removed. bdrmap then uses an efficient variant of traceroute to trace the path from each VP to every routed prefix observed in BGP. Alias resolution techniques are then applied to infer routers and point-to-point links used for interdomain interconnection. This data is used to assemble constraints, which guide the execution of heuristics to infer router ownership. The border mapping process aims at obtaining sufficient information about the links observed from the AS of the VP toward every other AS to constrain our subsequent border router inferences [43, 168, 262]. For validating the bdrmap output we first check the inferred links against public datasets [128, 218, 220, 254]. The probe hosts are emailed for cross-checking when our results are in contradiction with those public datasets. Four of the six involved VP hosts replied to our queries. They also gave us more insights into the setup of links for which the neighbors of the VP's AS and their respective AS relationships had been rightly discovered by the border mapping process [168] (Section 3.2.2.4.2). This cooperation allowed us to better analyze the collected data. At this end of this process, on average 96.2 % of neighbors of the VP networks are correctly discovered.

Following that, we periodically probe both sides of each discovered IP link every five minutes using TTL-limited probes set to expire at the near and far ends of the link. Regarding ethical considerations, we ensure that our measurements would not adversely affect the VP network by

using a low probing rate (small packets sent at the rate of 100 packets per second). Moreover, the targets of the probing traffic (both ends of each mapped IP link) do not put the VP hosts at risk. Further, our probes do not collect traffic data or any information, which may be considered sensitive due to privacy reasons. Our measurements lasted a year from February 22, 2016, to March 03, 2017. Note, the *bdrmap output* is frequently updated so that new interdomain links are immediately taken into account. We detail how congestions events are detected from the analysis of these TSLP measurement outputs in Section 3.2.2.3.2. In case we detect repeated occurrences of congestion on a link (Section 3.2.2.3.2), we measure packet loss on those links by probing both ends of the said links at a higher rate, *i.e.*, one packet per second, and compute the loss rate over every batch of 100 probes. These were run from July 19, 2016, to April 01, 2017, roughly five months after latency measurements, since we made sure the targeted links are all suffering from repetitive congestion events before launching them. The outputs reveal the losses experienced by communications going through the measured links.

3.2.2.3. Data analysis

3.2.2.3.1. Evolution of number of discovered links For each VP, we identify the links discovered from that VP that are at the IXP, since some VPs are hosted by an IXP member, while others are in the content network of an IXP (Section 3.2.2.1). To achieve this, we categorize the links having any of their IPs belonging to the (peering or management) prefix of any studied IXP as links established at those IXPs. After that, we validate the *bdrmap output* (Section 3.2.2.2) with the corresponding IXP operator through mails and inspect the evolution of the number of neighbors of the VP's AS over time (Section 3.2.2.4.1). We then geolocate both IPs of each link using the Netacuity Edge database [66] and hints in Reverse DNS outputs [124, 225] as added checks to be more confident that those links were indeed established at the IXPs.

3.2.2.3.2. Analysis of congestion cases We begin by gathering the time series collected in Section 3.2.2.2 per VP and discovered neighbor. We then apply an algorithm to detect *level-shifts* in the measured time series, which indicate that the router queue at the interdomain link was filling up, possibly due to the link being congested. The level-shift algorithm [285] identifies changes in the direction of the rank-based non-parametric statistical cumulative sum (CUSUM) test as evidence of a level-shift. It is tuned to use five-minute latency samples and to detect level-shifts that last at least 30 minutes. The magnitude of a level-shift that results from congestion corresponds to the size of the router buffer. We impose a threshold on the minimum magnitude of the level-shifts that we label as potentially caused by congestion. The objective of this threshold is to eliminate false detections that result from noise in the RTT times series or slow ICMP response generation from the routers. Next, we show that this objective is achieved reasonably well with a threshold of 10 ms. We inspect the sensitivity of selecting 10 ms as opposed to 5 ms, 15 ms, or 20 ms by analyzing the variation in the number of inferred congestion cases.

For each value of the threshold, we obtain the links flagged as potentially congested (Table

3.9) and manually check whether the collected TSLP data for those links are following a persistent diurnal pattern indicating peak-hour congestion. We flag 11.2 % more links as potentially congested when using 5 ms; however, the number of links for which we identify a recurring diurnal pattern is the same as that with a 10 ms threshold. In contrast, we flag 50 % fewer links with recurring diurnal patterns when using a 15 ms or 20 ms threshold. Finally, the IXP operators are contacted to confirm whether 10 ms is a reasonable threshold: we received two responses, all of whom stated that they considered 10 ms a reasonable threshold.

Table 3.9: Sensitivity analysis of the threshold value used for labelling potentially congested links in our datasets.

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

To analyze the flagged links that presented recurring diurnal patterns, we ensure that we detect no level-shift on the near side, which would mean that the observed congestion was not at the targeted link. In this step, we also tag for further analysis links showing unclear patterns, *i.e.*, RTTs to the far end present a diurnal waveform, whereas those to the near end are inconclusive. To make robust inferences about whether any observed congestion was at the targeted links, we use the Record-routes method [155, 167] to check path symmetry, thereby ensuring that an increase in RTTs from a near to a far router is solely due to traffic on that link.

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

3.2.2.4. Results and discussion

In this section, we summarize our measurements per IXP and specified the number of observed links, which experienced congestion during the study (Table 3.10). We then shed light on the evolution of the number of discovered links, AS neighbors and peers of the AS hosting each VP. After that, we perform an in-depth analysis of the most interesting results per VP, characterizing whether the congestion was sustained or transient, the impact on packet loss rate, and the causes of the observed phenomenon.

3.2.2.4.1. Evolution of number of discovered links Table 3.10 summarizes per VP the total number of discovered IP links, inferred IP peering links, as well as AS neighbors, and peers obtained from the border mapping process (Section 3.2.2.2) when considering three snapshots. Its column "Discovered IP links" gathers all router-level links found to connect the VP network to that of any of its neighbors. Inferred IP peering links correspond to the subset of discovered IP links having any side that belongs to the IXP prefix (Section 3.2.2.3). The number of neighbors and peers of the AS host are the highest (1,215 and 197 respectively) for our VP in Liquid Telecom that peers at KIXP. We notice that the number of neighbors and peers decreased from 13 on March 17, 2016, to 7 on November 15, 2016, for AS30997 (GIXA): this drop is due to the commercialization of the content network of the IXP (Section 3.2.2.4.2), causing the disconnection of non-registered members. Meanwhile, AS37228 (RINEX) and AS33791 (TIX) have a roughly constant number of peers over our measurement period.

Table 3.10 also presents the number of inferred congested links. *Congested links* are those for which RTTs to the far end show a recurring diurnal pattern, whereas those to the near end stay constant. A congestion case, which is later mitigated is described as being *transient* in the rest of this thesis; otherwise, we refer to it as *sustained*. While for the first four probes, we find one or more cases of congested links, no case is detected for the last two (VP5 and VP6). In fact, the fraction of observed links having experienced any congestion is at most 7.7 % for VP1, 3.3 % for VP2, 0.6 % for VP3, and 33 % for VP4. In total, 2.2 % of the discovered peering IP links have experienced congestion. Therefore, there is not any evidence of widespread congestion. That said, we analyze in depth, in Section 3.2.2.4.2, striking congestion cases observed from VP1 and VP4, highlighting their causes and consequences.

3.2.2.4.2. Analysis of congestion cases

Cases seen from VP1 deployed at GIXA Only two of the links mapped by VP1 hosted at GIXA [104] experienced congestion: those to GHANATEL (ex-Vodafone, GH) and KNET (GH).

GIXA – GHANATEL

The waveform registered for the first link presents different amplitudes over a total of roughly five months. First, RTTs to the far end sometimes peak at 20 ms and 50 ms at other times, while those to the near end remain low and constant during the first 3.5 months (March 03, 2016 to June 14, 2016) termed *phase 1*. Figure 3.18 illustrates part of phase 1. Our analysis of the Record-routes (RR) probes during that period gives us some confidence that the route is symmetric. Since the RR probes show symmetry, then the peak on top of the peak depicted by the shape of the red curve of Figure 3.18 is interesting: it likely indicates congestion in both directions on the link.

From the level-shifts that we find to last from March 15, 2016, to June 14, 2016, and whose existence confirms the occurrence of congestion, we infer as characteristics the average magni-

Table 3.10: Evolution of the number of discovered IP links, AS neighbors, and peers per vantage point.

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

tude A_w of the shifts to be 27.9 ms and Δt_{UD}, roughly 20 hours, implying long congestion events. While discussing with the IXP operator about the possible causes of such phenomenon, we were explained the following: GIXA peering and content networks are separated. The content network (hosting VP1) contains GGCs that need to be updated through transit links. In *phase 1*, GHANATEL was the ISP providing those required transit services through a 100 Mbps link, whereas its clients were served through its main peering link of 1 Gbps size. The 100 Mbps transit link is the one identified by our measurements as suffering from congestion. Thus, GHANATEL users [21] were likely not directly impacted during *phase 1*.

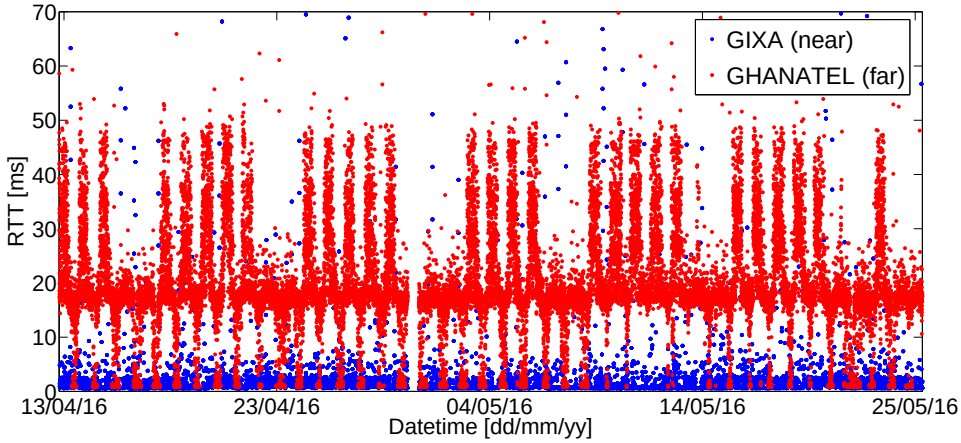

Figure 3.18: RTTs AS30997 (GIXA) – AS29614 (GHANATEL) in part of *phase 1*

Further, we remark that the amplitude of the waveform then dropped to 10 ms from June 15, 2016, to August 06, 2016 (the date from which our latency probes to the far end were unsuccessful): we term this *phase 2* (Figure 3.19a). The beginning of this period coincides with the shutdown of the transit service. The IXP operator explained that GHANATEL shut off the transit service to force the IXP to pay for it. The ISP then used that link for peering till early October 2016, leaving the GGCs non-functional. We still observe a diurnal pattern confirmed by the loss rate increase during that *phase* (Figure 3.19b). Though Figure 3.19b depicts losses reaching only 25 %, our measurements run from July 21, 2016, to August 06, 2017 reveal that losses kept varying between 0 % and 85 % of the packets traversing the link. We conjecture that during *phase 2*, GHANATEL end-users may have been affected by the congested peering link; in addition, all end-users of GIXA peers may have also been affected by the detour of their packets while accessing Google content, which was no longer cached at the IXP.

In early October 2016, GHANATEL stopped using the problematic link. This corresponds to the change of the transit provider by the IXP to an intercontinental ISP, which set up a higher capacity link of 620 Mbps: the IXP is paying for the transit services and members of the IXP are required to register for accessing content. This policy change led to the decrease in the number of peers connected to the content network mentioned in Section 3.2.2.4.1 and Table 3.10.

Finally, we notice that in both phases, the elevation in far end RTTs correlates with days of

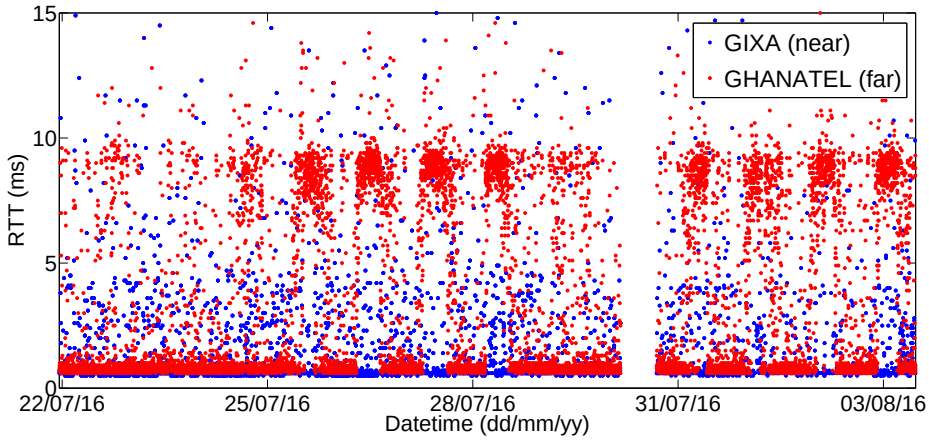

(a) *RTTs to both sides of the link in phase 2 hinting congestion*

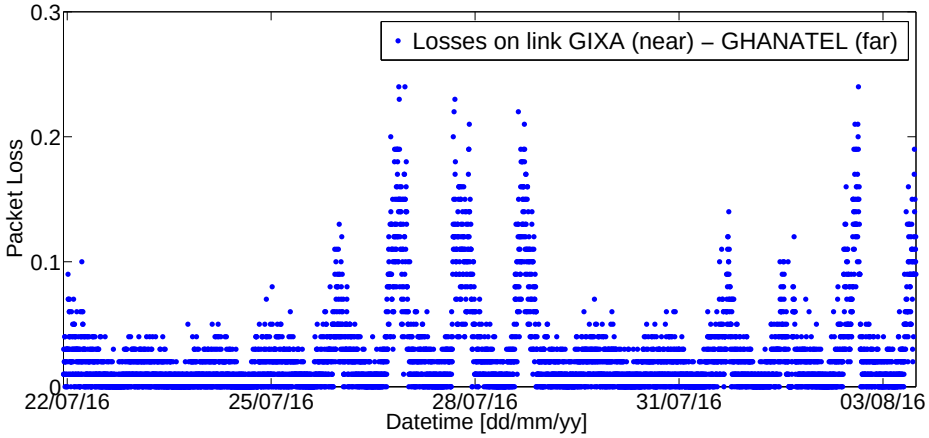

(b) *Packet loss on the studied AS link in phase 2*

Figure 3.19: RTTs and losses AS30997 (GIXA) – AS29614 (GHANATEL)

the week. For *phase 1*, five high spikes correspond to the business days, whereas the rest, to those of the weekend (Figure 3.18). Since congestion events occurred till the shutdown of the link, we deduce the congestion is sustained.

GIXA – KNET

Let us consider the link GIXA – KNET, for which Figure 3.20 presents RTTs to both ends of the link, along with the loss rates. To begin with, KNET delivers high-quality video, data, and voice solutions throughout West and Central Africa [154]. Its link with GIXA was mapped by the VP hosted in that network on June 29, 2016. From August 06, 2016, values of RTTs to the far end present a diurnal waveform, while those to the near end remain constant and stay below 1 ms (Figure 3.20a). Until the end of our measurements, we consistently observed the same pattern for a total of approximately 8 months. Besides, the analysis of RR probes during that period provides evidence of route symmetry in the measurement duration. Further, we evaluate the characteristics of the waveform to find that A_w is 17.5 ms, while Δt_{UD} is of 2 hours 14 min after level-shifts sanitization, *i.e.*, a single congestion event lasts roughly 2 hours.

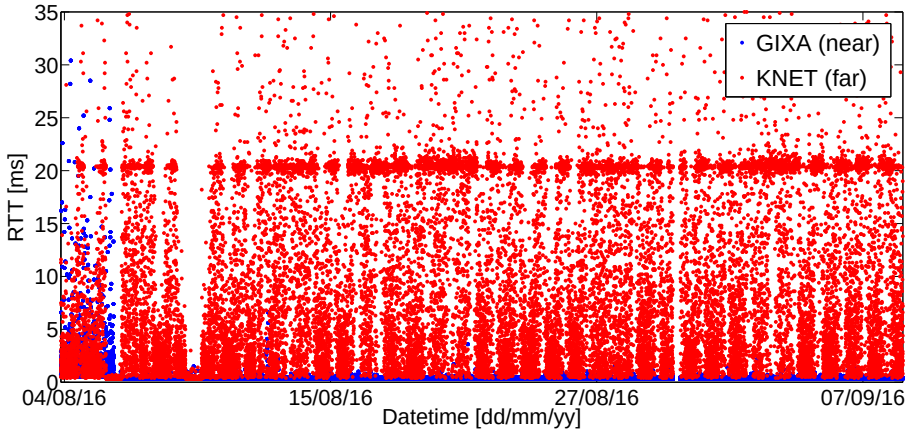

(a) *RTTs to both sides of the studied link hinting congestion.*

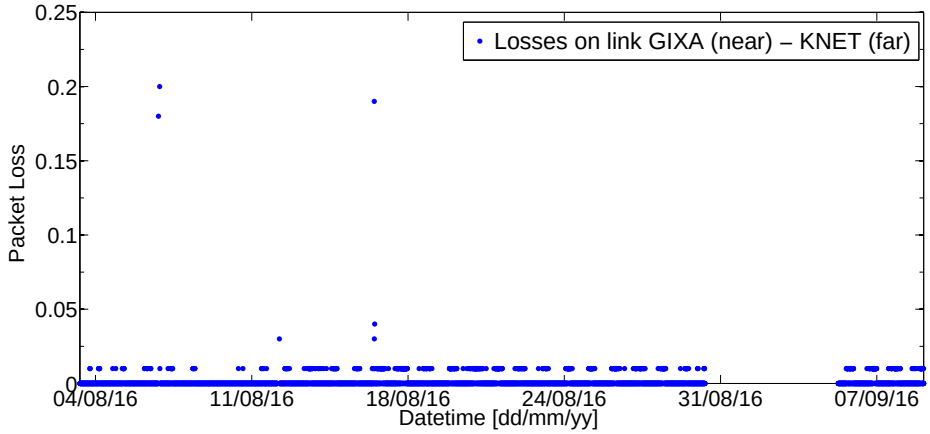

(b) *Packet loss on the AS link in the same period as above.*

Figure 3.20: RTTs and losses AS30997 (GIXA) – AS33786 (KNET).

One might assume, since we started seeing evidence of congestion on the GIXA – KNET link on the same day (August 06, 2016) as the link GIXA – GHANATEL disappeared, that there is a causal relationship between the two events. Further investigation showed that this was not the case: although KNET has a regional footprint, it does not provide transit. On October 06, 2016, during the GIXA operator interview, we were informed that the KNET port at the IXP was not congested. In such a context, two other reasons needed to be investigated internally by KNET: *(i)* whether the router on the far end is overloaded at peak times, resulting in slow ICMP responses or *(ii)* whether the link with the GIXA content network is congested.

On May 05, 2017, KNET informed the IXP that they are not experiencing congestion as opposed to our results and expressed that they have not received any complaints from their customers accessing content. The lack of complaints may be explained by the fact that the average loss rate measured on the link from July 21, 2016, to March 29, 2017 (Figure 3.20b) is low (0.1 %). The observed pattern is the same regardless of the type of the day (business or not). It shows an obvious decrease every day around midnight, an increase at various times of the day, and a constant RTT value around 20 ms in the afternoon. As this pattern is observed till the end of the campaign,

we believe the congestion is sustained.

3.2.2.5. Case seen from VP4 in QCell at SIXP: QCell – NetPage

VP4 is hosted within QCell, a SIXP member. From our previous measurement results presented in Section 3.2.1.5.2, we have found that in August 2014 (a month after the launch of SIXP) that RTTs between QCell and NetPage were constant around 1.5 ms (Section 3.2.1.5.2). However, we notice that the RTTs across that peering link showed repeating diurnal patterns from February 29, 2016, to April 28, 2016 (*phase 1*, shown in Figure 3.21a) indicating congestion on the link. The diurnal waveform then disappeared from April 28, 2016, to March 30, 2017, and most RTT values are below 10 ms (*phase 2*).

While interviewing the SIXP operator, we were told that during *phase 1*, the demand to access the GGCs (for which QCell provide transit) from NetPage was huge: NetPage's engineers noticed that a high bandwidth dedicated to the Google traffic from their users was degrading interdomain performance and causing congestion. They thus asked for an upgrade of their link with SIXP from 10 Mbps to 1 Gbps. After the upgrade (done on April 28, 2016, according to our raw data), the congestion events disappeared and are not evident until the end of the measurement campaign (Figure 3.21b). We believe NetPage's end-users may have been affected by these events, however.

As for the characteristics of the waveform, the average magnitude A_w of the level-shift during *phase 1* is 10.7 ms, while the period of the waveform is of roughly 1 day. Moreover, congestion events last on average a third of the duration of those registered during *phase 1* for the link GIXA – GHANATEL (Section 3.2.2.4.2), since Δt_{UD} is 6 hours 22 min. Finally, we notice that the waveform is the same over weeks (Figure 3.21a) and that to each day corresponds a spike. The height of the spike reaches 35 ms in the week, whereas it stays around 15 ms during the weekends. The reasons behind this may be an intensive access to Google content for daily activities combined with a high amount of communications among clients of both ISPs during business days compared to weekends.

3.2.2.6. Implications of our results

We now highlight the takeaways from our work and discuss the implications for research and network operations. A key takeaway is that we have observed congestion on only a small fraction of the monitored links during this measurement period. However, we have also noted that the IXP ecosystem is highly dynamic in Africa, as ISP presence at IXPs, policies adopted by the IXPs, and the presence of CPs can change over time. With the push for peering in the African region, it is likely that the IXP substrate will become more mature in the future, supporting more peering between interconnected networks and hence increased traffic volumes. All these factors motivate the need for longitudinal measurement and monitoring of this evolving infrastructure.

We have shown that the TSLP technique can detect congestion without requiring access to data from network operators. However, we emphasize that judicious interpretation of the *causes*

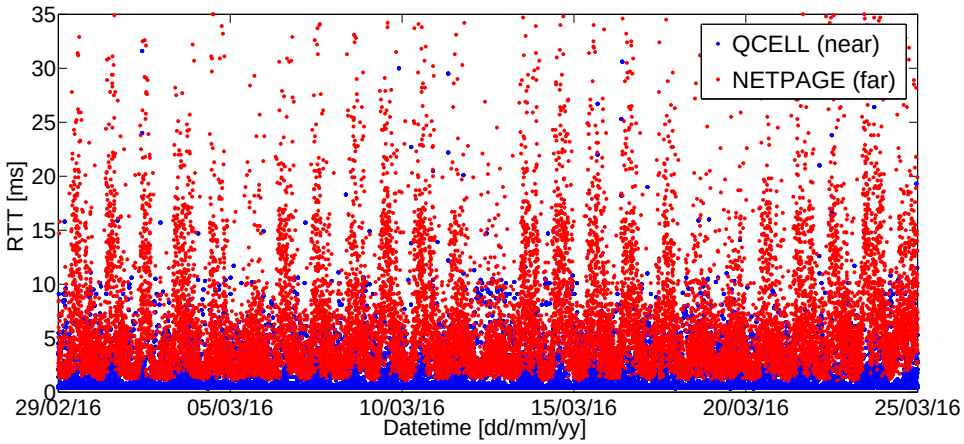

(a) *RTTs to both sides of the link in phase 1, hinting congestion.*

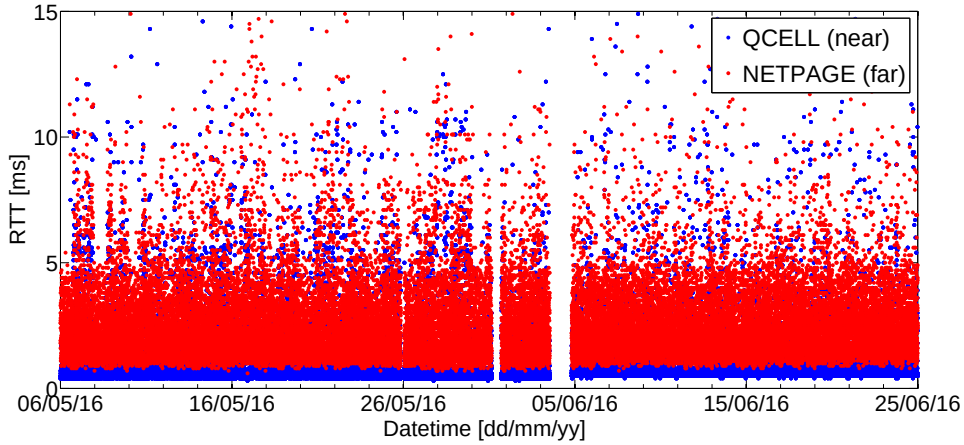

(b) *RTTs to both sides of the link in phase 2: the diurnal pattern disappeared and the congestion events are not evident anymore.*

Figure 3.21: RTTs AS37309 (QCell) – AS37323 (NetPage).

of the observed congestion events requires the collaboration/validation of the stakeholders, as these are often related to hidden events that are not made public. We have learned that congestion occurred on a link used to update Google caches hosted at the IXP, on a link used by an ISP to peer at the IXP (*e.g.*, VP1), or on an under-provisioned link connecting a Google cache host to one of the IXP peers (*e.g.*, VP4). High demand appears to be the main cause in the last scenario: since we had no access to data on traffic traversing the studied IXP, we could not cross-check with our findings to analyze whether streaming video was the principal source of this high demand. Such a cross-correlation of TSLP data with data from the operators is retained as a future work. In the two first cases, congestion was sustained; in the case of the link GIXA – GHANATEL, there was a dispute between the two parties, while in the case of the GIXA – KNET link, the low packet loss on the link likely meant that end-users were not severely impacted and hence the ISP did not upgrade the link. Further, it is worth mentioning that even though our findings regarding the causes of congestion at IXPs may apply to IXPs in other regions, we have preferred not to attempt to generalize them beyond our observations and validations with the operators.

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

3.3. Passive measurements

3.3.1. A Route-collectors Data Analyzer for monitoring the growth of peering in an Internet region: Case study of AfriNIC

Since Internet connectivity appears to be a lever of development in connected areas, there is an increasing interest from the Internet community in continually characterizing local interconnection in under-connected regions to efficiently help improve it. Meanwhile, ISPs are more and more interested in acquiring updated details about the current situation to identify potential positioning opportunities in those geographical areas [61, 79, 81, 117, 170, 202]. In this perspective, this work aims at designing and implementing a system able to profile the set of IXPs in an Internet region.[8] As part of the ISOC strategy to allow the Internet community to monitor and understand the evolution of the IXPs in a particular region, we developed a route-collector data analyzer tool and afterward we deployed and tested it in AfriNIC, which represents the Internet frontier due to its low Internet penetration (*cf.* Section 1.1). Moreover, this study is in line with the need for a longitudinal measurement and supervision of its evolving IXP infrastructure, mentioned in Section 3.2.2.6. In fact, the African peering ecosystem has been the subject of much attention over the last four years with the goal of meeting the traffic localization challenge: the efforts of stakeholders in this direction have led to the setup of more Internet exchanges [6]. There are 38 IXPs in Africa (hosted in 29 countries), of which 20 have been set up since 2009 [292] as of September 2017.

However, 21 years after the launch of the first IXP, the monitoring and measurement infrastructure of the region still challenges the evaluation of the progress made on traffic localization. Assessing the impact of related activities, such as policy implementation and infrastructure developments, is quite challenging, considering that very few IXPs provide publicly accessible data on current traffic statistics or colocation data. As noticed in Table 3.5, PeeringDB and PCH public datasets on IXP colocation are not up-to-date when it comes to IXPs in Africa, because some IXP members do not register in those datasets or do not add their prefixes (Section 3.2.1.5). Besides, locally useful data essential to support the growth of peering in the region is unavailable: this is particularly important in regions such as AfriNIC or LACNIC, where the hidden Internet topology complicates the analysis of the expansion possibilities [29, 81, 85]. Further, other increasingly

[8] It has notably been conducted in collaboration with Vìctor Sánchez-Agüero, Ph.D. Student at IMDEA Networks Institute and Universidad Carlos III de Madrid (UC3M, Spain, ES)

useful measurement resources (*e.g.*, the RIPE Atlas network [2, 34, 248, 250]) still offer limited visibility in the African IXP substrate, since only 17.5 % local networks host a RIPE Atlas probe despite intensive deployment efforts [81, 85].

The ISOC then decided to help offset the lack of progressive, visual, and near real-time information on the status of networks operating in a given Internet region, developing in a joint effort with UC3M a methodology that enforces the collection, collation, and publication of useful data points, from an internal VP. In fact, automating these tasks will easy the monitoring and reporting of the progress being made on interconnection and traffic exchange in the said region. We actively contributed to the definition of that methodology as well as to the design of a system, which automatizes it and whose implementation for the AfriNIC region led to the African Route-collectors Data Analyzer (ARDA). ARDA is an open-source tool with a web interface that constantly collects raw routing data from route-collectors (existing at IXPs in Africa) with a peering viewpoint of the Internet (defined in Section 2.1.3). It then inspects this data from various angles to assess peering evolution in the region. It was built in 18 months from December 2015 and is freely available at `arda.af-ix.net` [87], *i.e.*, hosted in the domain of the African IXP Association (Af-IX, `www.af-ix.net`).

Such a compass is intended to: *(ii)* provide *network operators* with supporting information for peering decisions, *(iii)* provide empirical data to support business investment decisions and opportunities in the region (*Internet business development*). Besides, this tool will *(iv)* inform development organizations and policy-makers on gaps and state of interconnection in the region (*Internet community)*, and *(i)* help *researchers* undertake interconnection studies or to complement measurement studies that use other data sources, such as RIPE Atlas network [248], Ark measurement infrastructure [40], etc. Needless to say, it will contribute to achieving the objectives n^{o} 1, 3, 4, and 7 of this thesis enumerated in Section 1.2.1.

We present our definition of the techniques automatized by the route-collectors data analyzer as well as our design and implementation of the ARDA platform, including the key algorithms used to analyze data, analysis results from the BGP data, and use cases showing their value for the Internet ecosystem. As it will be seen, ARDA is built so that it can easily be applied to other regions in the future. It is still a living project that the ISOC keeps on supporting, and some of its possible features and expansions will also be commented in Section 6.2.

The methodology adopted to achieve our purposes constitutes the remainder of this section. After discussing the related work in Section 2.1.3, we define the requirements of the route-collectors data analyzer in Section 3.3.1.1. We then introduce its architecture in Section 3.3.1.2. Next, we present the different steps of the data collection and storage process (Section 3.3.1.3.1), the data analysis (Section 3.3.1.3.2), highlighting our technical choices for the implementation of the ARDA platform from which arose some striking (visualization) results that underline its relevance for the Internet community (Section 3.3.1.3.3).

3.3.1.1. Requirements of the route-collectors data analyzer

In this section, we better highlight the main visible outputs expected at the end of this work. They can be listed as follows:

1. **IXP growth and business potential**: the route-collectors data analyzer should constantly provide graphical views of the visible networks at each IXP, help IXPs market their features, and help end-users identify sub-regions that are connected to a particular IXP.

2. **Interconnection development progress and gaps**: the route-collectors data analyzer should monitor local and regional interconnection growth, help identify IXPs that are facing potential challenges, as well as track local and regional policy and regulatory impact on interconnection development.

3. **Technical support**: the route-collectors data analyzer is expected to report on the networks that are likely to have routing inefficiencies at each IXP.

We next determine the main aspects around which the designed system can be centered. The route-collectors data analyzer must depend on a reliable system, which locally collects and stores in a common format, the historical and current routing data previously fetched by passive measurements at the IXP. Pre-defined statistics, termed *metrics* in the rest of this section, are then expected to be computed based on this collected data and presented under the following three views: *(i)* the *IXP View* whose metrics are per IXP *(ii)* the *National View* whose metrics involve the set of IXPs in the same country, and *(iii)* the *Regional View* for which the provided metrics cover all IXPs in the region. Further, the designed route-collector data analyzer should have the ability to integrate private route-collectors deployed by local IXPs and the ability to be configured for other regions. Regarding the implementation of the designed system in the AfriNIC region, it is essential to select a suitable location on the Internet, where this system could be hosted so as to be delivered with a high QoS to its potential users: IXP operators, Internet developmental institutions, current and potential peers (network operators, CPs), etc.

3.3.1.2. Proposed architecture of the route-collectors data analyzer

The architecture of the route-collectors data analyzer is composed of three modules (Figure 3.22), which have been defined given the above-listed tasks. First, the *data collection module* is in charge of automatically identifying existing route-collectors, their type, and location in the studied region. It then ensures the concurrent download and parsing of BGP data from those sources to extract entries corresponding to those of our data structure. Not only this module collects historical BGP data in the background, but also it downloads the latest available routing data hourly or daily. Second, the *data storage and metrics computation module* ensures the storage of the key information from among those previously extracted and their usage to compute our metrics using data for the last month, the last year, or the whole period of the dataset. The results

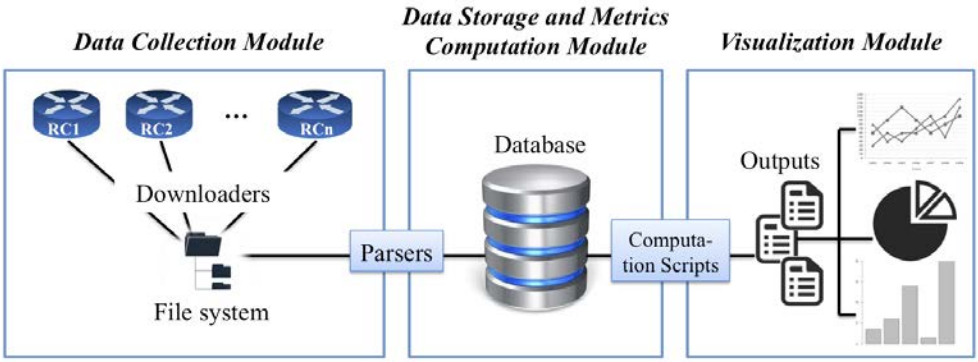

Figure 3.22: Architecture of the route-collectors data analyzer. RC stands for route-collector

are divided into weeks, months, and years respectively. Optimized algorithms, parallelism for fast computations are essential to delivering those results in real-time. This module thus contains *numerous scripts* (playing distinct functionalities) of which any set are concurrently launched by *an orchestrator* to satisfy the need to update the values corresponding to each metric on time. Finally, the *visualizations module* generates in real-time and presents in the most appealing way dynamic graphs depicting the evolution of the previously computed metrics. Those charts are classified depending on the three views mentioned in Section 3.3.1.1. As one can notice, each module logically relies on the results obtained by the previous ones and on their good functioning. The functioning and our implementation of these modules in the ARDA platform are detailed in Section 3.3.1.3.

3.3.1.3. Implementation of ARDA and results

3.3.1.3.1. Data collection The selection of the suitable DSes is critical for successfully meeting the requirements while implementing ARDA's data collection module. For the geolocation of any new route-collector, we retained four DSes (OpenIPMap (*OIM*) [245], Maxmind (*MM*) [187, 188], reverse DNS lookups outputs (*RDNS*), and Team Cymru (*TC*) [286]). These DSes are cross-checked as explained in [81, 85] and in Section 3.2.1.3. When all DSes having an entry do not return the same CC for the detected route-collector, ARDA does not suggest any location and lets its administrator manually add it. Based on the CC, the route-collector can then be tagged as deployed or not in Africa.

To make ARDA give a broad view of the IXP substrate in the AfriNIC region, we chose to design it so that it combines data from existing RouteViews collectors with those from PCH and IXPs private route-collectors to perform its analysis. However, both the peering and transit links of Liquid Telecom and Network Platforms LTD (two JINX members) are captured by the JINX RouteViews, as shown in Figure 3.23a: this route-collector thus does not have a peering viewpoint. Consequently, we removed it from the set of route-collectors. For similar reasons, the RouteViews collector deployed at NAPAfrica (ZA) on September 7, 2017, cannot be considered either (*cf.* Figure 3.23b). Since the Internet eXchange Points JINX and NAPAfrica host PCH

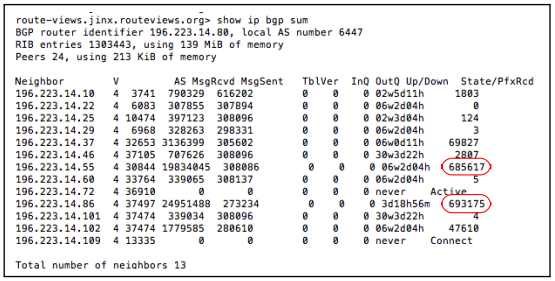

```
route-views.jinx.routeviews.org> show ip bgp sum
BGP router identifier 196.223.14.80, local AS number 6447
RIB entries 1303443, using 139 MiB of memory
Peers 24, using 213 KiB of memory

Neighbor        V         AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
196.223.14.10   4  3741  790329  616202        0    0    0 02w5d11h      1803
196.223.14.22   4  6083  307855  307894        0    0    0 06w2d04h         0
196.223.14.25   4 10474  397123  308096        0    0    0 02w3d04h       124
196.223.14.29   4  6968  328263  298331        0    0    0 06w2d04h         3
196.223.14.37   4 32653 3136399  305602        0    0    0 06w0d11h     69827
196.223.14.46   4 37105  707626  308096        0    0    0 30w3d22h      2807
196.223.14.55   4 30844 19834045  308086        0    0    0 06w2d04h    685617
196.223.14.60   4 33764  339065  308137        0    0    0 06w2d04h         5
196.223.14.72   4 36910       0       0        0    0    0 never    Active
196.223.14.86   4 37497 24951488  273234        0    0    0 3d18h56m    693175
196.223.14.101  4 37474  339034  308096        0    0    0 30w3d22h         4
196.223.14.102  4 37474 1779585  280610        0    0    0 06w2d04h     47610
196.223.14.109  4 13335       0       0        0    0    0 never    Connect

Total number of neighbors 13
```

(a) *"sh ip bgp sum" on the JINX Routeviews collector. AS30844 (Liquid Telecom, UK) and AS37497 (Network Platforms LTD, ZA) advertise their full view of the Internet to the route-collector given the number of IP prefixes received from those ASes.*

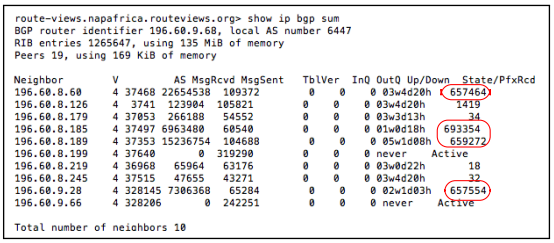

```
route-views.napafrica.routeviews.org> show ip bgp sum
BGP router identifier 196.60.9.68, local AS number 6447
RIB entries 1265647, using 135 MiB of memory
Peers 19, using 169 KiB of memory

Neighbor        V         AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
196.60.8.60     4 37468 2265453B  109372        0    0    0 03w4d20h    657464
196.60.8.126    4  3741  123904   105021        0    0    0 03w4d20h      1419
196.60.8.179    4 37053  266188    54552        0    0    0 03w3d13h        34
196.60.8.185    4 37497 6963480    60540        0    0    0 01w0d18h    693354
196.60.8.189    4 37353 15236754  104688        0    0    0 05w1d08h    659272
196.60.8.199    4 37640       0  319290        0    0    0 never    Active
196.60.8.219    4 36968   65964   63176        0    0    0 03w0d22h        18
196.60.8.245    4 37515   47655   43271        0    0    0 03w4d20h        32
196.60.9.28     4 328145 7306368   65284        0    0    0 02w1d03h    657554
196.60.9.66     4 328206       0  242251        0    0    0 never    Active

Total number of neighbors 10
```

(b) *"sh ip bgp sum" on the NAPAfrica Routeviews collector. AS37468 (Angola Cables, AO), AS37497 (Network Platforms LTD, ZA), AS37353 (Macrolan, ZA) and AS328145 (Lyca Digital, ZA) advertise their full view of the Internet to the route-collector given the number of IP prefixes received from those ASes.*

Figure 3.23: Outputs of *"sh ip bgp sum"* run on JINX and NAPAfrica RouteViews collectors as of October 22, 2017, showing that they capture routing information received via both peering and transit links by some of their peers.

route-collectors with a peering viewpoint and a more significant set of peers, not considering their respective Routeviews collectors has little impact on the quality or the scope of the data collected.

Table 3.11 summarizes the type and number of route-collectors per IXP covered by our dataset and their corresponding country host. In that table, RCs means route-collectors, PCH stands for Packet Clearing House and RV stands for RouteViews. Table 3.11 also specifies the year of the launch of each IXP and the date of the deployment of the first route-collector of each type; using these dates, we compute the gap period needed to better point out the dataset limitations. IXPs private route-collectors have not yet been included in our DSes. In total, ARDA involves data from all (41) route-collectors of the region, which could be taken into account for this work. These are deployed at 24 IXPs in 18 African countries located in four African sub-regions out of five (Section 1.1.2).

Further, we used RIRs datasets [11, 20, 23, 160, 252] to extract information related to ASNs and prefixes assignments. Finally, we selected DSes from APNIC's routing table analysis [263] for any comparison between routing information at the IXPs and those appearing on the Internet.

3.3.1.3.2. Data storage and metrics computation We enumerate below some metrics evaluated by ARDA and detail the algorithms used for their computations, showing how they fit into the three main aspects listed in Section 3.3.1.1.

IXP growth and business potential

To evaluate the growth of each involved IXP, the number of visible prefixes, origin ASNs, and peering ASNs are quantified per week, month, and year.

The number of visible prefixes at an IXP represents the number of distinct prefixes seen at all its route-collectors. While computing it, bogon prefixes are separated from those routable on the Internet to help identify IXPs at which peers announce more bogon prefixes. Similarly, the distinct

Table 3.11: List of the 24 African IXPs and the corresponding 41 route-collectors subject of this study.

CC	Country	IXP	Year of IXP launch	Type (#) RCs	1st date of RC deployment	Gap period (in years)
BJ	Benin	BENIN-IX	2013	PCH (1)	29/07/2015	2
BW	Botswana	BINX	2005	PCH (1)	08/07/2016	11
EG	Egypt	CAIX	2002	PCH (2)	17/10/2011	9
GM	Gambia	SIXP	2014	PCH (1)	20/02/2015	1
KE	Kenya	KIXP	2002	RV (1)	07/10/2005	3
				PCH (3)	06/08/2010	8
		MSA-IX	2014	PCH (1)	10/02/2017	3
LR	Liberia	LIBERIA-IX	2015	PCH (1)	13/01/2016	1
MG	Madagascar	MGIX	2016	PCH (1)	15/03/2016	0
MW	Malawi	MIX	2008	PCH (2)	11/07/2013	5
MU	Mauritius	MIXP	2008	PCH (1)	25/05/2015	7
MZ	Mozambique	MOZIX	2002	PCH (2)	21/07/2010	8
NA	Namibia	WHK-IX	2014	PCH (1)	17/06/2015	1
NG	Nigeria	IXPN	2007	PCH (2)	30/01/2015	8
RW	Rwanda	RINEX	2004	PCH (1)	11/05/2015	11
SD	Sudan	SIxP	2011	PCH (2)	10/12/2014	3
ZA	South Africa	JINX	1996	PCH (3)	19/07/2013	17
		DINX	2012	PCH (2)	21/02/2014	2
		CINX	1997	PCH (2)	21/07/2010	13
		NAPAfricaCT	2012	PCH (3)	18/04/2013	1
		NAPAfricaDB	2012	PCH (1)	22/09/2015	3
TZ	Tanzania	AIXP	2006	PCH (1)	15/06/2015	9
		TIX	2004	PCH (1)	06/06/2015	11
TN	Tunisia	TUNIXP	2011	PCH (3)	09/12/2014	3
UG	Uganda	UIXP	2001	PCH (1)	13/06/2016	15
Total		24 IXPs	From 1996	PCH (39) RV (1)	From 2010 From 2005	0 — 17

origin/peering ASNs visible in the routing data collected at each IXP are listed. While the origin AS (whose identifier is the last ASN from the left) of a given AS path is the network originating the prefix, the peering AS (first ASN from the left) is that connected to the IXP route-server.

The evolution of those numbers highlights how popular is a local IXP compared to others and how fast it has been growing. It also helps identify IXPs with the highest/stable number of peers or reachable networks in the region/each sub-region, as well as those, which are not functional for a while and the corresponding malfunction period.

With routing data covering the last four weeks, the percentage of prefixes (assigned to each country in the world), which are seen at any local IXP, is then computed. Towards this end, the set of prefixes allocated by an RIR to its countries members is fetched from each RIR database. ARDA then verifies if any of the prefixes visible at an IXP overlaps any such allocated prefix. The percentage of prefixes assigned to a given country that are visible at the considered IXP, therefore, represents the ratio of the *number of prefixes seen at the IXP that overlapped those assigned to the country* to the *total number of prefixes assigned to that country*.

Such statistics will give IXP members and prospects an accurate knowledge of which countries or regions they are/will be able to reach while/after peering at any IXP in Africa. They are

intended to help prospects compare those IXPs by their ability to allow them to reach countries of their interests. The results are presented in Section 3.3.1.3.3.

Next, ARDA compares the percentage of IPv4 to that of IPv6 blocks assigned to the country hosting a given IXP, which are seen or not at that IXP. To achieve this, all IPv4 and IPv6 blocks allocated to the country host of the IXP are identified. ARDA then checks if any prefix seen at the considered IXP overlaps any such blocks. The ratio of the *number of visible prefixes at an IXP found to overlap the assigned IPv4/IPv6 blocks* to the *total number of assigned IPv4/IPv6 blocks* is then computed.

Interconnection development progress and gaps

The metrics listed in Section 3.3.1.3.2 are evaluated at the national and the regional levels to monitor interconnection development growth and gaps. While the national level gathers data from all IXPs in a given African country, the regional level presents data from all IXPs located on the continent.

Technical support

ARDA also reports on networks, which are likely to have routing inefficiencies at each IXP. First, the number of prefixes of various length announced over time is quantified. Second, the behavior of IXP members on aggregation and de-aggregation when announcing their prefixes at the peering points is compared to that at their upstream. To inspect this, all assigned prefixes are fetched and individually cross-checked with the set of prefixes visible on the Internet[9] available at [263], thereby identifying the allocated blocks, which match those announced on the Internet. The length of the latter prefixes is then contrasted with the length of those, visible at each IXP. The goal of this comparison is to identify prefixes whose announcement at the public peering fabric are shorter, match exactly (best practice), or are longer. Performing such an analysis aims at raising awareness amongst IXP members, which are not applying the best practice.

3.3.1.3.3. Visualizations and Results Before presenting some showcases of its functionalities, underlining their usefulness, and revealing striking results that demonstrate how ARDA can help profile the African IXP substrate in real-time, we specify technical details related to its implementation.

Technical choices

For local content to be hosted locally and be as close as possible to most potential users, the server destined to host ARDA was planned to be deployed within the infrastructure of an African IXP. The JINX infrastructure (in ZA) was selected, given the stability it has acquired as the oldest IXP (Table 3.11) in Africa and since several networks are connected to it (Table 3.5).

[9] We involved 3 international looking glasses APNIC's router in Washington, US, APNIC's router at DIX-IE, Japan, and Bhutan Telecom's router at LINX, London

The hardware destined to host the web server was then selected for a high availability service (64 GB of RAM, two Intel Xeon 2.4 GHz processors, redundant power supplies, 18 TB of disk space composed of hot-swappable hard drives, etc.). These choices were also made considering the number of concurrent clients (expected to reach thousands of people), the loads of answering their requests, and the computation tasks that the server will have to support. Next, a Linux – Apache – MySQL – PHP (LAMP) server was built. We only included open source technologies so that anybody can interact with the scripts without expenses, once the code is released.

End-users that will interact with ARDA were classified into two categories: the common users and the administrator. The common user can be an IXP member/operator, an ISP engineer, a decision-making institution, a member of the Internet community, or a researcher. The administrator is responsible for ARDA maintenance and management.

To avoid long waits while end-users are accessing the results, pre-computing the numerical values of each metric was preferred to computing them upon requests. The corresponding set of python scripts, therefore, uses the needed raw data to frequently compute the metrics listed in Section 3.3.1.3.2 and deliver up-to-date information to the visualization module. Those outputs are then directed to text files, which are re-used by the PHP and Javascript scripts to display the graphs. Another measure taken to achieve this goal was to physically separate the computation from the visualization module (Figure 3.24). Three virtual machines (VMs) are thus hosted on the server. The first one termed *Pulse*, which is the most powerful, is destined to the computations. The second one termed *Front-end*, less powerful, plays the role of the web server. The last one, even less powerful, hosts a monitoring system that supervises the three VMs and the host machine. All of them host the OS Ubuntu 14.04.3-server.

Every 15 mins, the outputs of metrics computations are transferred from *Pulse* to *Front-end*, under text files formats, some of which can be downloaded upon requests. The adopted technical architecture (Figure 3.24) allows *Front-end*, and thus, ARDA to still be functional with end-users accessing the results of the last computations, even if *Pulse* were to experience a failure. It will also enable caching (*Front-end* at diverse locations) in the near future.

Further, a MySQL database for hosting the raw routing information, a database for hosting the RIRs assignment data, and another one for user-related information (Figure 3.24) were built. The former was indexed for more efficiency in the data storage and their provision to our scripts. The main information composing its data structure are the type of the route-collector, the route-collector name, the AS path, the origin ASN, the network, etc. Details related to route-collectors are stored in the same table. Any IXP at which a new route-collector is later deployed has its information automatically added in that table and is included in the next series of computations.

To avoid overloading *Pulse*, the maximum number of computation scripts running simultaneously was set to 8, given their individual workload. In addition, the historical or real-time data downloaders are always running in the background. An *orchestrator* was then designed to play the role of tasks scheduler *i.e.*, it identifies per view, every four hours, the script whose end of execution date is the oldest and relaunches it when the maximum number of scripts is not exceeded

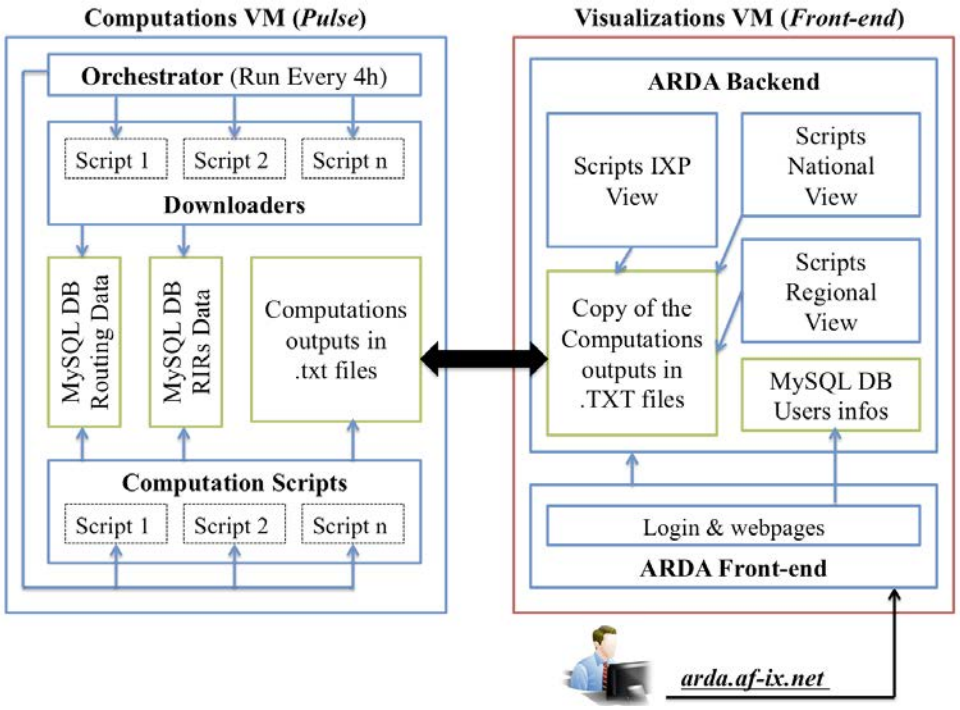

Figure 3.24: Simplified ARDA technical architecture

(Figure 3.24). By doing so, it ensures that every 15 days, most of ARDA's analysis results are updated at least once.

Regarding data collection and storage, IPv4 and IPv6 RouteViews real-time data are hourly fetched using BGPStream [39, 209] since June 2016. This operation, combined with the data parsing and storage, usually ends within the first 15 mins of each hour. Meanwhile, IPv4 and IPv6 snapshots [213] are daily fetched from PCH website: it is unfortunately the only way to get this information as of this writing, since there is no API to access this data. On average 8 min are needed per day for downloading and storing the data from the 40 PCH route-collectors of Table 3.11, while pausing in between any two of them for a random period. Further, we defined the format of the outputs of private route-collector data sources as being the same as those of PCH so that a similar treatment can be applied to both inputs.

It is worth mentioning that IPv4 and IPv6 historical data were downloaded for the period 2005 to end of May 2016. As of April 2017, all PCH and RouteViews historical data were fully downloaded, parsed and stored. The size of the database as of October 25, 2017 is 380.1 GB, and it increases at a rate of roughly 0.6 GB per week, when storing only daily snapshots.

Nevertheless, some issues arose during the implementation of ARDA. As an example, the PCH website was constantly evolving forcing us to often rewrite our downloaders. Moreover, PCH route-collectors are not publicly associated with an IXP. Upon request, we were provided by PCH with this information. Managing the simultaneous run of computations scripts to keep the displayed results always up-to-date was also challenging.

Showcases of the relevance of ARDA

We explore some key features offered by two of the views of ARDA (Section 3.3.1.1): the *IXP view* and the *Regional view*. To begin with, the user can download the detailed list of values obtained for each metric, used to plot the graphs displayed in those views.

IXP View Table 3.12 summarizes the values obtained as of April 15, 2017 and September 18, 2017 for the metrics presented in Section 3.3.1.3.2.

ARDA provides a lower boundary of how many networks are peering at each African IXP and identifies those networks. Table 3.12 indeed shows that the African IXP having the highest number of members connected to its route-collectors is NAPAfrica Cape Town (124 members in April and 144 in September 2017) located in South Africa (ZA). The smallest number of members is 2, registered for SIxP (Sudan (SD)) and AIXP (Tanzania (TZ)) regardless of the month. On average, 21 members are peering at the studied IXPs as of April 2017; this number has increased to 24 in September 2017. Further, at 78.3 % IXPs, notably NAPAfrica, JINX (ZA), TIX (TZ), etc., almost all IXP members are connected to the deployed route-collector: it does not imply that each member peers with everyone at the IXP, however. For instance, one can notice that the number of peering ASNs at KIXP (Kenya, KE) found by ARDA in April 2017 (30) is close to that on KIXP website [287] (32). However, this is not the case of IXPN (Nigeria, NG) for which the number of detected peering ASNs (6 in April 2017) is really low compared to the 36 members listed on IXPN website [151]. Peers at IXPs in similar cases (whose names are not followed by a $\star$ in Table 3.12) need to remedy this situation. We remark that this was later corrected by the peers at IXPN in June 2017; consequently, the number of peering ASNs at the said IXP is 37 (identical to the ground truth) as of September 18, 2017.

ARDA also gives an insight into the origin ASNs seen at an IXP. As of April 2017, while for the category "peering ASNs," JINX is the runner-up IXP with 63 ASNs, it appears as the top local IXP for the category "origin ASNs" with 22,659 ASNs (Table 3.12). This number corresponds to roughly the 2/5 of the total number of networks composing the Internet during that period (57,015 ASNs according to CAIDA's inferred AS relationships [44]). Five months later, the highest number of origin ASNs (28,466) is seen at NAPAfrica (Table 3.12), which thus becomes the top local IXP in terms of peers (144). The highest amount of visible prefixes is also registered at that IXP in both April and September 2017 (with 160,418 and 212,885 prefixes, respectively).

We then compare the number of local ASNs (*i.e.*, origin ASNs assigned to the country hosting the IXP) to the number of external (*i.e.*, origin ASNs assigned to the country different from that hosting the IXP). We find that the percentage of local ASNs is low at IXPs where there are many visible networks (*e.g.*, roughly 0.76 % at JINX), and high where there are few (*e.g.*, roughly 70.1 % at CAIX).

Another functionality offered by ARDA is the ability to match ASNs visible at an IXP with reachable countries worldwide. Note that all IXPs selected in the following examples can be considered as mature Internet markets in the region, given their launch date. Let us split into

8 categories, the set of origin ASNs visible at KIXP (launched in 2002) and JINX (1996) as examples: *local AfriNIC ASNs*, which gather ASNs assigned to the country hosting the IXP; *external AfriNIC ASNs i.e.*, ASNs assigned to African countries different from the country hosting the IXP; *private ASNs*; *reserved ASNs*; *RIPE NCC ASNs*; *ARIN ASNs*; *LACNIC ASNs*; and *APNIC ASNs*. Figures 3.25a (left) and 3.25b (left) show that the percentage of KIXP-visible ASNs that belong to the category external AfriNIC ASNs (64.9 %) is higher than that seen at JINX in April 2017 (1.7 %). It is due to the considerable amount of origin ASNs from other regions (external origin ASNs) visible at the latter IXP compared to that of KIXP. Five months later, as the number of external origin ASNs visible at KIXP has increased from 365 to 6,690 (Table 3.12), one can notice by comparing Figures 3.25a (right) and 3.25b (right) that the fraction of external AfriNIC ASNs at KIXP has dropped to 2.8 %. Meanwhile, the percentage of KIXP-visible ASNs that belong to the category *ARIN ASNs* has drastically increased from roughly 5 % to 42.5 % of all origin ASNs seen at that IXP. The dynamics of the African IXP ecosystem are also noticeable at the above-listed IXPs when considering the evolution of the fractions corresponding to other categories of origin ASNs seen at the IXPs. As an example, the fraction of KIXP-visible ASNs that belong to the category RIPE ASNs has increased from 6.5 % in April 2017 to up to 18.5 % in September 2017. In the meantime, the percentage of KIXP-visible prefixes belonging to the local AfriNIC ASNs category (13.6 %), higher to that seen at JINX (1 %) as of April 2017, has drastically decreased to 0.2 %.

Comparing the pie charts of KIXP and JINX – Figures 3.25a and 3.25b – to those of CAIX (Egypt, launched in 2002), TIX (Tanzania, 2004) – Figures 3.25c and 3.25d –, and RINEX (Rwanda, 2004) hinted the existence of some policy issues at CAIX. In fact, no external AfriNIC ASNs are visible as origin ASNs at CAIX (noticed from August 2016 [86] to September 2017) contrary to the other IXPs, although KIXP, TIX, RINEX, and CAIX were launched in approximately the same period. After discussing with the CAIX operator, our hypothesis was confirmed. We were informed that CAIX *does not allow any member not operating in Egypt to peer at the IXP*. ARDA shows how this policy sadly limits the scope of CAIX (*cf.* Figure 3.26b (left)).

In addition, ARDA matches origin ASNs visible at each IXP to the countries they have been assigned to by their respective RIRs, and colors those countries depending on the percentage of allocated ASNs seen at the IXP. Such a feature could be of strategic importance when helping the Internet community to understand the reach of networks connected to a given IXP. Figure 3.26a highlights the results obtained in the case of KIXP (413 visible origin ASNs in April 2017 and 6,756 visible origin ASNs in September 2017). Figure 3.26a (left) shows that 50 % African countries had no ASN seen at the IXP in April 2017. Further, no ASN allocated to a country in North Africa (*NAf*) was directly peering or seen at the IXP: this may be due to the closeness of *NAf* to larger IXPs in Europe. The top five countries whose origin ASNs were visible at KIXP are ZA (69 ASNs), KE (48), the nearby countries TZ (35), Uganda (UG, 20), and finally Brazil (BR, 19). They represent respectively 20.8 %, 60 %, 59.3 %, 62.5 %, and 0.4 % of the ASNs assigned to the said countries. Few networks assigned to European and North American countries were

Table 3.12: IXP View: Overview of some metrics evaluated by ARDA per African IXP in the dataset as of April 15, 2017 and September 18, 2017. IXPs at which almost all members are peering with the route-collectors are followed by a ⋆. N/A stands for no data in the route-collector for the considered period.

IXPs (CC host) involved in the dataset	**Some metrics evaluated by ARDA, whose computations are described or whose values are referred to in Section 3.3.1.3**							
	#visible peering (origin) ASNs at the IXP		#visible local (external) origin ASNs at the IXP		#visible prefixes at the IXP		%IPv4 (%IPv6) blocks assigned to the country	
	15/04/2017	18/09/2017	15/04/2017	18/09/2017	15/04/2017	18/09/2017	15/04/2017	18/09/2017
Benin-IX⋆ (BJ)	5 (8)	6 (9)	3 (5)	6 (3)	176	186	45.8 % (0 %)	48.3 % (0 %)
BINX (BW)	6 (20)	6 (22)	10 (10)	13 (9)	210	212	64.9 % (0 %)	64.1 % (0 %)
CAIX⋆ (EG)	3 (67)	2 (65)	47 (20)	49 (16)	3,363	3,078	73.7 % (25 %)	72.8 % (21.4 %)
SIXP (GM)	6 (9)	6 (9)	7 (3)	6 (3)	66	68	60 % (0 %)	60 % (21.4 %)
KIXP⋆ (KE)	30 (413)	29 (6,756)	48 (365)	66 (6,690)	3,888	50,126	70 % (38.2 %)	68.5 % (25.4 %)
LIBERIA-IX⋆ (LR)	4 (8)	4 (9)	4 (4)	6 (3)	88	94	50 % (0 %)	57.2 % (0 %)
MGIX⋆ (MG)	5 (8)	6 (9)	2 (6)	3 (6)	183	576	50 % (0 %)	72.7 % (0 %)
MIX⋆ (MW)	N/A (N/A)	N/A (N/A)	N/A (N/A)	N/A (N/A)	N/A	N/A	N/A (N/A)	N/A (N/A)
MIXP⋆ (MU)	9 (12)	N/A (N/A)	7 (5)	N/A (N/A)	204	N/A	24 % (14.3 %)	N/A (N/A)
MOZIX⋆ (MZ)	12 (23)	13 (25)	12 (11)	14 (11)	339	861	62.5 % (0 %)	65.4 % (0 %)
WHK-IX (NA)	4 (8)	4 (8)	5 (3)	5 (3)	97	105	53.1 % (0 %)	53.1 % (0 %)
IXPN⋆ (NG)	6 (109)	37 (188)	77 (32)	120 (68)	1,503	2,264	49.8 % (0 %)	63.5 % (0 %)
RINEX⋆ (RW)	12 (85)	13 (93)	8 (77)	11 (82)	660	804	77.3 % (0 %)	75 % (10 %)
SIxP⋆ (SD)	2 (8)	2 (8)	7 (1)	6 (2)	675	692	70.4 % (0 %)	63.3 % (0 %)
JINX⋆ (ZA)	63 (22,659)	68 (25,063)	172 (22,487)	294 (24,769)	140,967	162,936	56 % (49.3 %)	61.1 % (47.9 %)
DINX⋆ (ZA)	15 (165)	20 (312)	58 (107)	166 (146)	1,263	2,462	14.4 % (7.6 %)	40.2 % (17.1 %)
CINX⋆ (ZA)	19 (464)	23 (451)	148 (316)	211 (240)	4,629	4,685	53.2 % (22.2 %)	51.2 % (20 %)
NAPAfricaCT⋆ (ZA)	124 (18,022)	144 (28,466)	171 (17,851)	258 (28,208)	160,418	212,885	46.4 % (29.9 %)	48.6 % (47.5 %)
NAPAfricaDB⋆ (ZA)	44 (401)	53 (445)	124 (277)	197 (248)	3,669	3,703	28.4 % (11.1 %)	29.9 % (15.4 %)
AIXP⋆ (TZ)	2 (42)	2 (42)	17 (25)	23 (19)	352	348	32.2 % (42.3 %)	28.2 % (34.1 %)
TIX⋆ (TZ)	36 (169)	37 (183)	39 (130)	56 (127)	1,324	1,496	78.3 % (50 %)	80 % (43.9 %)
TUNIXP⋆ (TN)	2 (24)	4 (29)	9 (15)	13 (16)	1,250	1,290	98 % (14.3 %)	98 % (10 %)
UIXP⋆ (UG)	24 (238)	21 (280)	17 (221)	24 (256)	2,287	2,495	72.1 % (18.8 %)	74.2 % (14.3 %)

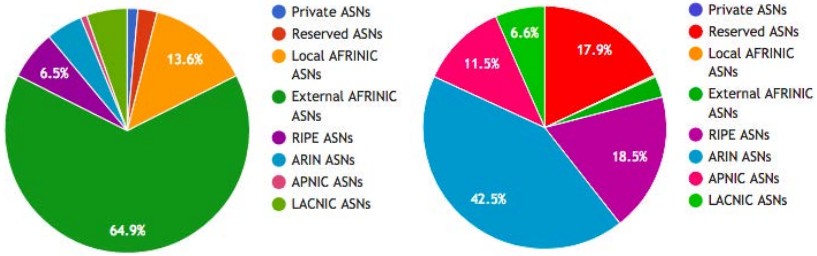

(a) *Percentage of ASNs per category visible as origin ASNs at KIXP as of April 15, 2017 (left) and September 18, 2017 (right). The differences between those two graphs are explained in Section 3.3.1.3 - IXP View*

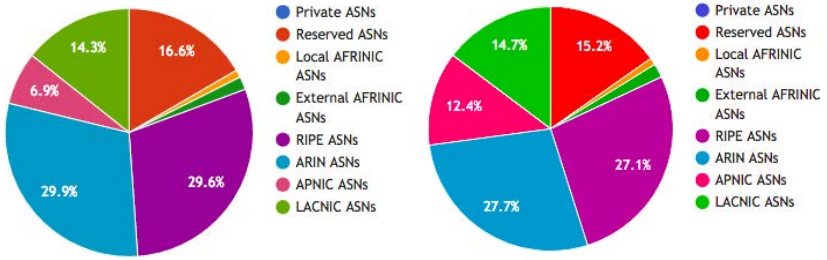

(b) *Percentage of ASNs per category visible as origin ASNs at JINX as of April 15, 2017 (left) and September 18, 2017 (right)*

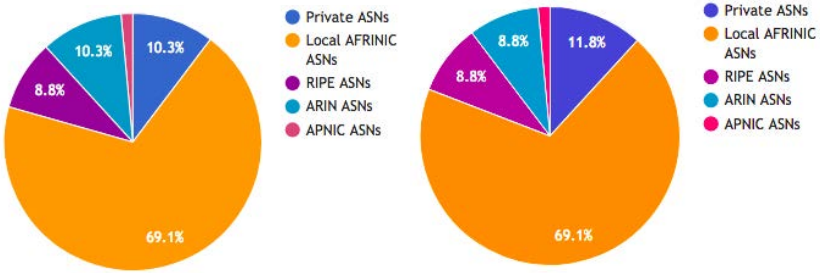

(c) *Percentage of ASNs per category visible as origin ASNs at CAIX as of April 15, 2017 (left) and September 18, 2017 (right)*

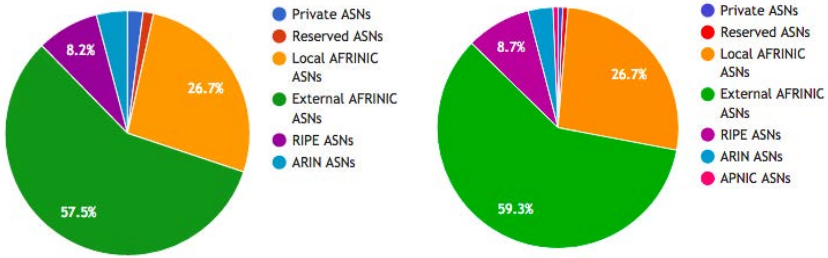

(d) *Percentage of ASNs per category visible as origin ASNs at TIX as of April 15, 2017 (left) and September 18, 2017 (right)*

Figure 3.25: Percentage of ASNs assigned by each RIR visible as origin ASNs at JINX (South Africa, launched in 1996) and KIXP (Kenya, 2002), CAIX (Egypt, 2002), and TIX (Tanzania, 2004) as of April 15, 2017 and September 18, 2017.

also visible, which may be of particular interest to progressive out of region networks looking to expand into East Africa (*EAf*). Figure 3.26a (right) shows that in September 2017, the reach of the networks peering at KIXP has improved significantly due to the increase in the number of origin ASNs seen at the IXP (Table 3.12). We can also specify that most of the new ASNs seen at KIXP are those allocated to the US: they correspond to 31.2 % of the ASNs visible at the IXP and 34.1 % of the number of ASNs assigned to the US. Contrary to April 2017, ASNs allocated to countries in *NAf* are seen as well; only three African countries still have no allocated ASN seen at KIXP (Chad, Eritrea, and Ethiopia).

Figure 3.26b (left), which highlights the reach of networks peering at CAIX (Egypt (EG)) contrasts with Figure 3.26b (right), that of IXPN (NG). 70.5 % of networks seen at CAIX are ASNs assigned to EG: no ASNs assigned to another country in the *NAf* region is seen at the IXP, consequence of the policy adopted by the IXP, which we discussed above. By contrast, networks assigned to countries in *WAf* and to ZA in *SAf* are seen at IXPN: this does not prevent 61.9 % of the ASNs allocated to NG to be visible at IXPN, but it enables regional interconnection.

Figure 3.26c compares the reaches of networks peering at UIXP (UG) and TIX (TZ), to show how well countries in the *EAf* and *SAf* regions are interconnected. After comparing to Figure 3.26b, one can deduce that CAIX and IXPN need to increase their marketing toward networks operating in other African sub-regions and different continents to expand their reach to such parts of the world. In case this strategy is well implemented, the reach of those IXPs will be similar to those of NAPAfrica and JINX (ZA), depicted in Figure 3.26d.

Regional View Our results show that, as of April 15, 2017, only 17.2 % of the ASNs assigned by AfriNIC are directly peering at any African IXP involved in this study. Meanwhile, 58.2 % of assigned ASNs are visible as origin ASNs at any such IXP. As of September 18, 2017, these fractions are 17.2 % and 57.6 % respectively. The remaining 41.8 % of ASNs correspond to African networks or content providers that are currently transiting local traffic. *Considerably reducing this percentage should constitute a concern for the Internet community*. Overall, our findings show that the level of regional interconnection has remained static over five months despite the growing and highly dynamic peering ecosystem in specific countries, such as KE and ZA.

Last, but not least, the Internet community may wonder (with the rising of concerns about the penetration of IPv6 due to the exhaustion of IPv4 blocks [242]) which percentage of IPv4 and IPv6 blocks assigned by AfriNIC are seen at any IXP in the region. We learn from ARDA that as of September 2017, in total 51.9 % IPv4 blocks assigned to any African country are seen at one or more local IXPs, whereas only 20.4 % of IPv6 blocks are seen.

Impact of ARDA on the Internet community A week after its launch [194], ARDA counted 389 users connecting from 155 ASes and located in 56 countries worldwide. Table 3.13 gives more details about the number of distinct IP addresses, which connected to it from April to September 2017, their corresponding ASes and CCs. Notably during the operator's meetings African Internet

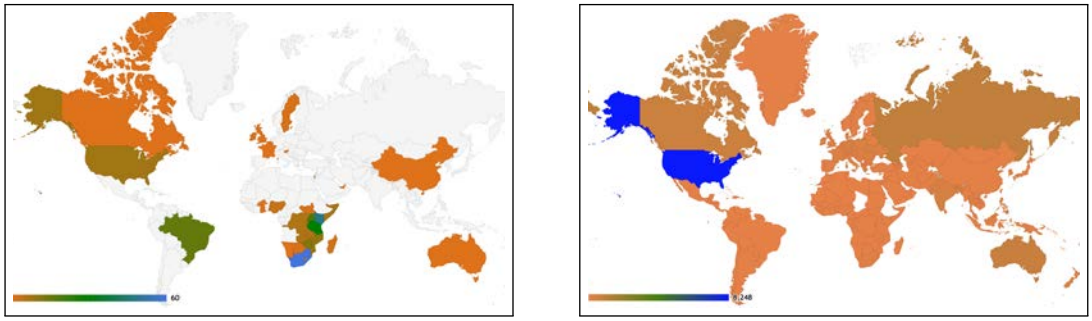

(a) *Percentage of ASNs assigned to each country by its RIR Visible at KIXP as origin ASNs as of April 15, 2017 (left) and September 18, 2017 (right)*

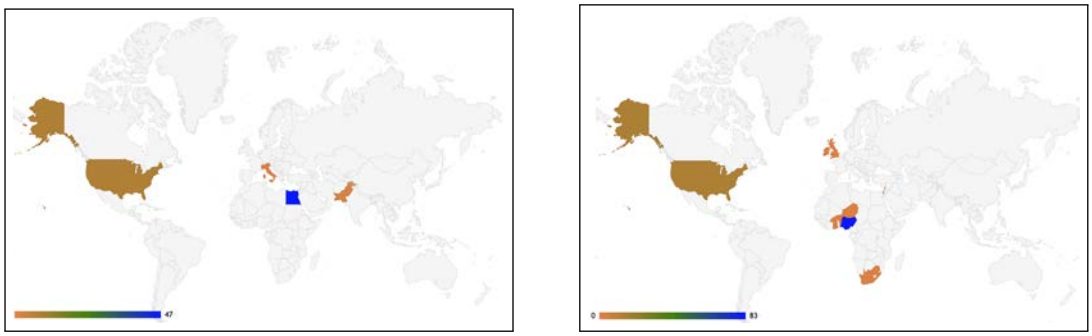

(b) *Percentage of ASNs assigned to each country by its RIR visible as origin ASNs at CAIX, EG (left) and IXPN, NG as of September 18, 2017*

(c) *Percentage of ASNs assigned to each country by its RIR visible as origin ASNs at UIXP, UG (left) and TIX, TZ (right) as of September 18, 2017*

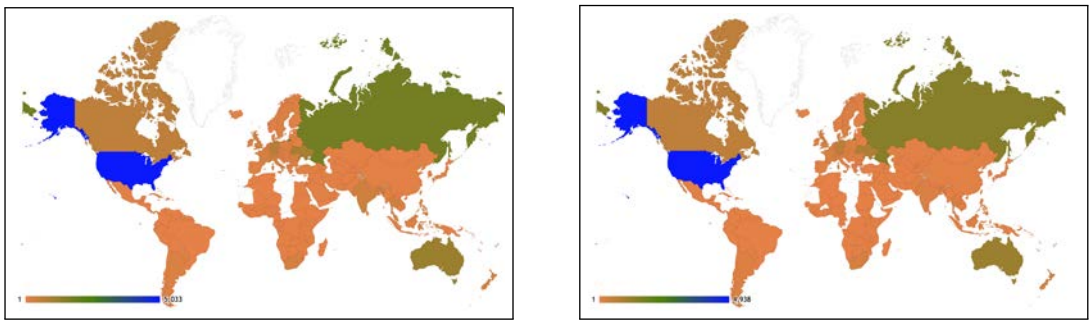

(d) *Percentage of ASNs assigned to each country by its RIR visible as origin ASNs at NAPAfrica, ZA (left) and JINX, ZA (right) as of September 18, 2017*

Figure 3.26: Percentage of ASNs assigned to each country (worldwide by its corresponding RIR), which is visible as origin ASNs at selected African IXPs.

Summit (AIS) 2017 and AfPIF 2017, several IP addresses from the local ASes offering connectivity and from the countries hosts were frequently connecting to the platform: these explain the peaks in the number of users in May-June and August 2017, respectively. More checks need to be done on those IPs and their activities on ARDA for separating bots from real end-users.

Table 3.13: Number of distinct IPs accessing ARDA from April to September 2017, their ASes and CCs

Month/Year	# distinct end-users' IPs accessing ARDA	# distinct ASes hosting end-user's IPs	# distinct CCs in which end-user's IPs are geolocated
04/2017	377	151	55
05/2017	458	168	50
06/2017	478	146	48
07/2017	423	131	48
08/2017	484	157	66
09/2017	489	123	45

3.3.1.4. Takeaways

As part of the ISOC strategy to allow the Internet community to monitor and understand the evolution of IXPs in a particular region, a route-collector data analyzer tool has been designed, and afterward it has been implemented, deployed, and tested it in AfriNIC. We have thus obtained the "African" Route-collectors Data Analyzer (ARDA), an open source web platform for analyzing publicly available routing information collected since 2005 by all PCH and RouteViews collectors with a peering viewpoint. ARDA provides metrics, which picture the status of the interconnection at local, national, and regional levels. Upon provision of their BGP feeds to a route-collector, local IXP participants are automatically taken into account. We have found that a small proportion of the ASNs assigned by AfriNIC (17 %) are peering in the region. Through them, roughly 58 % of all African networks are visible at one IXP or more. We have also noticed that these values have been static from April to September 2017. Next, we have shown how ARDA can help detect the impact of a policy on the growth of local IXPs, notably in the case of CAIX (EG). We believe that this tool will be a helpful compass in the quest for a better traffic localization or new interconnection opportunities in a given Internet region since it maintains in real-time, detailed and updated information on its IXP substrate.

Chapter 4

African Web Ecosystem

The surge in the deployment of IXPs [6, 292] at a significant rate, noticed in Section 3.2.1, and that of edge connectivity [1, 73, 111, 150, 163, 265] give insights into the rapid development of the Internet infrastructure in Africa. Despite these, Africa is far from achieving the online capacities enjoyed in the West, mainly because of the poor provisioning of content infrastructure in the region, which forces end-users to often fetch website content from the other side of the world [157]; there is little existing evidence to quantify this, however. It is thus essential that researchers and engineers begin to place more focus on *not only* underlying connectivity *but also* content infrastructure (*e.g.*, web servers, caches) in the region.

In spite of the several recent works measuring global web infrastructures [37, 76, 117, 123, 157, 210, 276, 315], none of them have (*i*) focused on developing regions like Africa; or (*ii*) explored if worldwide results apply to these areas. This leaves critical questions unanswered, primarily driven by the unusual make-up of African Internet infrastructures. First, the Internet in Africa is at a very different stage of its evolution: sub-optimal topology and peering configurations can make communications (*e.g.*, protocol behavior) very different [315]. Second, standard practices used for content provision (*e.g.*, placement of caches at IXPs) are difficult to apply due to the lack of IXPs that fulfill the requirements of CPs [81, 85, 117]. Third, hosting services are not as ubiquitous in the region, potentially making the management of web content much more complex [157]. Fourth, due to the lower level of Internet penetration and disposable incomes [139, 149], there are fewer (medium term) business incentives for optimizing web delivery. Again, the depth, veracity, and severity of this reasoning remain unproven. It is therefore essential to explore some of these factors, in an attempt to improve future ISP and CP deployments.

This section, which results from a deepening of our previous work [89], aims to offer a thorough understanding of the web infrastructure of Africa. We employ several measurement methodologies for exploring CPs and the configurations of network operators (Section 4.1.1). We start by analyzing traffic from a large European Internet eXchange Point (IXP) to determine a lower bound of the amount of traffic failing to be localized in the African continent (Section 4.2). We find that Africa still performs poorly with this measure. Despite the geographical distance, sig-

nificant amounts of African traffic are transited through Europe (even when the destination is another network in Africa). To help explain this, we focus on one of the largest CPs in the world: Google. After substantially improving our earlier geolocation methodology (presented in Section 3.2.1.2 and in [89]), we show that Google has made notable deployments in the region (Section 4.3.1). However, unlike their operations in Europe and the United States (US) where 90 % of caches have been mapped to Google's own AS [37], in Africa, they have primarily partnered with local network operators to deploy their caches. We find 1,067 functional caches in Africa hosted in 59 Autonomous Systems (ASes) and geolocated in 27 countries. Despite this achievement, roughly 48.3 % AfriNIC IPv4 prefixes still rely (exclusively or not) on North America for access to Google content (Section 4.3.1.1). By measuring redirections, we discover that local network operators tend not to serve each other. Significant inter-AS delays (caused by poor peering) mean that it is often actually more efficient to contact North America or Europe. It is particularly the case for countries in the *CAf* sub-region, which contain no Google Caches (GGCs). We further investigate other reasons for sub-optimal performance to find that various ASes have inefficient DNS configurations, using distant public resolvers that introduce significant delays to web fetches because of sub-optimal redirects and high-resolution delays (Section 4.3.2).

We then broaden our analysis to cover other popular global and regional CPs. Most are far behind Google in their support for African users (Section 4.3.3). Those popular providers, which include regional ones, have a very limited presence in Africa. Even the top local websites host their front-end services outside of the continent. Our traceroutes and Hypertext Transfer Protocol (HTTP) measurements, performed with RIPE Atlas probes scattered in the region, show how these decisions have severe performance implications for all web providers under-study. Our results leads us to highlight key lessons learned, as well as suggesting recommendations for improving future deployments (Section 4.3.4).

4.1. Active measurements and IP geolocation methodologies

4.1.1. Data collection

We begin by presenting our methodology used to analyze the nature and availability of content infrastructure. It involves three key steps: (*i*) collecting all IP prefixes for African networks; (*ii*) discovering all the content servers/caches that serve these networks; (*iii*) mapping the underlying path characteristics between users and the content infrastructure. All our measurement data is publicly available at [88] with the corresponding dates of their collection from 2015 to 2016. We further augment this data with traces taken from a large European IXP (*cf.* Section 4.2.1).

4.1.1.1. AfriNIC prefixes

To map content delivery infrastructure in Africa, it is necessary to compile a comprehensive list of the IP addresses and networks within the continent. To achieve this, we parsed the AfriNIC

IPv4 prefixes assignment and allocation files from 2005 to 2015 [11]. We could thus extract the list of IP ranges allocated by this RIR to local networks, as well as that of the countries to which they have been allocated. Among 3,488 available IPv4 prefixes, 3,082 of various lengths are assigned or allocated as of April 30, 2015. These are the prefixes we consider in this study; we term them *AfriNIC prefixes*.

4.1.1.2. EDNS0 client-subnet probes

Next, we collected a list of content caches that serve these *AfriNIC prefixes*. Since it would clearly be impossible to discover *every* cache, we focus on Google Caches (GGCs). Note that `www.google.com` is the top Alexa[1] website across the world and most African countries [18]. GGCs operate in a traditional CDN fashion: Whenever a client fetches a Google webpage, it is simply redirected, via DNS, to a nearby GGC.

To measure this, we used the EDNS0 client-subnet extension [210]. It has been developed to improve the accuracy of DNS-based redirections when a client is using a remote public resolver (*e.g.*, open DNS). The extension allows clients to include their network prefixes in DNS queries (the prefix length is determined by the recursive resolver). By doing so, CPs can redirect users to the correct server (rather than a location nearby to the public resolver).

We exploited this feature to launch EDNS0 queries with the client-subnet set to each of the *AfriNIC prefixes*, following a similar methodology to [37]. We could thus collect information on which GGCs end-users from across Africa are redirected to. We performed three EDNS0 crawls for `www.google.com`, using a variety of resolvers. First, we sent every hour on March 06, 2015, EDNS0 queries through Google public DNS (8.8.8.8). Second, we directed our queries through their name servers `ns1.google.com`, `ns2.google.com`, and `ns3.google.com` (all support EDNS0) every hour on April 12, 2015. Third, we sent again EDNS0 queries through `ns1.google.com` from April 23, 2015, to May 09, 2015 every hour. This revealed 3,011 unique GGC IP addresses, which we term the *EDNS0 probes* dataset.

4.1.1.3. RIPE Atlas DNS probes

A limitation of the above methodology is that we cannot be sure that the results returned via EDNS0 are equivalent to those that would have been returned to an actual client. To verify this, we augmented our dataset with a second set of DNS measurements. Towards this end, we used the RIPE Atlas infrastructure, as it is the largest open measurement infrastructure in the region. As of June 5, 2017, it has 527 VPs deployed in 231 ASes across 45 African countries (out of 58 African countries and neighboring islands) [248, 249].

We repeatedly launched, in parallel, six DNS requests of type A from all the available IPv4 RIPE Atlas probes in Africa to `www.google.com`. This was kept running for 7 days (from

[1] The platform `www.alexa.com` [18] is well-known for ranking existing websites worldwide and by region (as content providers can offer different services from a region to another).

March 24 to March 30, 2015). The active probes performed the query three times each, roughly every 60 s. We obtained 28,387,226 DNS queries.

Since not all the probes were online during the whole measurement campaign, our DNS lookups involve a total of 225 probes hosted in 38 African countries. 988 ASes have been allocated by AfriNIC as of May 07, 2015. After removing all the requests that have been performed by probes in Africa hosted in non *AfriNIC prefixes*, our DNS probes cover 111 AfriNIC ASes (11.2 %), and 146 *AfriNIC prefixes* (4.7 %). This constitutes the widest vantage on Google's infrastructure in Africa available yet. From this campaign, we obtained 1,917 GGCs IPs, which we term the *RIPE Atlas DNS* dataset.

4.1.1.4. Filtering inactive caches and private DNS resolvers

In total, we discovered 3,428 GGC IP addresses via our RIPE Atlas DNS and EDNS0 campaigns, since some IPs were in the outputs of both methods. Following the above, we performed 10 ICMP pings to each discovered cache to verify that it was active. We also issued HTTP requests towards all GGCs to check which ones were alive. These tests were performed from both Spain (ES) and the United Kingdom (UK) over multiple runs to ensure correctness (on March 09, April 09, April 13, and May 18, 2015). After discarding IP addresses that did not respond to either pings or HTTP requests, 3,120 IPs remained. We call this set of IPs the *functional GGCs*. RIPE Atlas probes also allow us to discover which DNS resolvers are used by African ISPs. We collected the IP addresses of all (239) default resolvers used by the probes. 70 of those IPs are RFC1918 private addresses (*e.g.*, 10.0.0.1); we discard these for the rest of this Chapter.

4.1.1.5. Measuring path characteristics

The above provides a comprehensive set of GGCs and DNS resolvers in Africa. Alone, it does not provide insight into the path cost for users though. We therefore launched from February 18 to May 22, 2015 a paris-traceroute campaign from all the RIPE Atlas probes in Africa to each of the GGCs IPs. A traceroute between each probe and each GGC IP is issued at five randomly defined timestamps during the said period. We used the User Datagram Protocol (UDP) [59]. The measurement campaign resulted in a total of 1,309,151 paris-traceroutes. It is important to emphasize that contrary to Gupta *et al.* [117] who performed traces towards GGCs in Kenya (KE), Tunisia (TN), and South Africa (ZA), our traceroutes targeted all the GGCs around the world, previously found to serve AfriNIC IP ranges. This provides a topology showing the routes and delays taken from African networks to the caches that serve them.

4.1.2. IPs geolocation

Before analyzing the collected dataset, it is essential to geolocate all discovered IPs. This is not trivial and is particularly difficult in Africa, which has seen less attention from mainstream geolocation research. Hence, our approach aims at gaining accurate location insight on all GGCs

and DNS resolvers. The first step of this approach to geolocating IP addresses gathers the methods used in Section 3.2.1.2.

4.1.2.1. Geolocation databases

We begin by using the traditional approach of geolocation data sources (DSes). To avoid problems found with individual geolocation Databases (DBs) [225], we use 10 public DSes to find the location associated with each IP. Similarly to Section 3.2.1.2, we selected: OpenIPMap (*OIM*) [245], MaxMind GeoIP2City (*MM*) [187], Team Cymru (*TC*) [286]; the delegated files of AfriNIC (*AF*) [11], APNIC (*AP*) [20], ARIN (*AR*) [23], LACNIC (*LAC*) [160], and RIPE NCC (*RP*) [252], as well as the RIR database *WHOIS*, and Reverse DNS lookups (*RDNS*) from which we infer the geolocation of an IP based on country codes (CCs), cities/airports names, or airport codes embedded in the reverse names. 1,357 GGCs and 103 DNS resolvers return a domain via a Reverse DNS lookup. Only 11.5% of the 3,120 GGC IPs have an airport or city code in their name. The remainder (88.5 %) contain no RDNS geolocation info and is composed of 14.6 % IPs with their names under the format of either `cache.google.com` or `google.cache.com`; 21.5 % IPs do not have any airport or city code in their name, whereas 63.8 % of IPs have not been resolved.

When all the DSes with an available entry for an IP give the same result, we use that country code (CC). But when this is not the case, we choose five random RIPE Atlas probes in each of the possible countries and perform three user-defined ping measurements towards the considered IP. We assume that the IP is located in the country with the lowest round trip time (RTT). For 42 % of GGC IPs, all the available DSes return the same country code. Amongst the remaining (1,812) IPs, only 1.1 % show an inconsistency of three countries, whilst the rest have an inconsistency of two. The delay tie-breaking approach allows us to geolocate a further 57.6 % of the GGCs. At the end of both steps, 99.5 % of functional discovered GGCs are geolocated. As for the DNS resolvers, all the available DSes return the same country code for only 15 IPs (9.5 %). We then apply the tie-breaking process for the remainder, thereby geolocating 91.7 % IPs.

We summarize the results of this first step in Table 4.1. The coverage column shows the percentage of IPs for which a Data Source (DS) has given an answer (*i.e.*, a valid CC). The Trust column shows the percentage of IPs for which the considered DS entry is equal to the country that we finally selected for that IP. Overall, the DSes are surprisingly accurate with many attaining a Trust above 0.9. That said, there are some significant outliers. *LAC* has no coverage, while some DSes such as *OIM*, *AP*, *RDNS*, *RP*, and *AR* have a very low coverage (*e.g.*, 10 % and below). *RP* and *WHOIS* are particularly poor. We notice, for instance, that 16.8 % of the answers from *RP* are "EU", whilst the final location is either in Ghana (GH), Tunisia (TN), or the Netherlands (NL). Similarly, although it has a high coverage (97.9 %), over half of the geolocations provided by *WHOIS* are inaccurate. These results highlight a key point: using these DSes in isolation would be very unwise in Africa.

Combining several DSes with latency-based measurements was sufficient to achieve accurate

Table 4.1: Comparison of geolocation data sources for both Google caches (GGCs) and DNS resolvers IP addresses as of October 2015. N/A stands for Not Applicable.

DSes	3,105 GGCs IPs		144 DNS resolvers	
	Coverage	Trust	Coverage	Trust
OIM	0.4 %	100 %	0 %	N/A
RDNS	8.3 %	93.8 %	0 %	N/A
MM	98.3 %	89.5 %	100 %	98.6 %
RP	10 %	75.3 %	12.5 %	88.9 %
AF	35.8 %	93.1 %	81.2 %	94 %
AP	2.6 %	100 %	0.7 %	100 %
AR	10.7 %	98.5 %	22.9 %	87.9 %
LAC	0%	N/A	0%	N/A
TC	98.97%	90.34%	100%	95.13%
WHOIS	97.93%	47.41%	94.44%	8.82%

geolocation in Chapter 3, which investigates the core of the Internet infrastructure in the region. However, it may not be the case anymore in this study that deals with the web infrastructure for which the addressing is different. We, therefore, undertook three more steps to verify the accuracy of our results, thus leading to a four-step geolocation approach. These are (*i*) speed of light sanity checks, (*ii*) multilateration geolocation, and (*iii*) final speed of light filtering.

4.1.2.2. Speed of light sanity checks

As a next step, we seek to filter any geolocations that show signs of discrepancies. We follow a similar strategy to [37] for filtering incorrect geolocations based on speed-of-light violations. Towards this end, we repeatedly launched from August 28 to October 18, 2016, (instantaneous) ping measurements from 100 RIPE Atlas probes randomly selected worldwide towards the geolocated GGC and DNS resolver IPs. In total, 2,217 IPs replied, resulting in 480,849 latency measurements. From these, we then extract the lowest RTT for each probe-IP pair, termed $Measured_{RTT}$.

Knowing that the signal is transmitted at the speed of $2c/3$ through optical fiber [235], we compute the minimum possible delay Min_{RTT} from each probe to the IP location as $3D/2c$. Note, D is the great circle distance between the coordinates of the probe (in km) and the geolocated IP; and c is the speed of light in the vacuum (in km/ms). In cases where $Min_{RTT} > (Measured_{RTT}/2)$, we consider the IP wrongly geolocated. Otherwise, the geolocation is (potentially) correct. 454 GGC IPs and 8 DNS resolvers IPs violated one or more of these speed of light checks, *i.e.*, about 20.8 % of the probed IPs.

In 87 % of the cases, the IPs whose geolocations were found to break/fail the speed of light test, were geolocated during the phase in which all DSes agree on the same CC for a given IP. The most common error is incorrect geolocation in the US: 385 GGC IPs out of 454 are wrongly geolocated in the US, while the remainder had been incorrectly geolocated in Mauritius (MU), NL, or the UK. Further, six DNS resolvers out of 8 are geolocated in the US, and the rest in MU. These findings illustrate how selecting the only available country code for a given IP can also

introduce discrepancies in the geolocation results.

4.1.2.3. Multilateration geolocation

The previous section highlighted a number of IPs that could not be correctly geolocated using geolocation databases (as shown via the speed of light checks). We next use multilateration with geographic distance constraints to address this [60, 116]. Multilateration is the technique adopted in the Global Positioning System (GPS), where satellites are used as landmarks. In our case, we consider all the RIPE Atlas probes (selected worldwide) involved in the previous latency measurements as landmarks (since we know their ground truth locations).

For each IP, our dataset contains a total of M landmarks sampled (*i.e.*, RIPE Atlas probes), ranging from 15 to 230, for all of which the GPS coordinates are known. We later check if the geolocation for each IP is the same by using M = 15, 16, 17, ..., 230 landmarks (randomly selected) in order to identify and remove cases of anycast IPs. We further report on the obtained results in the subsequent paragraphs. For all IP addresses, we compute the estimated physical distance D from each probe based on its measured RTT $Min_{RTTmeas}$. To this end, we use $(c \times Min_{RTTmeas})/3$. This produces an estimated radius, indicating the potential locations of the IP address (one radius per landmark). By then computing the centroid of the intersection of all radiuses from all landmarks, we can map the IP address to the corresponding CC [60, 116].

To perform this intersection and determine the geolocation of each IP, we first convert the GPS coordinates of all considered landmarks into Earth-Centered, Earth-Fixed (ECEF) coordinates. This information is stored into an $M \times 3$ matrice, P. We then compute the estimated physical distance (D) from each landmark to the IP with which we populated the $M \times 1$ matrice $Dists$. Next, we compute the least Squares solution of this $M \times N$ system to obtain the ECEF coordinates of the centroid [116]. After reconverting these ECEF coordinates into GPS ones, we can infer the CC of the IP.

To identify anycast IPs, we vary the number of landmarks M of each IP while running the aforementioned computation. Except for cases in which the IP is an anycast IP, or cases in which the intersection polygon is too large and covers many countries or islands, the CC obtained should be the same regardless of the number of landmarks. In cases where there is ambiguity, the IPs are removed from our data. 346 out of 2,217 IPs successfully pinged from our M landmarks have been geolocated using this methodology: The non-geolocated IPs correspond to cases in which the positions of the landmarks are not suitable for the circles to intersect. Amongst those 346 IPs, 171 are geolocated in only one country, regardless of the number of landmarks. We also noticed that, for example, Google DNS IPs "8.8.8.8" and "8.8.4.4" (both located by all geolocation DBs as being in the US) have different geolocations given the number of landmarks used, highlighting the fact that they correspond to anycast IPs.

Through this methodology, we have found 175 cases of wrong geolocations; we, therefore, removed these since they correspond to anycast IPs. Also, we corrected 69 previous wrongly geolocated IPs. At the end of this step, we could geolocate 2,732 GGCs and 151 DNS resolvers

IPs, corresponding to a total of 89.3 % of the discovered IPs.

4.1.2.4. Final speed of light filtering

As a final step, we repeated the speed of light checks using a separate testbed to identify any potentially erroneous geolocations from the previous section. We utilize three servers: *USserv*, which is known to be located in the US (California, San Diego), *ZAserv* in Africa (South Africa, Johannesburg), and *ESserv* in Europe (Spain, Madrid). From these three machines, we ping thrice all discovered geolocated IPs. We registered a total of 15,626 measurement outputs (2,219 IPs replied to our pings). As a last cross-check, we then apply the same speed of light test as that of Section 4.1.2.1. Next, we remove any GGCs and DNS resolvers that violate the new speed-of-light checks. 81 IPs are removed, leaving 2,654 GGCs and 148 DNS resolvers IPs. In total, we geolocate 86.8 % IPs of the discovered online GGCs and public DNS resolvers IPs. In the rest of this section, for any statistics related to only IPs and their ASes, we work with all (3,120 GGCs and 169 resolvers IPs) functional GGCs and DNS resolvers, while any statistics including geolocation results are computed for the portion of GGCs and DNS resolvers IPs that we could geolocate (2,654 GGCs and 148 resolvers IPs).

4.2. The need for a better traffic localization, seen from the VP of a large European IXP

Although there has been a wealth of studies looking at traffic from the vantage of European and US networks, we still know very little about the generation and treatment of African traffic. Thus, before diving into the nature of web infrastructure, we first inspect the *need* for improved Internet and web infrastructure in Africa by quantifying the amount of traffic that leaves the continent as seen from the vantage of a large European IXP data.[2]

4.2.1. IXP packet traces

The previous measurements are all active and give little insight into the traffic generated by African users. To address this, we augmented our data with packet traces collected from a large European IXP. Our goal is to explore (and exploit) the observation that large amounts of African AS paths traverse European IXPs [81, 85, 117]. We wish to verify this claim and quantify the potential benefits from localizing traffic within Africa. The collected traffic consists of almost 2 Terabytes of pcap captures from IPFIX records, covering five days worth of traffic (August 23 to 28, 2015). The IXP data is sampled 1 per 10,000 s and an approximation of the total traffic observed is given by multiplying the number of bytes in a flow by the inverse of the sampling interval [122]. In total, over 15 billion flows are seen.

[2] This traffic data analysis has notably been done in collaboration with Eder Leao Fernandes, Ph.D. Student at Queen Mary University London (QMUL, UK)

We then tag each flow with the specific RIRs that assigned its source and destination IP addresses [11, 20, 23, 160, 252]. Before doing so, we remove duplicates and overlaps (which are due to prefix transfers among RIRs or prefix resales among operators [242]) by considering that a given prefix is only operated by the last RIR to assign it. Clearly, this vantage point only provides us with a subset of African and regional traffic and, therefore, offers a biased sample point, notably due to the geographical location of the IXP (Europe), as well as the existence of several other large-scale IXPs in the same region. Nevertheless, it still provides a lower-bound vantage to underline the need for a better African traffic localization.

4.2.2. Does Africa have a traffic localization problem?

In Section 3.2.1.6, we argued that a major problem in Africa is the prevailing lack of peering, and the subsequent need for (Africa-to-Africa) traffic to be routed via remote transit networks, notably through European IXPs [81, 85, 117]. These results, however, were obtained using active traceroute measurements. We thus utilize our European IXP dataset to confirm the veracity of these assertions.

We compute, for comparison purposes, the total volumes of traffic exchanged between IPs allocated by each RIR as seen from this vantage point. Figure 4.1 shows the quantities of total traffic originated and destined to the same region traversing the IXP. This provides a crude measure of how efficient each Internet region is at localizing traffic, and avoiding intercontinental tromboning or remote peering [38], data that can only be used as a lower bound.

Unsurprisingly, it can be seen that the greatest traffic volume is exchanged between RIPE NCC (European) prefixes (*cf.* Figure 4.1). This is natural considering the physical location of the IXP. More unusually, we also observe a significant volume of ARIN-to-ARIN traffic (North America). Of more interest are the developing regions (AfriNIC, APNIC, and LACNIC), all of which can be seen to route non-negligible amounts of traffic through Europe in a circuitous manner. These observations confirm that there is a significant need for greater traffic localization, notably in inter-African networking. We find that African networks (1,273 ASes as of February 2017 [11]) could offload from intercontinental links at least 0.7 Gbps of traffic on average from this single IXP alone. This would lead to improved performance for end-users as well as significant transit costs savings, considering the expensive pricing of a 10 Gbps wavelength on major international routes linking Africa to Europe (US$112,500) compared to the pricing of those linking other continents [17, 208, 227].

4.2.3. Where is intercontinental African traffic destined to?

The above shows that the amount of Africa to Africa traffic traversing the studied IXP is not negligible. Before continuing, it is important to take a closer look at the destinations of traffic generated by AfriNIC prefixes. We next focus on the destinations of traffic originated and destined to IPs allocated by AfriNIC passing through the European IXP. We note that much of the physical

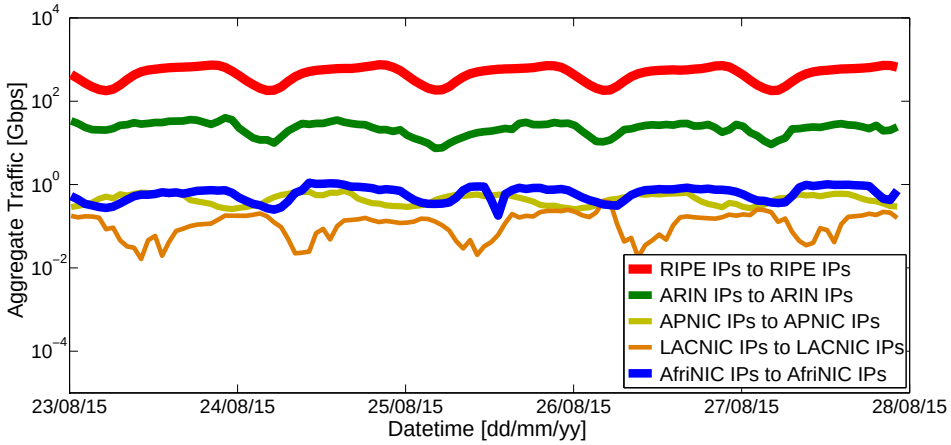

Figure 4.1: Volumes in Gbps of total traffic originated and destined to IPv4 and IPv6 addresses allocated by each RIR passing via the studied IXP.

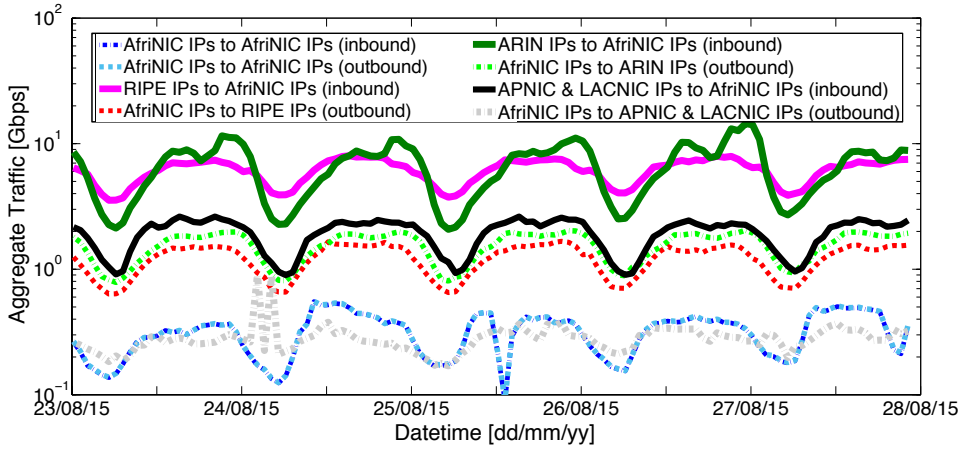

Figure 4.2: Volumes in Gbps of total traffic originated by AfriNIC IPv4 and IPv6 addresses and destined to IPv4 and IPv6 addresses allocated by each RIR (and vice-versa) passing via the studied large European IXP.

cabling connecting Africa to the world runs up through Europe [198], so it is safe to assume that our dataset contains a reasonable amount of traffic leaving Africa.

Figure 4.2 summarizes the results across the duration of the IXP dataset. The outbound traffic corresponds to the total traffic conveyed by the forward path, whereas the inbound traffic is that traversing the reverse path. As shown in the figure, the total volumes of traffic originated from and destined to AfriNIC IPs, which are exchanged via the IXP can be classified from the highest to the lowest in the order of the following RIRs: ARIN, RIPE, APNIC, or LACNIC, and AfriNIC IPs. The above shows that most traffic passing through the IXP (originated and destined to AfriNIC IPs) is actually exchanged with ARIN and RIPE IPs. Interestingly, despite the European location of the IXP, ARIN is the most popular destination. This is likely because of the bulk of web and service infrastructure hosted in the US [14]. Regardless, the analysis suggests that significant amounts of traffic and content consumed in Africa are sourced from outside of the continent. This observation indicates that the region could benefit greatly from more local content creation and

more local hosting of content and services. We largely inspect in the subsequent sections the current provisioning from an infrastructural perspective to understand the key deficiencies.

4.3. Deployment and utilization of the web infrastructure serving Africa

4.3.1. Exploring Google in Africa

Due to its scale and popularity, we start by mapping out the Google infrastructure used by African networks. The statistics presented in this section are computed based on the redirection of *AfriNIC prefixes* to any functional GGC from both our EDNS0 and DNS measurement campaigns.

4.3.1.1. Mapping Google cache locations

Overall we discover 3,120 functional GGCs serving Africa. However, when discussing CCs, we only use the 2,654 GGCs that we could correctly geolocate (contrary to the results presented in our previous work [89]). We first investigate the countries in which these GGCs are located, shown in Figure 4.3. We color code the locations: yellow markers represent GGCs hosted in RIPE NCC ASes, red ones are in ARIN, blue markers are in APNIC, and green ones are in AfriNIC ASes. The size of the marker is proportional to the number of IPs geolocated at that position. Table 4.2 also lists the top 10 ASes and countries in terms of cache numbers. The percentage between parentheses indicates the fraction of GGCs located in either the corresponding AS or country.

Figure 4.3: Geolocation of GGCs serving AfriNIC prefixes according to our refined geolocation methodology.

A range of ASes can be seen hosting GGCs. We discover 80 ASes in total, most of which are not owned by Google. 70.2 % of the ASes are allocated by AfriNIC, 22.6 % by RIPE NCC, 5.9 % by ARIN, and 1.1 % are APNIC ASes. However, most GGC IPs are in AfriNIC and ARIN

Table 4.2: Top 10 ASes and countries hosting GGCs IP addresses serving AfriNIC prefixes extracted from both DNS and EDNS0 methods. Parentheses contain the percentage of hosted GGCs.

Rank	AS (3,120 GGCs considered)	Rank	CC – Country (2,654 GGCs)
1	GOOGLE, US (37.2 %)	*1*	US – United States (31.8 %)
2	TMNET-AS-AP, MY (5.1 %)	*2*	MY – Malaysia (6.1 %)
3	YOUTUBE GOOGLE, US (4.7 %)	*3*	DE – Germany (5.5 %)
4	LEVEL3, US (2.6 %)	*4*	ZA – South Africa (5.2 %)
5	MEO-INTERNACIONAL, PT (2 %)	*5*	NL – Netherlands (4.9 %)
6	RETN-AS, UA (1.9 %)	*6*	EG – Egypt (4.5 %)
7	ROSTELECOM-AS, RU (1.5 %)	*7*	MU – Mauritius (2.8 %)
8	ETISALAT-MISR, EG (1.5 %)	*8*	IT – Italia (2.6 %)
9	TELECOM ITALIA, IT (1.5 %)	*9*	KE – Kenya (2.3 %)
10	MTNNS-AS, ZA (1.5 %)	*10*	NG – Nigeria (2.3 %)

IP ranges : Indeed, 40.2 % of the 2,654 functional GGCs belong to prefixes allocated by AfriNIC, whereas 32 % belong to ARIN. The rest (21 % and 6.5 %, respectively) belong to prefixes allocated by RIPE NCC and APNIC. African deployments have therefore deviated heavily from Google's prior setup in developed regions, which has seen Google hosting most (90 %) servers within its own networks [37]. Only 41.9 % of GGCs are hosted in Google ASes: 37.2 % in Google and 4.7 % in YouTube Google. All other caches are spread across third-party networks; prominently, TMNET-AS-AP has 5.1 %, and Level3 has 2.6 %. All other ASes contain under 2.5 % of the caches. We also find that many of the above ASes are based outside of Africa (≈30 %).

Compared to our results presented in [89], our new geolocation technique reveals there is a higher proportion of GGCs in Africa than in North America, while the percentages of GGCs in Europe and Asia have slightly increased (Figure 4.4a). Despite the efforts of stakeholders to keep local traffic local [6, 292], a large number of foreign caches are still relied upon though. 32 % of the 2,654 geolocated functional caches are in the US. As shown in Table 4.2, other prominent countries include NL, Malaysia (MY), and Germany (DE). Overall, 47 countries host a GGC: 27 in Africa, 12 in Europe, 3 in Oceania (Australia, New Polynesia, and New Caledonia), 2 in North America (the US and Canada (CA)), 2 in Asia (Malaysia (MY) and Bahrain (BH)), and 1 in South America (Peru (PE)). Africa contains only 40.2 % of all caches accessed by its users. Most are located in South Africa (ZA), Egypt (EG), Mauritius (MU), KE, and Nigeria (NG). An obvious reason for this setup is that Google's ASes seem to have only a marginal presence in Africa. We also highlight that, surprisingly, Africa is not particularly reliant on Europe for Google content. Only 21 % of caches are based in Europe, despite the closer geographic proximity than the US.

We also note that there are *no* caches in most countries of the *CAf* sub-region, *e.g.*, Democratic Republic of Congo (CD), Congo (CG), Gabon (GA), and Central African Republic (CF). Instead, caches are mostly based near the edges of the continent (as shown in Figure 4.3). This is likely driven by the expanding number of coastal submarine cables (inland cabling is much more expensive) [198, 277–279]. That said, we find that even well-meshed countries such as Angola (AO) and Namibia (NA) [198] have no GGCs. It is worth noting that not only our EDNS0 queries include all prefixes allocated by AfriNIC to the above-listed countries, but also some of the RIPE

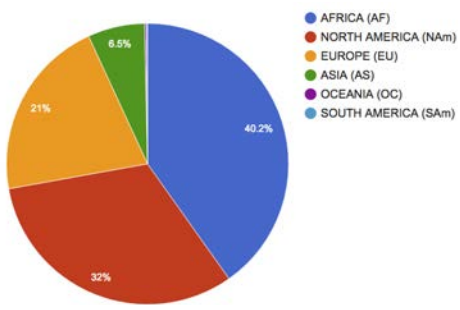

(a) *Locations of the 2,654 functional GGCs that are geolocated in Section 4.1.2.1*

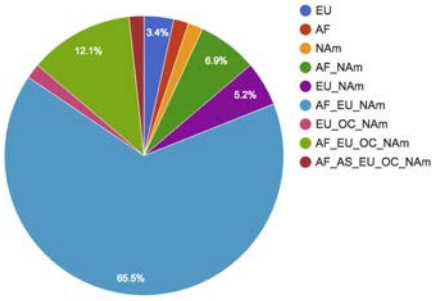

(b) *Percentage of African countries served by each continent*

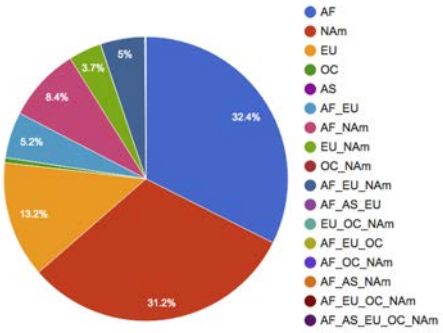

(c) *Percentage of AfriNIC IPv4 prefixes jointly served by each continent*

Figure 4.4: Statistics on Google redirections of AfriNIC IPv4 prefixes extracted from data collected through EDNS0 and DNS queries.

Atlas probes from which we launched our DNS queries are hosted in networks operating in those countries.

4.3.1.2. Mapping redirections

We next explore which caches end-users in Africa are redirected to: the presence of caches in North America and Europe is not important if they are only used occasionally. Figure 4.4 presents (*i*) the proportion of caches found in each continent, (*ii*) the percentage of countries that are served by various combinations of continents, and (*iii*) the percentage of *AfriNIC prefixes* served by various combinations of continents.

Figure 4.4a shows, as stated previously, that a significant number of GGCs are deployed in Africa (40.2 %). Nevertheless, 94.8 % of African countries are served by the US at least once in our dataset. In fact, Figure 4.4b shows that 65.5 % of countries spread their requests amongst Africa, Europe, and North America. This could be for many reasons, *e.g.*, using external caches to support "overflow," where demand exceeds local capacity. Figure 4.4b also shows that 12.1 % of countries are served by Africa, Europe, North America, and Oceania together. That said, we observe that 5.2 % of countries are exclusively served by North America and Europe. In fact, Mayotte (YT), though being an island nearby Comoros and Madagascar, is solely served by North

America, indicating that this is not caused by the need for an "overflow". In that case, YT does not host its own GGC, forcing it into using external caches. Ideally, end-users in that country would be redirected to other nearby African countries but, clearly, certain reasons (later explored in Section 4.3.1.4) prevent this.

Comparing Figures 4.4b and 4.4c also highlights some interesting properties. Whereas the bulk of requests on a per country basis are redirected to North America, Europe, and Africa, this is not the case on a per network basis. Only 1.7 % of *countries* solely use North American caches. In contrast, 31.2 % of *networks* solely rely on North America. Further, while *only* 1.7 % of countries are exclusively served by African caches, we find that 32.4 % of networks are. In other words, redirection is primarily based on specific networks rather than countries. This means that many networks fail to gain access to caches located in Africa, even though others in their country can do so. *Choosing the "right" ISP, therefore, seems particularly important in this region.*

4.3.1.3. Cache sharing

We next inspect in what circumstances countries and networks share their caches with others. It is particularly pertinent in Africa, as recent work has highlighted that network operators are often resistant to cooperate [117]. Note that sharing is a product of both individual network policy and redirection strategies employed by Google. Figure 4.5 compares the number of caches within each country against the number of African countries that use those caches. It includes the percentage of other countries that the GGCs are shared with, when considering only the top 35 countries hosting a GGC: African GGCs host countries are in green, whilst GGCs host countries on other continents are in black. Theoretically, if cache deployment were ubiquitous, each country should only need to serve requests from its own residents. In such a case, the number of countries mapped to a GGC (*i.e.*, the blue line) should always be 1. Figure 4.5 shows, however, that this is not the case. In total, 60.6 % of countries found to host GGCs share their caches with at least one other country. Indeed, 57.9 % of African countries (hosting GGCs) share their caches with other countries, whilst this percentage is 81.8 % for those outside Africa.

Unsurprisingly, the most extreme is the US (845 caches), which serves almost all African countries (54). This is dominated by Google's US-based ASes. Similarly, in Europe, 48 African countries are served by DE (147 caches). As shown by red squares in Figure 4.5, Italia (IT) serves 32 African countries with its 69 caches, while NL serves 16 countries with its 130 caches. Countries outside Africa share their caches, on average, with 15 other countries, compared to just the half by African countries. In Africa, sharing is primarily performed by more developed states, *e.g.*, ZA (serves 14 countries with 139 caches), MU (serves 13 countries with 75 caches), and KE (serves 5 countries with 62 caches). In contrast, many less developed countries have very different trends. There are countries, which host a large number of caches, yet only serve one other country: *e.g.*, Zimbabwe (ZW), which contains 45 caches, Mozambique (MZ) 30, and Cameroon (CM) 30. Meanwhile, countries such as TN, Morocco (MA), Algeria (DZ), Tanzania

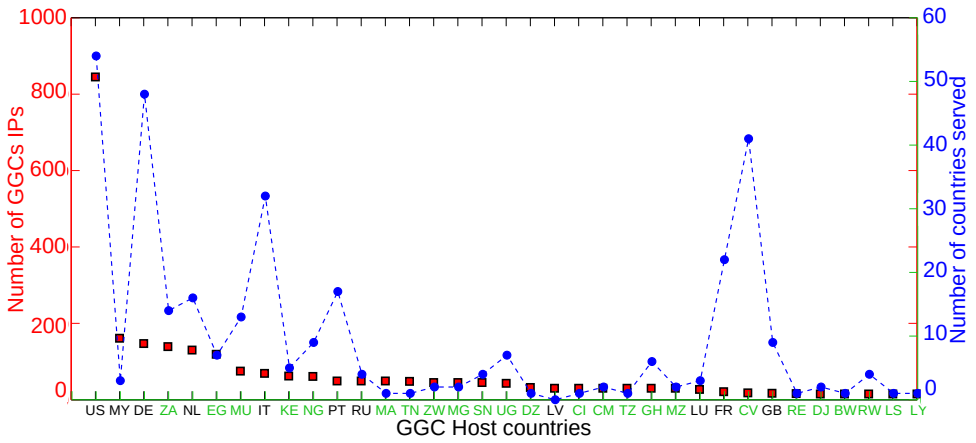

Figure 4.5: Distribution of GGCs serving AfriNIC prefixes across countries.

(TZ), and the Ivory Coast (CI) never serve a user in another country.

Table 4.3 compares the percentage of GGCs in a country against the percentage of requests redirected to that country (last column). Given the fact that we suppressed any IP suffering from problematic geolocation in Section 4.1.2, 77.4 % of the outputs of our EDNS0 probes and 49.3 % of the DNS queries are covered. A proportional and cooperative redirection strategy would result in the percentage of GGCs in a country and the percentage of requests redirected to that country being identical. This, however, is not the case. Clear trends can be seen, with 31.8 % of caches in the US receiving 33.6 % of our requests from Africa when considering EDNS0 probes. We notice that caches in DE (5.5 % of caches) receive 12.7 % and 25.4 % of requests for EDNS0 probes and DNS queries, respectively. Caches in these countries, therefore, serve a disproportionately large number of requests. In contrast, SC and ZA are the only African countries that service about 10 % of the requests. The rest service low proportions (5.5 % and below). Hence, despite wide deployment, African caches do not receive a fair proportion of requests.

Of course, the lack of sharing among caches in Africa while servicing requests from the continent that we highlighted above is driven by individual networks, rather than entire countries. 15.1 % of the networks containing our RIPE Atlas probes host a cache. Only 63.1 % ever share their caches with others. For instance, in the collected dataset, ASes Utande Internet Services (ZW), Ubuntunet (TZ), GULFSAT-AS (MG), and RAYA Telecom (EG) never serve other networks. It is impossible to concretely state the reason; however, we conjecture that it is a combination of both well reported inter-AS performance issues [37, 81, 85] and network operator policy. We analyze the former in Section 4.3.1.5, but the latter highlights a key problem faced in Africa, where it is often challenging to initiate cooperation across organizations and countries [30, 207].

4.3.1.4. Understanding disincentives for sharing

The above raises questions about *why* caches in Africa are not typically shared across networks. Our analysis suggests that a key reason is that many African networks still remain disconnected from local IXPs [81, 85]. Sharing cache capacity would, therefore, generate transit costs,

Table 4.3: Percentage of total redirections towards GGCs in top 10 countries hosting caches, computed based on outputs from EDNS0 probes from all AfriNIC prefixes and DNS queries from RIPE Atlas probes.

Rank	CC	Country	% caches hosted	EDNS0 probes	DNS queries
1	US	United States	31.8 %	33.6 %	14.8 %
2	MY	Malaysia	6.1 %	0.07 %	0.04 %
4	DE	Germany	5.5 %	3.6 %	25.4 %
5	ZA	South Africa	5.2 %	12.1 %	11.3 %
3	NL	Netherlands	4.9 %	1.9 %	0.8 %
6	EG	Egypt	4.5 %	3.7 %	0%
7	MU	Mauritius	2.8 %	5.3 %	2.1 %
8	IT	Italia	2.6 %	1.7 %	4.8 %
9	KE	Kenya	2.3 %	3.5 %	0.3 %
10	NG	Nigeria	2.3 %	8%	0.008 %

suffer from high inter-AS delay and, consequently, reduce the probability of a CDN redirection algorithm selecting a non-peered neighbor. In order to explore this, we collect information on IXP peering from IXP websites, PeeringDB and Packet Clearing House (PCH) [218, 220, 292].

The said piece of information reveals that most networks sharing caches are peered at IXPs. For example, 99.9 % of the requests served by DE caches are redirected to networks peering at DE-CIX in Hamburg; all redirects to the UK go to Google's own AS peered at the LONAP IXP; and 99.7 % of redirects to NL go to third-party networks peering at AMS-IX. Similarly, 99.9 % of redirects to the US go to peers of one of 33 US IXPs. In these cases, sharing cache capacity is straightforward, as IXP membership allows low-delay, low-cost interactions between networks. To explore this in Africa, we use our paris-traceroute dataset to check if the African networks sharing their caches are peered at IXPs. We find that all African ASes connected to an IXP share their caches. The top two networks for sharing are in ZA (MWEB and InternetSolutions). Unfortunately, only 18.6 % of African ASes found by our measurement outputs to host a GGC are peered at an IXP. This means that for the remainder, sharing their caches would generate transit costs. Further, the higher inter-AS delays would drive Google's redirection algorithms away from selected non-peered networks. Nearly all redirects that stay within Africa are between networks peered together at an IXP. This strong correlation suggests that *the main barrier to unlocking significant web performance improvements in Africa is actually to enable cache sharing via peering.*

4.3.1.5. GGC performance

Finally, we wish to quantify the performance of Google in Africa by measuring the delay between the RIPE Atlas probes and the GGCs (Section 4.1.1.5). As three RTT values are recorded per latency measurements, we extract the minimum RTT for each probe to measure the best case scenario. Figure 4.6a shows a CDF of the minimum RTTs to the GGCs measured over each probe in our dataset. Remarkably, the web requests to caches in Africa attain a mean of 223.7 ms (and

a median of 193.9 ms) (*cf.* Figure 4.6b) compared to caches in other regions for which registered RTTs are lower. As an example, RTTs to caches in South America have the lowest mean RTT of 89.9 ms (median of 53.4 ms). These confirm that CDN redirection algorithms are right to avoid sending users to other African networks, regardless of their geographical closeness.

Delays to Europe are high (with an average of 124.2 ms and a median of 137.2 ms), but lower than those to African caches. Only caches in Asia (284.1 ms average and 297.4 ms median) perform worse than those in Africa while serving African end-users. The key exceptions to these observations are African networks that host their own cache, which are thus reachable by their end-users with an average minimum RTT of 179.1 ms (median of 75.5 ms) compared to 251.4 ms for those without (median of 201.6 ms). This confirms that *the sub-optimality found in African topologies [37] impacts the ability of caches to be locally used/shared within a reasonable delay bound.*

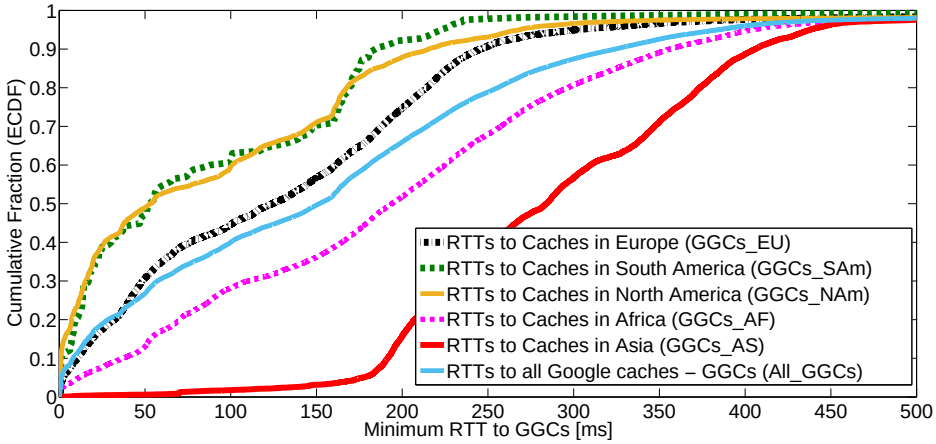

(a) *CDF of the minimum RTT distribution*

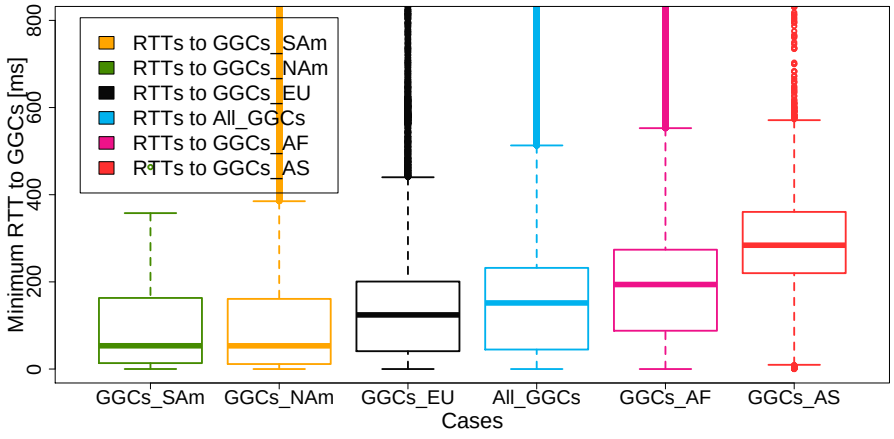

(b) *Boxplot of the minimum RTT distribution*

Figure 4.6: Delay distribution from different sets of RIPE Atlas probes in African networks to serving GGCs. The cases listed in Figure (b) correspond to those in the legend of Figure (a) and their respective colors are identical.

4.3.2. DNS in Africa

A critical part of web behavior is DNS (which is typically used by CPs for redirection). Hence, we explore the DNS configurations used by African networks.

4.3.2.1. Mapping DNS resolver locations

The RIPE Atlas probes allow us to discover which DNS resolvers are used by African ISPs. We collect the IP addresses of all (239) default resolvers used by the probes. 70 are private addresses (*e.g.*, 10.0.0.1); we discard these for the rest of this section. We then geolocate 87.6 % of the remaining resolvers using our methodology presented in Section 4.1.2. Our results show that the majority are based within Africa (as expected); however, 2.1 % located outside of the continent.

It has previously been found that non-local resolvers can adversely impact CP performance [210]. In total, 83.8 % of resolvers are hosted within the same network as the probe. This is ideal for CP redirection, as the CP would be able to effectively locate the client (using the IP address of the DNS resolver). Nevertheless, 16.2 % of unique resolvers are hosted within different networks. Furthermore, 34.6 % of all the probes share these resolvers located in different networks, showing that many ISPs utilize third-party resolvers by default. We observe that these ISPs use DHCP to automatically configure clients to use third-party resolvers. The reason for ISPs adopting this behavior is generally easier management — clearly attractive in the African region.

It, however, comes at the cost of performance for CPs [53], since their clients would appear as if they were in a different network (where the resolver is). In 32.5 % of cases, the third-party DNS resolver is not even in the same country. This is reflected in the geographic distances observed between our probes and the resolvers. On average, the third-party resolvers are 13,690 km away from the probes they serve (distances ranging from 996 km to 18,116 km). In contrast, ISPs using local resolvers have distances ranging from just 0.07 km to 3,554 km (average 325 km).

4.3.2.2. DNS resolver performance

By using distant DNS resolvers, it is possible that significant start-up delays may be introduced for web fetches. Third-party resolvers hosted in other countries have an average delay of 129 ms compared to just 25 ms for resolvers hosted by the ISP. To explore this further, we split the DNS queries into two categories: those sent to resolvers in the same country (67.5 %) and those sent to resolvers in different countries (32.5 %).

The first category is composed of DNS queries sent to (*i*) ISP resolvers located in the same country (86.1 %); and (*ii*) open resolvers in the same country (13.9 %). The second category is composed of DNS queries sent to (*i*) open DNS resolvers (0.8 %); (*ii*) open resolvers in different countries (4.1 %); (*iii*) ISP resolvers located in different countries (15.1 %); and (*iv*) Google DNS (80 %). Figure 4.7 presents the resolution delay distributions.

The average response time of third-party resolvers in different countries is 132 ms. Meanwhile, the average response time of local resolvers in the same country is 25 ms. The best performance is naturally attained by resolvers in the local ISP, with marginally worse performance provided by third-party resolvers in the same country. The most significant drop in performance is introduced by public resolvers such as Google DNS. Although they are presented as methods to improve performance, this does not work in Africa due to the lack of public resolver infrastructure on the continent. For instance, around 50 % of Atlas probes suffer from an addition of over 100 ms delay when redirected to distant Google DNS resolvers located in the US. Some African operators are therefore outsourcing not only the hosting of web content but also the operation of key infrastructure such as DNS.

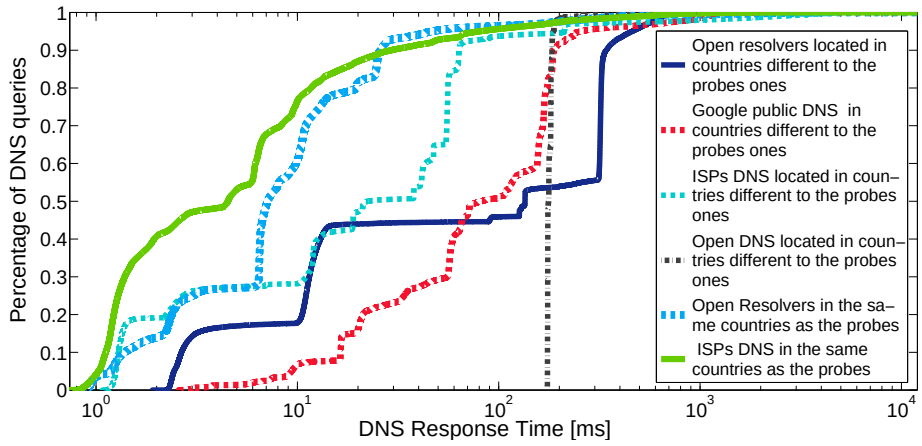

Figure 4.7: Cumulative distribution of DNS resolution delays.

4.3.3. Expanding to other Content Providers (CPs)

So far, Google has been focused on. Next, we expand our analysis to a variety of other popular websites.

4.3.3.1. Measuring top websites

To compile a list of popular websites, we took: (*i*) the top 10 global Alexa websites, (*ii*) the top 15 Alexa websites in Africa, (*iii*) the top 15 most popular websites in Africa listed by `afrodigit.com`, and (*iv*) `iroking.com`, a well-known video content provider on the African continent. We included websites from Afrodigit because we noted that the top Alexa websites were biased towards websites in certain countries (*e.g.*, South Africa, Nigeria, Egypt). We also added `iroking.com` to gain an understanding of video websites in Africa (because there are no local videos content websites in either the top Alexa or Afrodigit websites). Again, we utilize DNS to discover their front-end infrastructures. We concurrently issued DNS queries from RIPE Atlas probes to each of the domains over a four day period on a per hour frequency (May 23 – 26, 2015). This allowed us to observe the location of front-end servers hosting the websites using

our method from Section 4.1.2. In total, 566,994 DNS queries were launched.

Table 4.4 compares the sizes, the server geolocation, and the networks hosting the websites: note that the websites are classified by their content type. Surprisingly, only five websites from the 18 regional ones actually operate their servers in Africa. This is probably attributable to the more reliable and cheaper foreign hosting available [157]. It can also be explained by the significant inter-AS delays, due to which it is often actually more efficient (in terms of delay/QoS but not in terms of cost) to contact North America or Europe. The five sites hosted in Africa are in ZA, within four ASes. The remainder are in the US or Europe, with common platforms like Amazon and CloudFlare dominating. In terms of hosting practices, all of the African websites we measured (from the vantage of the 146 *AfriNIC prefixes* hosting our probes) used a single AS to host their content.

In contrast, the top global Alexa websites seen from our probes have a more distributed infrastructure. They are generally hosted in multiple countries and ASes. That said, we do not see any others achieving the distribution of caches that Google has in Africa. For instance, `facebook.com` only reveals five front-end IP addresses serving content for our probes (all hosted in Facebook's AS). Unlike Google, Facebook does not host within African networks, instead placing their infrastructure at their own points of presence [123]. Similar results are found across all global Alexa websites. For instance, `yahoo.com` serves our probes located in Africa from the UK and the US (both hosted in Yahoo's AS), and `amazon.com` serves our probes from the US (via Amazon and LimestoneNetworks). That is, the deployment of Google in Africa is *not* the norm. An interesting case was `taobao.com`, which we found to serve our probes from 15 caches hosted in three countries, namely ZA, CN, and the UK. They were found to belong to four ASes of which a South African AS, Vodacom (ZA); the remaining ASes were Level3 (US), Chinanet (CN), and CHINA169-BACKBONE (CN).

4.3.3.2. Website performance

We next expand upon the previous delay measurements (Section 4.3.1.5), to explore the HTTP performance characteristics of all websites studied. To gain a comparative benchmark, we augmented our African RIPE Atlas probes with 242 extra probes randomly chosen from Europe. We launched HTTP requests every 12 hours during the period June 2 – 5, 2015 from every probe to every website's homepage. To reduce the impact of differences in page size and third-party objects, we only fetched the homepage HTML; we did *not* request images, adverts, javascript, etc. This results in a mean page size of 169 KB, with a standard deviation of just 166 KB (we include website size in the figures). Figure 4.8a shows the minimum time to fetch the global Alexa websites from each probe (measured by the length of the TCP connection). Again, we take the minimum to observe the best case scenario for each probe.

We first inspect Google, which obtains very different page loads across the probes: load times vary from 2 ms to 1,250 ms with a mean of 200.9 ms. This is partly caused by the existence of GGCs in a subset of the networks hosting our probes. The median load time in networks hosting

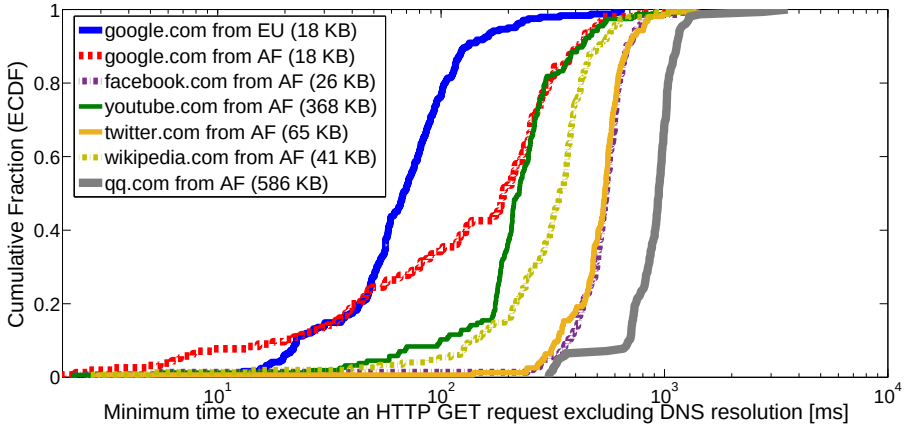

(a) *Distribution of minimum time to execute an HTTP GET request per probe (ms) from Europe (EU) and Africa (AF) to top global Alexa websites.*

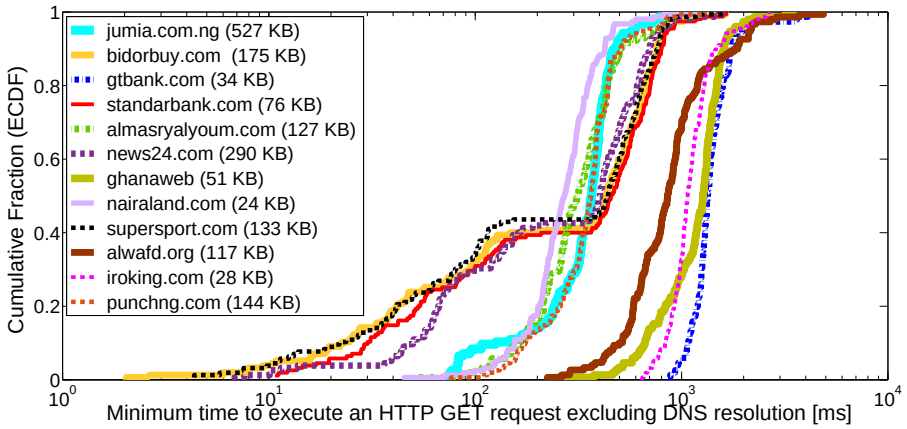

(b) *Distribution of minimum time to execute an HTTP GET request per probe from Africa to selected top local Alexa & Afrodigit websites.*

Figure 4.8: HTTP fetch time for top global and top regional websites from RIPE Atlas probes (website sizes are in parentheses).

a cache is just 148 ms compared to an overall median of 190.2 ms. Moreover, 60.7 % of probes in ASes hosting GGCs have a delay that is below the average for the continent. However, overall, only 26.2 % have a delay that is below that of the median seen in Europe (67.6 ms), and only 32 % have an HTTP performance below its mean (84.6 ms). This is not simply caused by the high DNS resolution times previously reported. Even when ignoring the DNS resolution times, we notice that only 35 % of probes in Africa fetch `google.com` in under 100 ms; this value is 78 % in Europe. Furthermore, the average of the HTTP performance from Europe to Google is more than twice that experienced from Africa. For medians, it is thrice.

In comparison, the other websites seen from Africa on Figure 4.8a have greater density around the mean (indicated by a sharp upturn in their CDF). This is because their infrastructures are not as well distributed in the region as that of Google. Consequently, most end-users in Africa have similar performance to each other. The median of the HTTP requests performed by the RIPE Atlas probes hosted in African networks is 223.8 ms towards `youtube.com`, 339.8 ms

Table 4.4: The sizes and locations of the infrastructures of the top 15 websites in Africa (by Alexa & Afrodigit), and top 10 global sites (Alexa).

Top 15 sites in Africa (by Alexa & Afrodigit)	Type	#IPs caches	CCs host caches	ASes	Top 10 global web-sites (by Alexa)	Type	#IPs caches	CCs host caches	#ASes
jumia.com.ng	E-commerce	1	DE	20546	amazon.com	E-commerce	4	US	2
konga.com	E-commerce	1	US	15169	taobao.com	E-commerce	15	ZA, UK, CN	4
bidorbuy.co.za	E-commerce	1	ZA	3741					
fnb.co.za	Financial services	1	ZA	17148	qq.com	Internet services	2	CN	2
gtbank.com	Financial services	1	US	26496					
absa.co.za	Financial services	1	ZA	3741					
standardbank.co.za	Financial services	1	ZA	10798					
almasryalyoum.com	News/media	5	NL, CR	13335	google.com	Search engine	924	18 (§ 4.3.1.1)	26
elkhabar.com	News/media	2	US	13335	yahoo.com	Search engine	4	US, UK	2
vanguardngr.com	News/media	1	US	14618	baidu.com	Search engine	1	HK	1
news24.com	News/media	1	ZA	10474					
punchng.com	News/media	1	IE	16509	wikipedia.com	encyclopedia	2	NL, US	2
iol.co.za	News/media	2	IE	16509					
ghanaweb.com	News/media	1	US	7859					
nairaland.com	Online community	5	US	13335	facebook.com	Social network	5	US, DE, NL	1
supersport.com	Sports	1	ZA	10474	twitter.com	Social network	7	US	2
alwafd.org	Politics	2	NL	13335					
iroking.com	Videos	2	IE	16509	youtube.com	Videos	41	SN, MU, US	3

towards wikipedia.com, 540 ms towards twitter.com, 549.1 ms towards facebook.com, and 943.41 ms to qq.com.

Figure 4.8a can also be compared to Figure 4.8b, which presents the same data for the top African websites (from Alexa and Afrodigit). We find that the top African websites get approximately equivalent performance to the top global websites, suggesting that these regional services have made little effort to optimize their local distribution on the continent. The regional websites on Figure 4.8b can also be separated into roughly three groups of varying load times. We note that the ones gaining highest performance are predominantly hosted on the continent, *e.g.*, supersport.com and standardbank.co.za, confirming the benefits that could be gained by services locally. In all cases, these websites are based in ZA, where infrastructure is well developed and affordable. Unfortunately, the worst performing local websites get even lower performance than the globally popular equivalents, indicating that they are not well provisioned. Unsurprisingly, they correspond to those that are based in either the US or Europe. An obvious takeaway message is that *these websites should aim to host their content locally*. In the future, as inter-AS connectivity improves, the increase of sharing caches across networks (via IXPs) could hopefully incentivize this (*cf.* Chapter 5).

4.3.4. Discussions

This section has explored the deployment of web infrastructure in Africa. Whilst we have measured the African interdomain routing in Section 3, we argue that this only addresses a subset of the challenges, as it does not take into account the web infrastructure.

We have shown that Africa is far from being self-sufficient regarding its hosting infrastructure. We have begun by inspecting packet traces from a large European IXP to witness notable amounts of traffic failing to still be localized in Africa. This has inspired us to study Google's deployment, which we have found to route significant amounts of Africa-destined traffic through Europe. Although we have discovered caches across half of the African countries, we have found that US infrastructure is regularly used. Unlike Google's global footprint, these African caches are largely based in third-party networks, which nearly always exclusively service their own subscribers. Only those connected via local IXPs (*e.g.*, JINX, CINX, TIX, or NAPAfrica) break this trend. Due to poor peering, we have found that, in many cases, reaching a geographically nearby African cache actually has a higher delay than contacting the US. As such, sharing cache capacity across networks can only work with improved operator cooperation [30, 207].

That said, we have found that Google is considerably more developed in Africa than other providers. We have then analyzed both global and regional websites to find that even local websites are hosted outside of the continent. In fact, only five out of the 18 regional website front-ends surveyed are hosted locally (all in ZA). The cheaper cost of hosting abroad and the significant inter-AS delays amongst African ASes are two possible reasons for this. In all cases, we have found clear trends showing that these hosting decisions have negative implications for performance. We have consistently observed higher HTTP load times for non-Google websites hosted

outside of the continent. For those hosted within the continent, we have seen a roughly consistent performance, although it is not yet equivalent to that seen in Europe.

There are a number of key implications from our work. We have clearly shown that improving connectivity in Africa is only one part of the equation; it is also necessary to ensure that services are appropriately provisioned. Thus, CPs should begin to improve their presence there. Intuitively, popular regional providers should be the front-runners in this effort. Although perhaps not immediately financially beneficial, this could act as a powerful catalyst for Internet uptake, which will result in revenues in the future. Combining the above, we can therefore propose some steps that should be taken by both network operators and web providers: (*i*) operators must improve peering between networks to enable cache capacity to be shared cheaply and with low delay; (*ii*) content providers must concurrently be encouraged to host caches at existing IXPs; (*iii*) network operators must correct their DNS configuration settings to rely on local DNS for resolution; and (*iv*) public DNS resolvers should be placed in Africa (*e.g.*, at some of the 38 African IXPs as of September 2017 [215, 216, 292]) to reduce the overheads for clients that continue to use them. These steps are complementary, with the ability of all stakeholders to encourage each other. For instance, if Google were to redirect more clients to GGCs hosted in Africa, network operators would be encouraged to increase peering to reduce the cost of these redirections.

Chapter 5

Topology and Infrastructure: A Look Towards the Future

In this chapter, we begin by identifying the interconnection challenges in the region, before looking towards the future of the African Internet while learning from its past, most notably the results of our previous longitudinal studies. We then present an option for enriching connectivity and incentivizing Content Providers (CPs) to establish presence in the region: an innovative interconnection framework to build a distributed Internet eXchange Point (IXP) layout spanning the continent and nearby islands.

5.1. Interconnection challenges in Africa and lessons learned from our previous studies

5.1.1. Interconnection challenges in the African region

Reasons for low penetration and low quality of Internet access in Africa are numerous: high Internet access costs inherent to energy instability, transit costs, network operation costs, lack of infrastructure in rural areas, lack of content hosted in the region, as well as the preference of end-users for popular Google, Facebook, or Youtube content mostly served from Europe and the US. Some of them have been thoroughly analyzed in the previous chapters. These lead to a constant loop (no local content no peering; no peering no local content). In such a context, I identified in [61] the key milestones for a better Internet access in Africa as follows:

1. A better energy provision to the industry: since power is essential for industry and therefore for Internet access, local governments or private companies need to make its provision stable and sustainable. Energy provision could be boosted by competition in this sector as well as an orientation of electrification politics from short to long-term towards the storage and provisioning of solar energy, renewable energy, gas, and even nuclear energy.

2. A climate of fairness and cooperation/partnership established by the regulations in the telecommunications market to secure and pave the way for massive investments.

3. The adoption of traffic engineering techniques and efficient routing by Internet Service Providers (ISPs), aiming at keeping local traffic local. Apart from those measures, increasing peering and adding services (DNS root servers, CCTLDs, looking glass, search engines, Internet portals, etc.) to local IXPs will significantly contribute to making ISPs save on transit costs. They could then use these saved costs or interests to invest in building the physical infrastructure.

4. The creation of local content and the stimulation of content hosting to boost local economies: it is worth mentioning that content developed in each country need to be attractive enough and need to have potential to be exported (*i.e.*, knowledge, culture, music, videos, activities specific to the country but well appreciated elsewhere) at least to other countries in its sub-region.

I also specified that although considerable efforts are being made on the continent to achieve these objectives, they need to be multiplied. More specifically, I listed the followings as essential conditions for achieving a better Internet in the region:

1. Affordable (cheap) international connectivity.

2. Cross-borders interconnections and regional transit networks.

3. Investments in terrestrial optical fiber within sub-regions, countries, and cities or investments in alternative technologies suitable for reaching more end-users.

4. Content produced by end-users (*e.g.*, students, universities, local government services, companies, etc.), hosted locally and highly available online.

5. Data-centers connected to local IXPs to host servers or government services (content produced locally, CPs caches, etc.).

One way of improving the African Internet is the bottom-up model. It consists for universities and National Research and Education Networks (NRENs) (*e.g.*, Tunisia NREN (TUREN), Moroccan NREN (MARWAN) [305], West and Central African Research and Education Network (WACREN) [299], Ubuntunet Alliance [19] – the alliance of NRENs of Eastern and Southern Africa, and so one) to build per country an academic network linking schools, universities, research centers, hospitals, etc. The existence of such networks would incentivize ISPs and governments to invest in cross borders connections and Internet provision. This option has been thoroughly inspected, a decade ago, by Pehrson *et al.* in their study [221], which targeted local educational and research institutions and suggested their interconnection at the regional level. In particular, [221] also gave some proposals for how to integrate pieces of the existing terrestrial

fiber being deployed in ongoing development programmes, proposing a regional infrastructure, which includes fiber deployed on and between nearby campuses, leased fiber, and capacity purchased at the wholesale level. Nevertheless, this interconnection is yet to be realized.

Another possible way is the top-down model which would need the regulations to facilitate fiber deployment, to ensure that all ISPs (private and incumbent) have the same rights on telecommunications market, as well as to enable both competition and partnership. In addition, this model would also require regulations to encourage infrastructure sharing for the welfare of the end-users, make "crossing borders" easy for the ISPs, make declarative the licensing procedure for ISPs or hosting companies, etc. Creating a distributed IXP layout spanning the continent may benefit from this environment and help fulfill the aforementioned critical points for enriching connectivity in the region. Along these lines, we investigate, in Section 5.2, whether such an IXP interconnection would be possible, and we estimate, if successful, the best-case benefits that could be realized regarding traffic localization and performance. The proposed interconnection framework, which aims at enabling ISPs present at isolated IXPs to interconnect and incentivizing CPs to establish a presence in the region, arises from the lessons learned from our previous studies.

5.1.2. Lessons learned from our previous studies

Technically speaking, the African Internet ecosystem is experiencing classic "growing pains": a few 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 IXPs as of March 2016,[1] period during which this work has been launched [90, 292]. However, only 29 of the 58 countries in the region (including nearby islands such as Sao Tome and Principe (ST), Mayotte (MY), 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,[2] and we have previously demonstrated the benefits that new IXPs can bring (Section 3.2.1) some local ISPs are still hesitant to peer at those IXPs as shown in Section 3.3.1. Adding to the difficulties, terrestrial fiber deployment remains fragmented [198, 264, 265], since fewer technical and political hurdles make submarine fiber cheaper to build than inland fiber [27, 277–279].

A major reason behind the stunted growth of the African Internet ecosystem is that the region suffers from a lack of local content (*cf.* Section 4 and [157, 270]). Content is mostly served from the US and Europe (EU), and even the most popular regional websites are hosted abroad, as investigated in Section 4. 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 and publish measurement trends for opportunities to improve Internet services at an affordable cost [61, 115, 270].

In Figure 5.1, we summarize the lessons learned from our previous studies and obtained re-

[1] Only one more IXP has been set up and is active as of September 2017, leading to a total of 38 IXPs in the region [292].

[2] 18 new IXPs were established in Africa from July 2014 to July 2015 [6, 217]

sults, which may represent the key features of a possible solution to reshape the African Internet. We will refer to these throughout the next section.

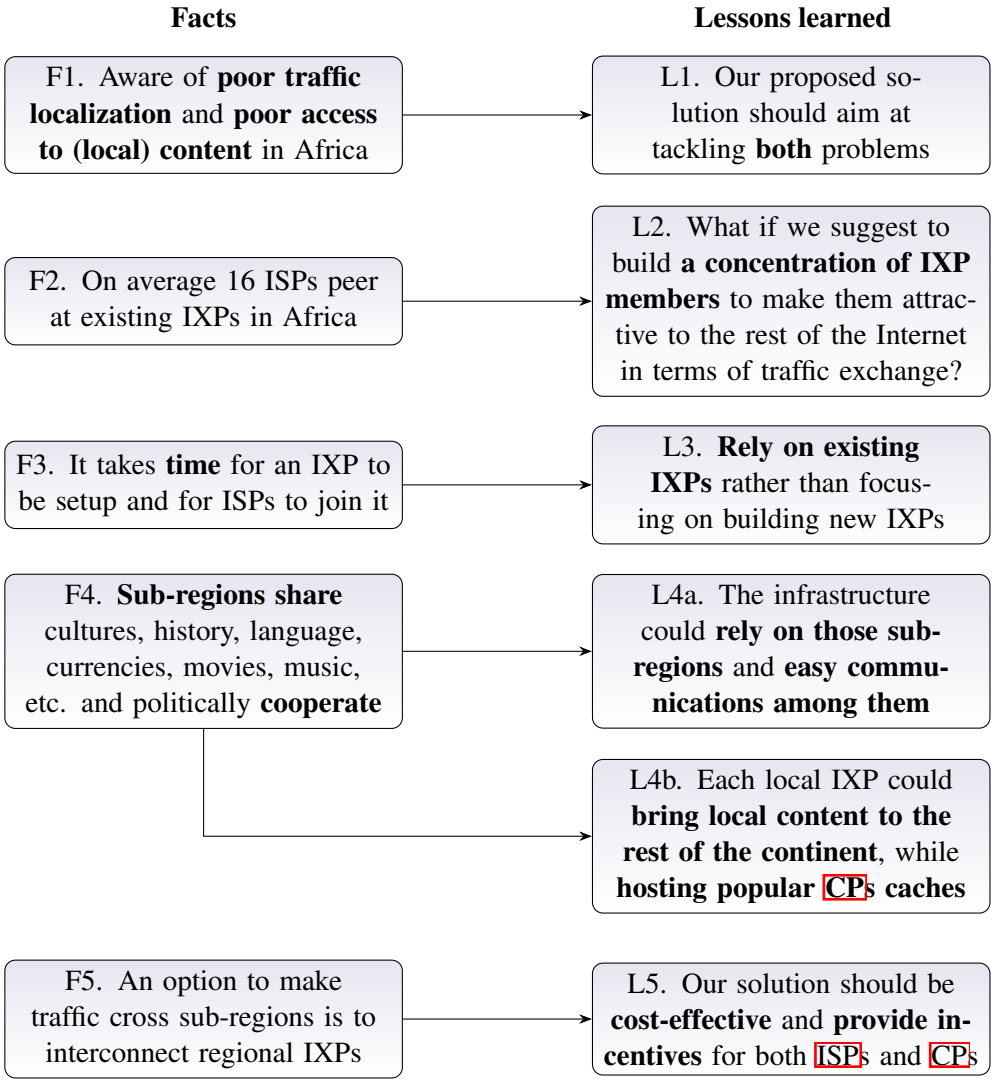

Figure 5.1: Identifying from our previous studies the key features of a solution to enrich connectivity in Africa

5.2. Reshaping the African Internet: from scattered islands to a connected continent

An option that could be considered to enrich connectivity on the African continent and incentivize CPs to establish a presence in the region is to interconnect ISPs present at isolated IXPs by creating a distributed IXP layout spanning the continent (*cf.* L2 and L5 in Figure 5.1). We are not the first to think about IXP interconnection as a way to achieve these goals [71, 72, 203, 211, 268, 275]. However, what is lacking is a concrete proposal for achiev-

ing IXP interconnection and a quantitative estimation of potential benefits from doing so. In this work, our main goals 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 5.2a shows an overview of the methodology we have adopted to create the distributed IXP layout and quantitatively estimate its benefits. Similarly to our studies [78, 81, 85, 89] presented in Chapters 3 and 4, 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 study, we have conducted a survey of the 37 African IXP operators to get their opinions on the feasibility of IXP interconnections that we report in Section 5.2.1.

We have then explored two simplistic approaches to solving this problem as a reference point for the rest of this section. 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 has 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 infrastructure (Section 5.2.2). We have, therefore, developed and evaluated a framework, which considers and parameterizes all these external factors using publicly available datasets (Section 5.2.4). Further, we have 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 (*cf.* L1 in Figure 5.1).

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 CPs could deploy caches (Section 5.2.3), as shown in Figure 5.2b. Our approach is novel in the following respects: *(i)* we make design choices that ensure that the solution is realizable right away (Section 5.2.3), *(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 (Section 5.2.4), *(iii)* we focus on a solution that requires the minimum possible investment in infrastructure (Section 5.2.6), another key feature identified in Figure 5.1 (L5), and *(iv)* we suggest three options applicable within/across sub-regions, given the interests of the stakeholders, to realize the interconnection scheme (Section 5.2.12.2.2).

We use extensive simulations with the open source BGP routing solver C-BGP [232–234, 260, 308] to evaluate the proposed solution and to quantitatively demonstrate the benefits that would be realized at each step (Section 5.2.6). 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 Alexa websites would decrease by more than their third.

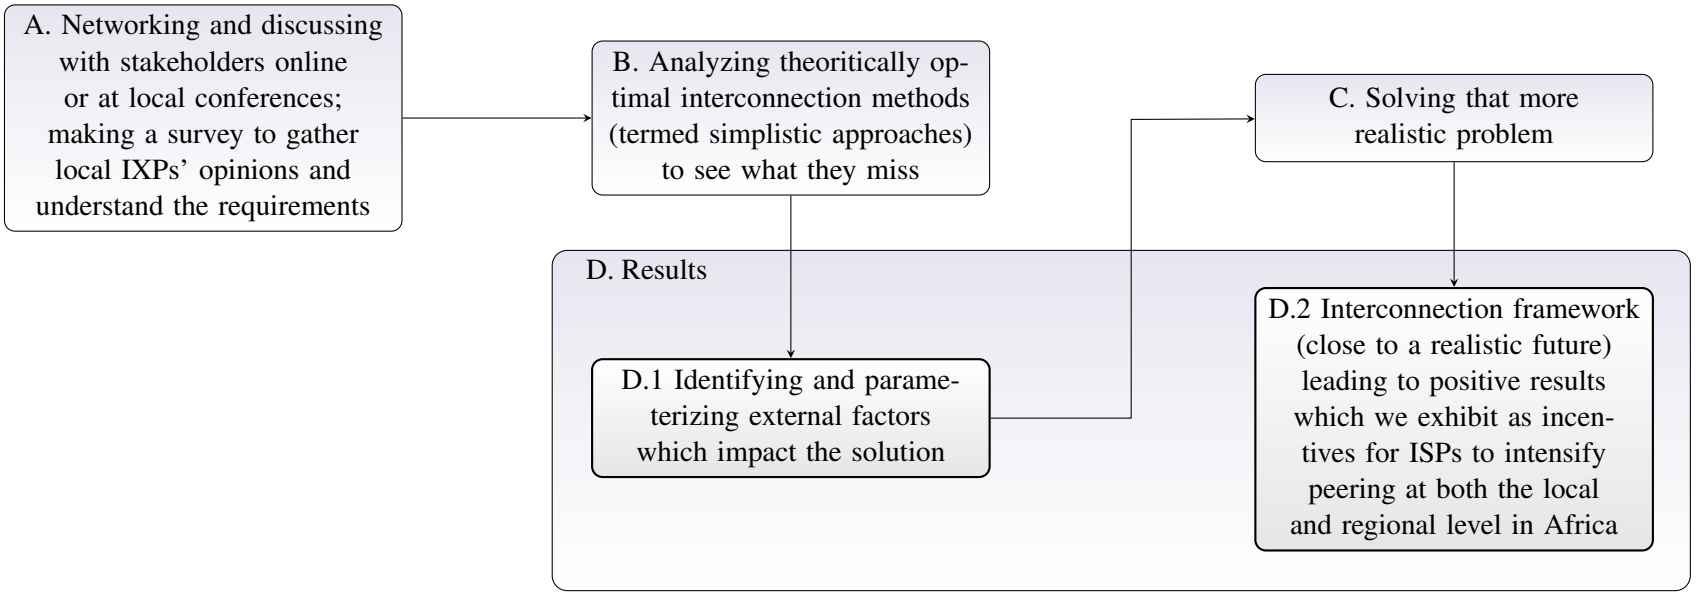

(a) Block diagram of the methodology followed in this work. Our approach to solve the problem (at step C) is detailed in Figure 5.2b

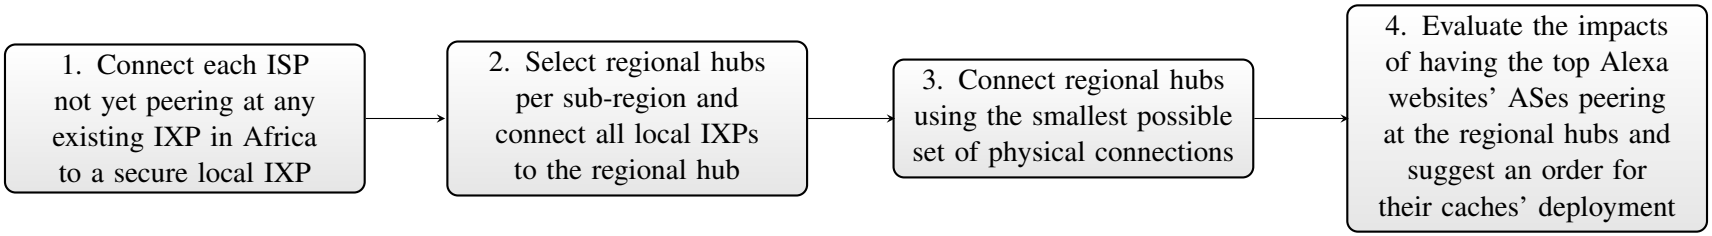

(b) Overview of our 4 steps of the proposed interconnection approach to build the distributed IXP layout. It details step C of Figure 5.2a

Figure 5.2: Block diagrams of the methodology followed in this work and our proposed approach to build the distributed IXP layout.

We hope such results will encourage local operators to increase peering and CPs to establish a presence in the region.

The scientific contributions of this work are four-fold: 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 a hierarchically organized IXP substrate. Further, we demonstrate the quantitative benefits of the framework regarding 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.[3] 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 paths going from African ASes to those hosting top global and regional content.

The remainder of this section is structured as follows. In Section 5.2.1, we perform a broad analysis of the region that consists of related work and the results of our survey of local IXP operators. In Section 5.2.2, we inspect simplistic approaches to the distributed IXP problem and briefly expose the reasons why they would not be feasible in practice. Next, we present in Section 5.2.3 an overview of our solution, a first attempt to interconnect existing IXPs in Africa. In Section 5.2.4, we present an overview of the data collection, the curation methodology, and the parameterization of the model. We then flesh out, in Section 5.2.6, each step of our approach and evaluate the benefits as compared to the initial AS topology. After that, we explore the sensitivity of our framework to variations in parameterization in Section 5.2.11, before discussing in Section 5.2.12 the limitations of our approach and its feasibility from a technical and political perspective.

5.2.1. Broad analysis of the region

5.2.1.1. Background of the region

As explained in Section 1.1.2, the 54 African countries can be classified into distinct *sub-regions* (North, West, East, Central, or Southern) as per the African Union [9, 306, 309]. The concept of African sub-regions is important while planning infrastructure in the region (*cf.* L4a

[3] C-BGP [232–234, 260, 308] is an open source routing solver that eases the investigation of changes in the routing or in the topology of large networks.

and L4b in Figure 5.1). Since countries within a sub-region already agree and cooperate[4] on various issues, this co-operation could be leveraged.

5.2.1.2. Survey of African IXP operators

To understand the viewpoint of African IXP operators about IXP interconnection, we surveyed the 37 local IXP operators as of early 2016 (Appendix B), 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, we propose a solution to this issue in Section 5.2.7.1.1. 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 two 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 tackle the first parameter by proposing the three first steps of our framework, while to the second one is dedicated its fourth step. A second IXP in *East Africa (EAf)* 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 (CAf)* replied that it would be interested in interconnecting to other IXPs to improve the interconnection options of its customers, 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 *EAf* described IXP interconnection as a controversial topic, since 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.

[4] The 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. [9]

5.2.2. Simplistic 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 [70], we use the Google maps Distance Matrix API [109, 112] 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 apply 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 [177, 198] to determine which physical links can be used to establish the spanning tree. Figure 5.3 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 (CI), Egypt (EG), Burkina Faso (BF)), rebellions (*e.g.*, DR Congo (CD), Nigeria (NG), Burundi (BI)), or terrorists attacks (*e.g.*, Sudan (SD), NG) over the last five years [46, 97, 293, 302–304, 309]. 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 (CG) and Chad (TD) in April 2016 [239, 240], or in the English-speaking areas of Cameroon (CM) from January to April 2017 [10, 56].

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 do not currently exist due to various economic and political reasons. As an example, CD and CG do not agree to let any fiber cross their shared border; due to regulatory disagreements, optical fiber deployed five years ago through the Congo river to interconnect both countries has still not been switched on [12]. 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 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 three closest IXPs in different sub-regions, and manually overlay the interconnection scheme with fiber maps [177, 198]. We find that the result is quite similar to Figure 5.3, with the main difference being that *CAf* now plays the role of a hub. Still, many of the links within sub-regions cannot be established. In *CAf*, we end up with not only the problematic physical link between CD and CG but also a terrestrial fiber between Kinshasa (CD)

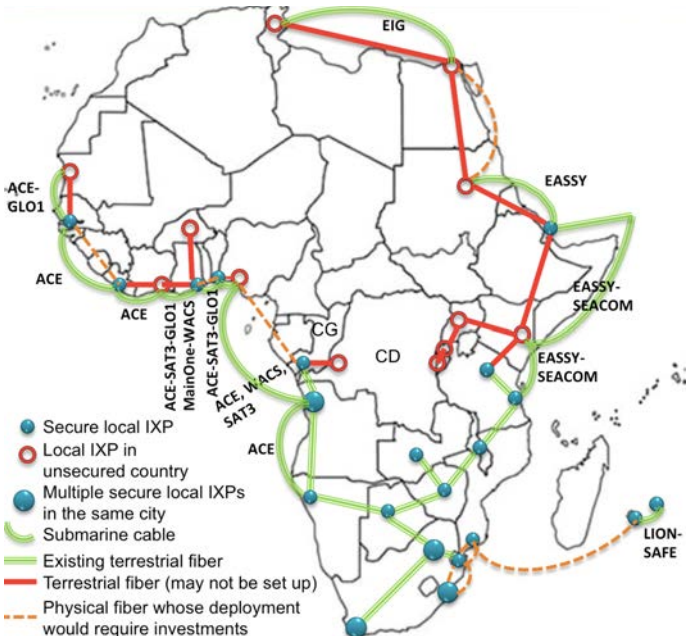

Figure 5.3: Interconnecting IXPs in Africa along the minimum spanning tree would be infeasible due to "unsecured" IXPs and the difficulty of fiber deployments along some links.

and Bujumbura (BI). Pehrson *et al.* [221] showed, a decade ago, that the best way to connect the East to the West of Africa is to cross CD 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 CD and at its border with Rwanda (RW) [46, 97, 309].

In summary, we have 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 solutions that account for socio-political and economic factors, which influence topology design in the region.

5.2.3. 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 Section 5.2.12.2.

Step-1: Connect each ISP not yet peering at any existing IXP in Africa to its closest 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[5] secure and easily accessible country. Note that 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 (*cf.* L3. in Figure 5.1). 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 Section 5.2.7.

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 [48, 64, 100]. 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 can peer with 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 Section 5.2.8) to help avoid 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 (*cf.* L5 in Figure 5.1). We present the details of how to interconnect regional hubs in Section 5.2.9.

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 content with low latency and the best possible performance. In step-4 of our proposed solution (Section 5.2.10), 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 providers CPs caches.

5.2.4. 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] Recall that we have deduced from our results in Section 3.2.1.6 that it is often better in terms of QoS for an ISP operating in Africa to peer at its closest IXP in Africa than at an IXP located on another continent.

5.2.5. Parameterizing geo-political and socio-economical contexts

Matrix depicting the geography of the African continent: We define M_{geo} as an $N \times N$ 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 an $N \times 1$ matrix to quantify the proportion of the IXPs located in Africa, which are hosted in an African 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}, an $N \times N$ matrix to denote whether one or more submarine cable systems, which belong to the same operator, land in a pair of countries. For example, the entry corresponding to the countries (Ghana, Ivory Coast) is 5, because five 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 [279]. 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 Section 5.2.7.1.1.

Matrix of terrestrial fiber deployment within or between countries: We define M_{tfib}, an $N \times N$ matrix that captures the presence of terrestrial fiber within or between countries. Specifically, since terrestrial fiber is often deployed along roads [70], 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 [109], following the road along which fiber is deployed [198, 265].

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 [198, 265] 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 [109, 198, 265]. 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 Section 5.2.7.1.1.

Matrix of African security or political realities: We define M_{pol}, an $N \times 1$ matrix that

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

$Pb = \{\}$ /* Initialization of Pb */
/* Label as best any unique country path */
for $c \in$ *P.keys()* **do**
  if *len(P[c]) = 1* **then** $Pb[c] = P[c]$
/* Identify the best path for the rest */
$current\_country\_path\_len = 2$
while $current\_country\_path\_len <$ *58* **do**
  for $c \in$ *P.keys()* $\mid c \notin$ *Pb.keys()* **do**
    /* Can we use submarine cables ? */
    $A_s = \{\}$ /* Sum # of common types of submarine cables per path */
    $C = \{\}$ /* Percentage of African IXPs in destination country */
    for $p \in$ *P[c]* **do**
      $i = 0$
      while *i* $<$ *len(p)-1* **do**
        $A_s[p]$+= $M_{sfib}[p[i], p[i+1]]$
        i += 1
      $C[p]$ += $M_{ixp}[p[i]]$
    if $\sqsupset p \mid A_s$*[p] = arg max* A_s*(x) and C[p] = arg max C(x)* **then** $Pb[c] = p$
    else if $\sqsupset p \mid A_s$*[p] = arg max* A_s*(x)* **then** $Pb[c] = p$
    /* What about terrestrial fiber ? */
    $A_t = \{\}$ /* Ratios of terrestrial cables deployment per path */
    $B_t = \{\}$ /* Investments in the countries on each path */
    $C = \{\}$ /* Percentage of African IXPs in destination country */
    for $p \in$ *P[c]* **do**
      $i = 0$
      while *i* $<$ *len(p)-1* **do**
        $A_t[p]$+= $M_{tfib}[p[i], p[i+1]]$
        $B_t[p]$+= $M_{se}[p[i]]$
        i += 1
      $B_t[p]$+= $M_{se}[p[i]]$
      $C[p]$ += $M_{ixp}[p[i]]$
    if $\sqsupset p \mid A_t$*[p] = arg max* A_t*(x) and* B_t*[p] = arg max* B_t*(x) and C[p] = arg max C(x)* **then** $Pb[c] = p$
    else if $\sqsupset p \mid A_t$*[p] = arg max* A_t*(x) and* B_t*[p] = arg max* B_t*(x)* **then** $Pb[c] = p$
    else if $\sqsupset p \mid A_t$*[p] = arg max* A_t*(x) and C[p] = arg max C(x)* **then** $Pb[c] = p$
  $current\_country\_path\_len$ += 1

identifies countries that have experienced political issues, insecurity (wars, terrorist attacks, riots, rebellions), and disputes with their neighbors [46,97,293,302–304,309] during the last five 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 Section 5.2.7.

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}, an $N \times 1$ 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 five years by each country in the telecommunications, transport, and energy sectors, respectively [294]. 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 Section 5.2.7.1.1. Further, we use R_{se} values in the five-year threshold sensitivity analysis (Section 5.2.11).

5.2.5.1. Collecting the Internet AS-level topology

AS relationship dataset: We used the CAIDA AS-level topology snapshot from March 2016 [44], 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 [42].

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 retained the Ark traceroutes data for the first two weeks of March 2016 [45]. 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 [85] composed, among others, of full mesh paris-traceroutes [24] 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 [18]: we added paris-traceroutes, previously collected in [89], 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 term 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 [164,220], PCH [214], TeleGeography Internet Exchange Map [227], and African IXP websites. We then asked local IXP operators to validate (Section 5.2.1) this dataset (from January to March 2016) for completeness, before using it in Section 5.2.6.

5.2.5.2. 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 RIRs [11,20,23,160,252]. For each IPv4 address block, we geolocated the IPs in the block using the Netacuity Edge database [66]. We are well aware of the limitations of existing geolocation databases [124, 225]; 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 [286] 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, 9,965 in North America, 2,281 in Asia, 3,351 ASes in South America, and 773 in Australia. We further classified ASes operating in Africa into the five sub-regions: 199 ASes in *WAf*, 296 in *SAf*, 66 in *CAf*, 83 in *NAf*, and 232 in *EAf*. In this section, we denote ASes that operate predominantly in the region as *African ASes*, while those operating predominantly outside the region are denoted *non-African ASes*.

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

Collecting data that shed 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 efforts can be alleviated. Our results in the subsequent sections (Sections 5.2.6 and 5.2.11) demonstrate, however, that the manual effort we invest here can have a large payoff regarding 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 IXP operators to validate and complete if necessary the inferred list of their IXP members obtained from publicly accessible datasets such as PeeringDB [164, 220] or PCH [214]. 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 accessible in the

technical report [90].

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.

5.2.6. 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 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 5.1 shows an overview of the metrics used to characterize the initial topology and the result of each step. Its 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 four steps. Figure 5.4 shows the distribution of estimated path RTTs for the initial topology and after each step: 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. We will refer to both Table 5.1 and Figure 5.4 throughout the remainder of the section.

5.2.6.1. Building the initial AS topology

5.2.6.1.1. 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 [234]. Simulating the entire AS-level Internet would be computationally inefficient with the resources we have available and is not necessary for our study. We implemented the following procedure to scale down the topology to a size suitable for simulation, without changing the possible outcome.

We start from every African AS (as defined in Section 5.2.5.2) and traverse *customer-to-provider* (*c2p*) links until we reach the clique of tier-1 providers [42]. 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 1,389 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 in Section 4.3.3 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 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

Table 5.1: Overview of topology characterization from each step of the proposed framework.

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 %	93 %
	% of AS paths with length of 2	1.5 %	9.2 %	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)	144.1 ms	58.9 ms	61.75 ms	61.1 ms	75.2 ms
	Interquartile range (Quartile 3 – Quartile 1)	162.1 ms	147.3 ms	115.9 ms	63.2 ms	32.1 ms
Paths from African ASes to non-African ASes	% of AS paths with length ≤ 4	50.8 %	53.9 %	53.9 %	54.4 %	54.2 %
	% of AS paths with length of 2	0.7 %	1.5 %	1.5 %	2.1 %	2.1 %
	Mode	4	4	4	4	4
Paths from African ASes to African ASes hosting popular content	% of AS paths with length ≤ 4	61.2 %	74.7 %	86.3 %	91.7 %	91.7 %
	% 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	71.1 %	70.2 %	73.2 %	73.8 %	74.3 %
	% of AS paths with length of 2	2.6 %	2.9 %	3.8 %	4.8 %	6.6 %
	Mode	4	4	4	4	3-4
	% of AS paths with mean RTT ≤ 100 ms	30.6 %	36.4 %	37.4 %	64.6 %	65.7 %
	% of AS paths with maximum RTT ≤ 1000 ms	22.82 %	22.81 %	23.03 %	60.83 %	87.5 %
	Median of mean RTTs (Quartile 2)	137.3 ms	137.5 ms	137.5 ms	82.5 ms	82.5 ms
	Interquartile range (Quartile 3 – Quartile 1)	162.1 ms	150.2 ms	148.7 ms	103.1 ms	103.1 ms
Sensitivity analysis (% best country paths affected by the change of the "insecurity" threshold)	Last year	4.4 %	3.7 %	6.9 %	0 %	N/A
	Last 3 years	1.8 %	3.7 %	6.9 %	0 %	N/A
	Last 10 years	4.4 %	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	US$73.4 million	US$439,849	US$0	US$0
	Higher boundary of total costs needed	N/A	US$1.8 billion	US$11 million	US$0	US$0

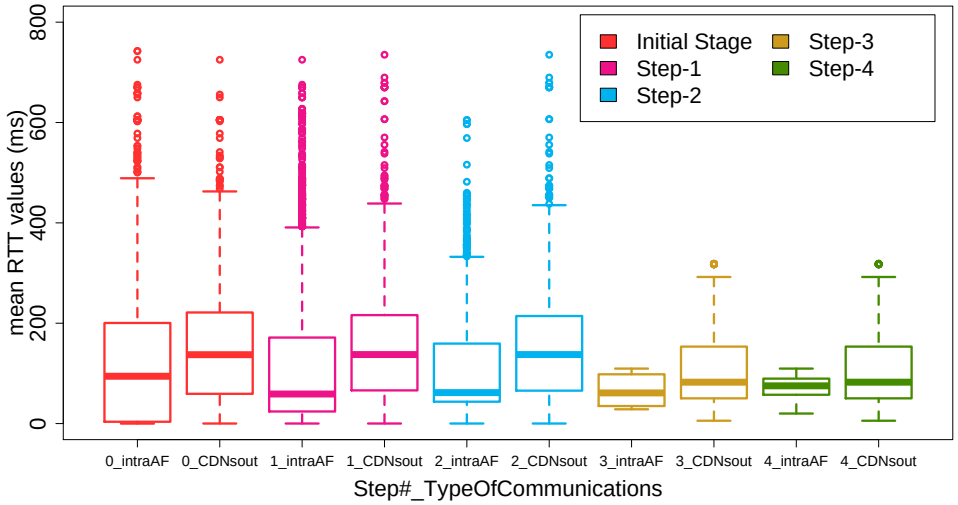

Figure 5.4: Boxplot of the estimated mean RTT distribution on AS paths at each step, depending on the type of path.

contains 1,725 prefixes.

After that, we obtain the preferred path from each AS in the retained topology to prefixes in $\mathcal{P}$, by simulating in C-BGP the whole AS graph, which consists of 53,537 ASes, 215,628 AS links from the *CAIDA AS relationship dataset* [44], 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 acknowledge this aspect, which could impact the results of this work and which is a matter of further study. However, preliminary analysis and geolocation of ASes at the router level for Africa has revealed that most of the African ASes have a national scope, so that, de facto, the impact of this simplification is likely less important for the outcome of this analysis. Finally, we represent an IXP by the set of peers and the peering links found between ASes according to our data as described in Section 5.2.5.1.

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 AS paths starting from AS4558 (known to host the KIXP Routeviews collector).

The BGP data from the JINX and KIXP Routeviews collectors for the first three 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 [90].

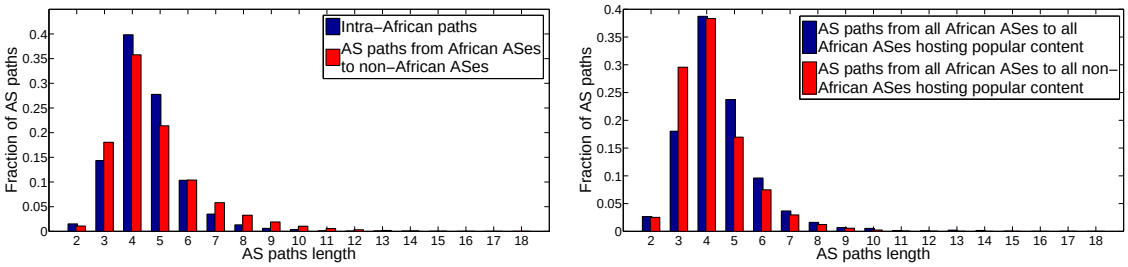

(a) *Distribution of AS path lengths for intra-African paths and for paths between African ASes and non-African ASes.*

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

Figure 5.5: 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.

5.2.6.1.3. Characterizing the initial topology We define an *intra-African path* as an AS path which originates and terminates at *African ASes*.[6] An *outside-African path* is a path from an *African AS* to a *non-African AS* (*cf.* Section 5.2.5.2). *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 (NAm), while 12 % traverse both EU and North American ASes. Figure 5.5 shows the path length distribution for both intra-African AS paths and outside-African AS paths. We find that the mode of path lengths is 4 AS hops 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 5.6 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 the CDFs for both types of paths to have similar properties; 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).

5.2.7. 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

[6] African ASes are those that predominantly operate in Africa, as defined in Section 5.2.5.2

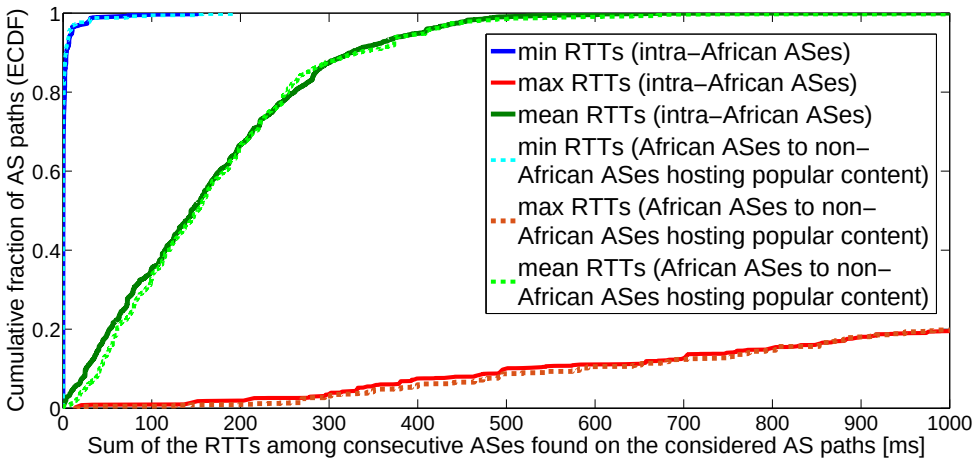

Figure 5.6: In the initial topology, CDF of the mean, minimum and maximum RTT estimates on intra-African AS paths and paths from African ASes to non-African ASes hosting popular content.

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.

5.2.7.1. Identifying secure IXPs and secure relationships between countries

We use the M_{ixp} and M_{pol} matrices (Section 5.2.5) to construct the matrix $\bar{M}_{ixp}$ representing secure local IXPs. For any country A, if $M_{pol}[A]$ is 1 (labeled *not secure*), then we set $\bar{M}_{ixp}[A]$ to 0. Table 5.2 provides details about the 25 secure local IXPs in 18 secure countries covering four African sub-regions as of March 2016: *North Africa* (*NAf*) does not have any secure IXP. The numbers of members ASes of Table 5.2 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: their number of members (in italics) were, therefore, collected from public datasets other than the IXP websites, where possible, or were estimated to the total number of ASes operating in the IXP host country. We next use M_{geo} and M_{pol} (Section 5.2.5) 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*.

5.2.7.1.1. 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

Table 5.2: List of (the 25) secure local IXPs in Africa as of March 2016 (with their number of members), classified by sub-region and country.

African sub-region	Country	#IXPs	#Members ASes
East Africa (*EAf*)	Djibouti	1	**5**
	Mauritius	1	**12**
	Reunion	1	16
	Tanzania	2	**33** - **6**
Central Africa (*CAf*)	Congo	1	**8**
Southern Africa (*SAf*)	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 (*WAf*)	Benin	1	**5**
	Gambia	1	**14**
	Ghana	1	**17**
	Liberia	1	*5*
Total	18 countries	25 IXPs	

"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 [27, 177, 198, 265, 267, 279].

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. 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 the best country path. These two cases accounted for 25 source countries. For each of the remaining 33 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: $A_s = \sum_{c \in \mathcal{C}} (M_{sfib}[c]/|\mathcal{C}|)$ and $C = M_{ixp}[d]$, where $\mathcal{C}$ 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 (TG) – Ghana (GH) via GLO1 or WACS submarine cables, to the path TG – Benin (BJ) via only GLO1. We also prefer the path CD – Angola (AO) via WACS and ACE cables and toward two IXPs, to the path CD - CG 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 path.[7] These are, for instance, the cases of Rwanda (RW) – Tanzania (TZ), Uganda (UG) – TZ 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 path.[8] Otherwise, if to a path correspond the maximum values for A_t and C, we select that path.[9] As an example, the country path Burkina Faso (BF) - GH is preferred to BF - BJ, because A_t is higher for the former and both BJ and GH 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 labeled unsecured (Libya, Egypt (EG), 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 EG connect to Djibouti (DJ) via EASSY and SEACOM. At the end of this step, all countries are assigned a best path, as depicted in Figure 5.7.

[7] Preference for country paths with considerable terrestrial fiber deployment, larger investments in telecoms, and more diversity in IXPs at the destination

[8] Preference for country paths with considerable terrestrial fiber deployment and characterized by larger investments in telecoms

[9] Preference for country paths with terrestrial fiber deployment and more diversity in IXPs at the destination

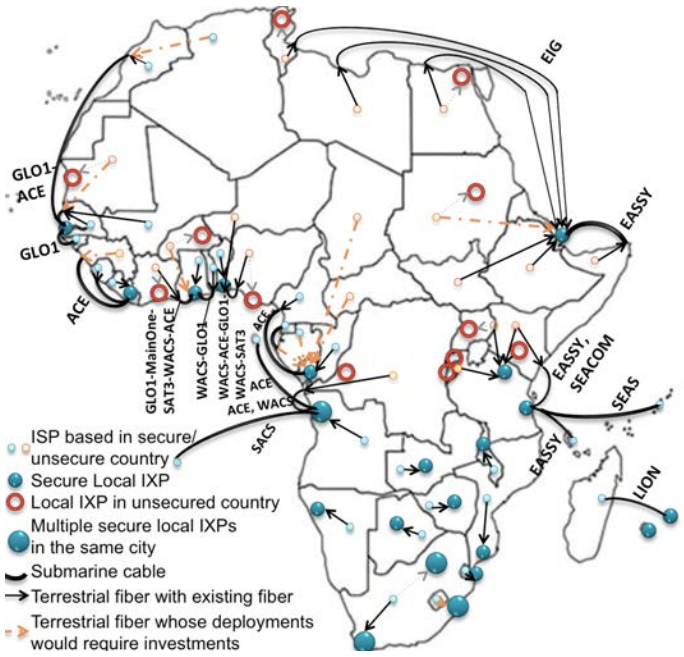

Figure 5.7: Result of step-1, where each ISP connects to its closest secure IXP.

5.2.7.1.2. 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 [65, 204], as it has no incentive to peer with networks it can reach via customer links. This consideration is consistent with our goal to evaluate the best possible scenario that can be realized and its impacts on AS path lengths and performance. In Section 5.2.12.2, 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 [44]. 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 [109]. 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 [226, 235], the $RTT(s,d)$ over the established link can be estimated as: $RTT(s,d) = \frac{2*C_h(s,d)}{2/3\mathbf{c}} = \frac{3*C_h(s,d)}{c}$, where $C_h(s,d)$ is the distance (km) between the cities following roads/railways [109], 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.

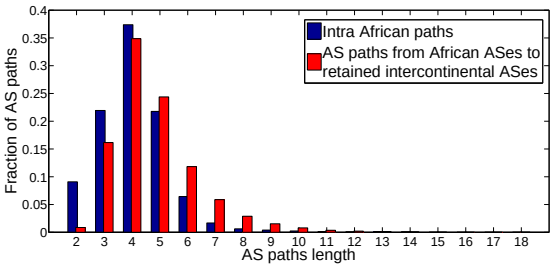

(a) *Distribution of AS path lengths for intra-African paths and for paths between African ASes and non-African ASes.*

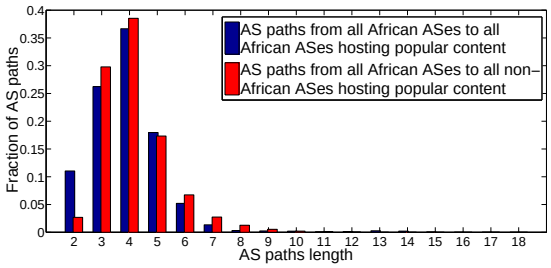

(b) *Distribution of AS path lengths for paths between African ASes and African/non-African ASes hosting popular content.*

Figure 5.8: 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.

5.2.7.1.3. Characterizing the resulting topology To simulate the effect of step-1, we propagate the 1,725 prefixes in C-BGP on a topology where all the new peering links have been established as described in Section 5.2.7.1.2. 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 9.7 % traverse ASes operating predominantly in both regions.

Figure 5.8a shows that the mode of intra-African AS path lengths is still 4. The fraction of such paths having a length of 4 or fewer increases from 56.9 % to 69.9 %, however. Similarly, Figure 5.9 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 (IQR) of 162.1 ms) in the initial topology to 58.9 ms (with an IQR of 147.3 ms) after step-1, as shown in Figures 5.4 and 5.9.

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 IQR 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.

5.2.8. 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

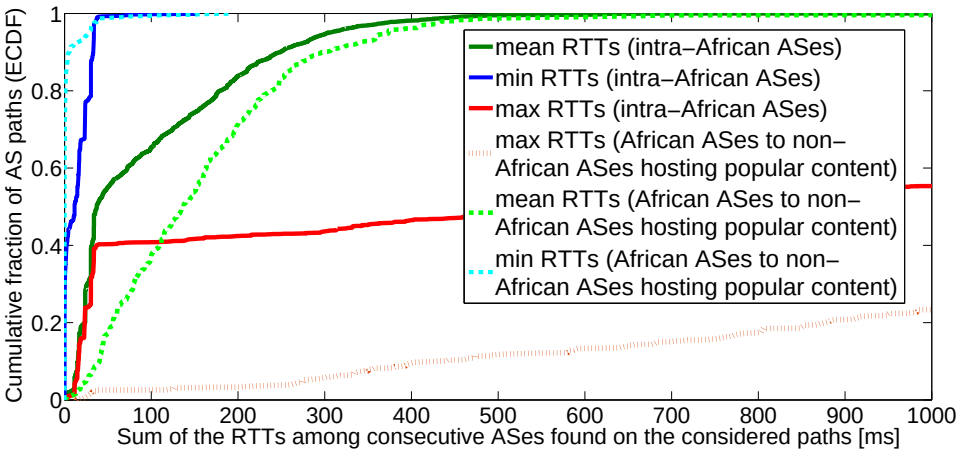

Figure 5.9: After step-1, CDF of the mean, minimum, and maximum RTT estimates on intra-African AS paths and paths from African ASes to non-African ASes hosting popular content.

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 5.2, the regional IXP hubs are TIX in TZ (33 members) for *EAf*, Congo Internet eXchange (CGIX) in CG (8) for *CAf*, NAPAfrica JB in ZA (141) for *SAf*, and GIXA in GH (17) for *WAf*.

Next, we need to connect each of the 33 remaining local IXPs[10] 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 Section 5.2.7.1.1) 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 three categories. First, among IXPs in *NAf* such as CAIX (EG), RIMIX (Mauritania (MR)), TUNIXP (Tunisia (TN)), and SIXP (Sudan (SD)), 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 (Section 5.2.7). Second, we found KINIX (CD) to have no secure path to its regional hub: again, we use the best country path from step-1. Finally, we connected IXPs located on islands, such as Mauritius-IX (Mauritius (MU)) and Renater-IX (RE) to their closest regional hub (TIX) via submarine cables. The results are shown in Figure 5.10.

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.

[10] 37 African IXPs minus the 4 regional hubs

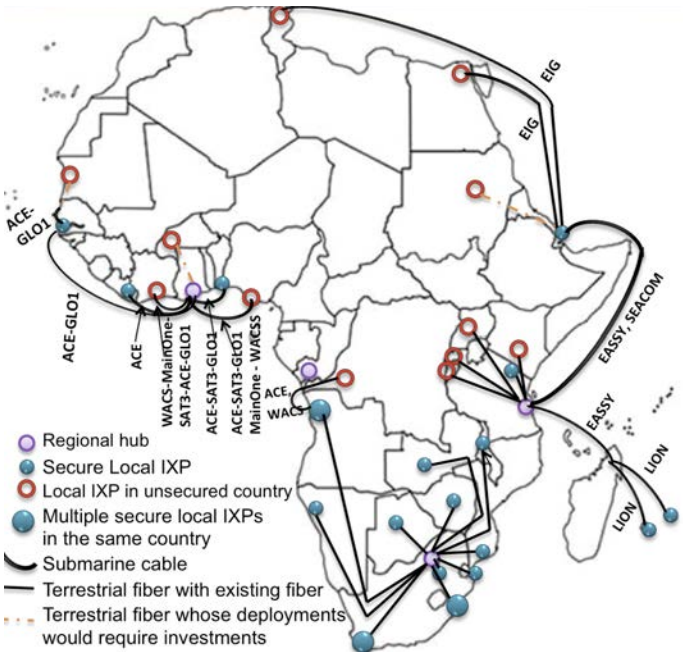

Figure 5.10: Result of step-2, where each IXP connects to the regional hub selected among the secure IXPs of each region.

5.2.8.1. 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 (61.7 ms) than that of step-1 with an IQR 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 5.1). 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 [89], have kept the same distribution as the initial stage.

5.2.9. Step-3: Interconnecting regional IXP hubs

After step-2, we are left with four 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 the best country path as in Section 5.2.7. In this case, however, instead of using the full $\bar{M}_{ixp}$ matrix as input, we use a sub-matrix of $\bar{M}_{ixp}$ 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

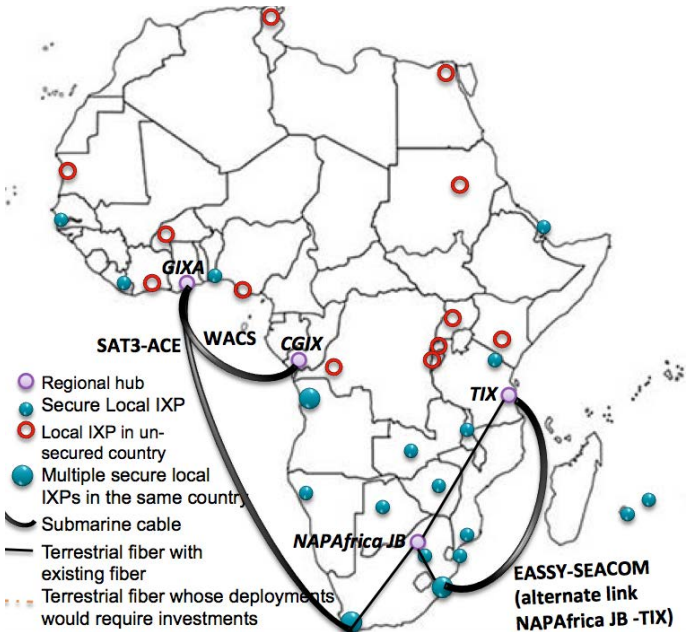

Figure 5.11: Result of step-3, where regional IXPs are interconnected with a minimum number of links.

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 (MZ), 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: CD, 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.

5.2.9.1. 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 5.11 summarizes the results. We find that three links are needed: the link between NAPAfrica JB and 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 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 Section 5.2.12.2.2.

5.2.9.2. 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. Regarding AS path lengths (Figure 5.12), 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 82.5 ms with an IQR of 103.1 ms, as compared to 137.5 ms with an IQR 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 5.1).

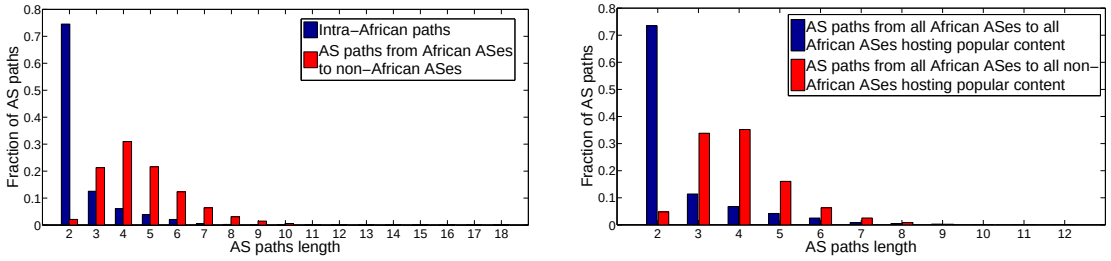

(a) *Intra-African communications and communications between African ASes and non-African ASes.*

(b) *Communications between African ASes and African/non-African ASes hosting popular content.*

Figure 5.12: 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.

5.2.10. Step-4: Incentivizing regional and global CPs 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 (CPs, as defined in Section 1.2.2) to contribute to its realization, we aim at emphasizing in this section what they might gain from participating in it. A typical CP 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 [273]. While CPs can vary in their technical operation (*e.g.*, whether they operate their own backbone network or not), we leverage the fact that all CPs 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 CPs 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[11] Alexa websites [18] mapped in [89] were to peer with networks present at the four 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 IQR (32.1 ms), as shown in Table 5.1. Meanwhile, the median of RTTs from African ISPs to popular content hosted outside Africa stays at 82.5 ms (Figure 5.4). These similarities are expected, as adding the presence of CPs 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.

While establishing CPs' 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 [21]; 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 3.2 %. 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, CPs should proceed in the order TIX, GIXA, NAPAfrica JB, and finally CGIX to have the largest impact (Figure 5.13).

[11] Content Providers (CPs) can offer different services from one region to another.

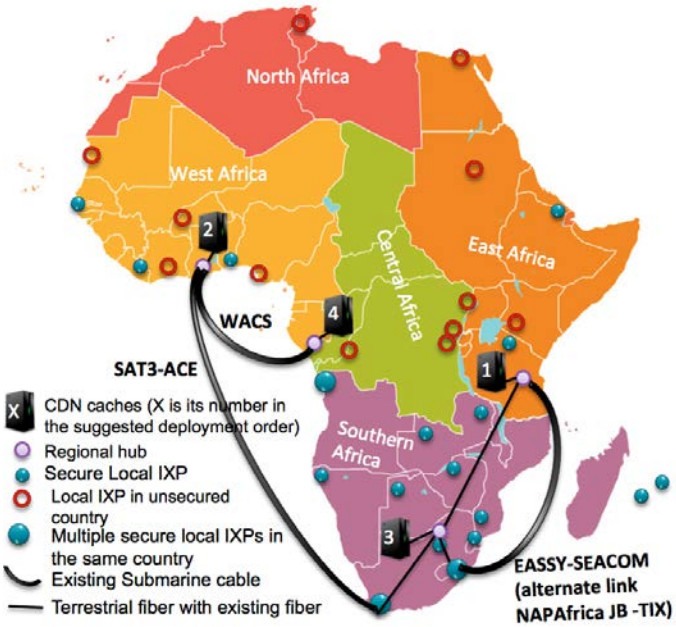

Figure 5.13: Result of step-4, where we suggest an order of CPs' caches deployment within the infrastructures of the strategic points represented by regional IXPs.

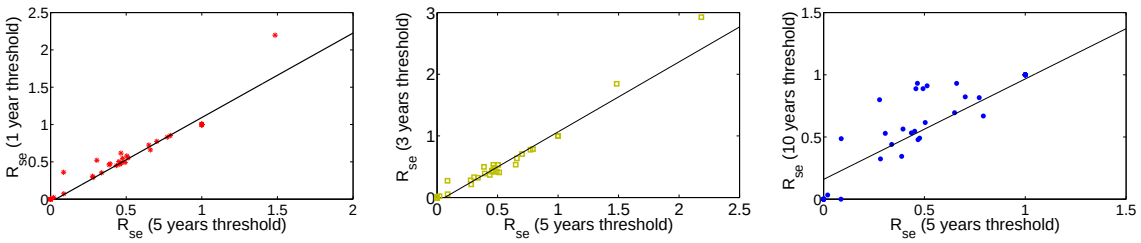

(a) *Correlation of 0.97 for* $(M_{se}(5years), M_{se}(1year))$

(b) *Correlation of 0.98 for* $(M_{se}(5years), M_{se}(3year))$

(c) *Correlation of 0.87 for* $(M_{se}(5years), M_{se}(10year))$

Figure 5.14: Sensitivity analysis: correlation between ratios R_{se} of the matrix M_{se} evaluated for different thresholds is found to be 0.972 for $(M_{se}(5years), M_{se}(1year))$, 0.979 for $(M_{se}(5years), M_{se}(3years))$, and 0.869 for the pair $(M_{se}(5years), M_{se}(10years))$.

5.2.11. Sensitivity analysis

An important consideration that drives the construction of the distributed IXP layout proposed in this section is the notion that a country is labeled "secure" or "unsecured" due to geo-political factors: In Section 5.2.5, we have chosen a period of five 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 five years, it is 20 for the last one year or the last three years, and 27 for the last ten years. Despite this difference, the list of secure local IXPs does not vary. Moreover, the regional hubs (selected at step-2 Section 5.2.8)

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 the row "Sensitivity analysis" of Table 5.1. Note, the values in the column "Initial stage" of that row 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 correspond to the percentage of the selected best country paths affected by that change. In fact, 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} (Section 5.2.5). 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 find the correlation coefficient r to be 0.972 for ($M_{se}(5years)$, $M_{se}(1year)$), 0.979 for ($M_{se}(5years)$, $M_{se}(3years)$), and 0.869 for ($M_{se}(5years)$, $M_{se}(10years)$). Figure 5.14 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 section 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 $M_{se}(10years)$.

5.2.12. Discussions

5.2.12.1. 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 have shown 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 have considered in this work (Section 5.2.5) that could affect the realization of the distributed infrastructure we propose. Nonetheless, this study was a first 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 have modeled each AS as a single router in our simulations and have not considered the internal topology of ASes, since the micro-factors that influence intradomain topology and routing are not the focus of this Chapter. 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 have kept this focus while designing our C-BGP model in order not to deviate from our primary goal. Finally, we have not included traffic data in our model, due to the lack of publicly available datasets about interdomain 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.

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

5.2.12.2.1. Peering economics In designing the distributed IXP layout, we have not at any stage suggested 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 favor of a more market-driven approach. We have 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, it is worth emphasizing that there are certain pre-conditions for our approach to be successful (as detailed in Section 5.1.1): 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 CPs to establish a presence in the region.

5.2.12.2.2. Suggested options for the feasibility of IXP interconnection After discussing with local IXP 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 [38,130]. Similarly, an ISP carrier can also provide transport between two regional hubs. Examples include Liquid Telecom [163], SEACOM [259], and MainOne [179] 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 CP with interests in the region (*e.g.*, Google, Facebook, etc.) could also invest in facilitating the interconnection. In this third category can be classified the efforts of Google for the last mile internet connectivity problem [1, 111, 231].

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 of IXP interconnection and investigating the feasibility of the proposed alternatives involves complex economic analysis, which is out of the scope of this work. 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 set up the IXP interconnection, new investments are only required in terrestrial fiber. In Africa, inland fiber deployment costs are mostly a function of labor 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 [120, 158, 196, 266]. 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 5.1). Details on the computations are available in the technical report [90]. In the last row of Table 5.1, we have also estimated the distance of terrestrial fiber to deploy per step and the corresponding costs. 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 work, 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 [163], SEACOM [259], MainOne [179], etc.), large CPs such as Google [1, 111, 231], local governments, Internet developmental institutions, or through regional projects setup by the African Union (AU). 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 [173, 174].

In the long run, stakeholders should consider making the proposed infrastructure redundant to improve its robustness to outages [70]. The first step would be to complete the set of links between the regional hubs,[12] 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.

[12] Recall that the solution in step-3 is a spanning tree and thus does not provide redundancy

Chapter 6

Conclusions and Future Work

In this chapter, which concludes this doctoral thesis, we first summarize our different studies, highlighting their main goals, as well as our contributions or key findings and their implications. Finally, we present our directions for future work.

6.1. Contributions of this doctoral thesis

6.1.1. African interdomain routing

The first step of our inspection of the African Internet has consisted of apprehending its interdomain routing. As a matter of fact, despite extensive studies on the Internet topology, much less was known about the AS level topology of the African Internet at the beginning of our research, especially when it comes to its IXP substrate. The main reason for this is the lack of Vantage Points (VPs) that are needed to obtain the proper information. Confronted with this near non-existence of measurements devices and the resulting lack of data on IP networks in the region, we have started by helping build its Internet measurement infrastructure (Section 3.1). We enhanced, from November 2013 to August 2016, the trustful and open RIPE Atlas measurement infrastructure in the region to shed light on both IPv4 and IPv6 topologies interconnecting local ISPs, while triggering the interests of other researchers to investigate this topic: we actively helped increase the number of VPs in Africa by 278.3 %.

This deployment effort has allowed us to perform, in the meantime, a four-year longitudinal study (Section 3.2.1) to understand the global African interdomain routing topology without bias towards any sub-region or country. This study has also aimed to reveal hidden topological changes (leading or not to communications performance improvements) and to identify practices of local ISPs that need to be encouraged or corrected. Overall, we carried out seven measurements campaign at random periods from 2013 to 2016, using all (or subsets of) the 324 probes hosted in 169 ASes operating in 40 African countries, the randomly selected 626 probes hosted in 380 ASes in 8 European countries and the randomly selected 329 probes in 195 ASes operating in the US. We have then adopted, as a best effort, a comprehensive method based on 10 data

sources combined with ping measurements to geolocate the IP addresses on the IP paths with high accuracy. While the IP to AS mapping with Team Cymru services has allowed us to obtain the corresponding AS paths, we have deduced the corresponding country paths with the geolocated IP addresses. We have then analyzed the collected datasets, proposing reproducible traceroute data analysis techniques suitable for the treatment of any set of similar measurements.

Our in-depth analysis has revealed a diversity of transit operators playing a role in the provision of both IPv4 and IPv6 African interdomain paths. Our inferred results, which depict the behavior of ISPs in the region, have shown evidence of the striking dependence of this large variety of ISP transit habits on socio-economic factors, such as the official language, the monetary region (specifically in West Africa), or the geographic location of the country in which the ISPs operate. We have also highlighted the prevalence of the dominant reliance on intercontinental ISPs for the establishment of continental connectivity. This leads to long AS paths and RTTs, sometimes among ISPs in the same country. We have shown a remaining lack of interconnection among African ISPs in IPv4 (South Africa being an exception) confirming the interest of initiatives to promote peering on the continent. We have then compared QoS within African countries, European countries, and the US to find that West African networks, in particular, need to promote investments in fiber networks and to implement traffic engineering techniques.

That said, we have shed light on traffic localization efforts made by stakeholders, as we have mapped, in our traceroute data, 62.2 % of the IXPs located in Africa and inferred their respective peers. In addition, we have highlighted the launch of recent IXPs or the usage of existing ones and quantified their impacts on AS path lengths and end-to-end delays. The study clearly demonstrates that to better assess interdomain routing in a continent, it is necessary to perform measurements from a diversified range of vantage points. It also raises the need for local ISPs to increase fiber deployment efforts (especially in the West), and intensify peering in the region.

An endemic phenomenon that may prevent existing or recently launched local IXPs from growing is interdomain congestion [314], notably in the context of increasing popularity of bandwidth-hungry applications such as streaming video, etc. The next step (Section 3.2.2) has thus aimed at inspecting the prevalence, investigating the causes, and measuring the impacts of congestion on peering links in the African IXP substrate. Towards this end, we deployed Ark probes (within networks peering) at six strategically selected African IXPs. Next, we run on those Vantage Points (VPs), the Time-Sequence Latency Probes (TSLP) algorithm, thereby collecting, every five minutes, RTTs to both edges of each mapped AS link for a whole year going from February 2016 to April 2017.

The thorough analysis of the collected dataset has allowed us to detect congestion events and quantify their corresponding periods and magnitudes at four IXPs. We have verified the events and investigated the causes by interviewing the IXP operators. Next, we have examined to which extent the existence of congestion negatively influences communications between a given AS and its neighbor. Our results have shown *no* evidence of widespread congestion: only 2.2 % of the discovered peering IP links have experienced (sustained or transient) congestion, which promotes

peering. We have then detailed the most interesting case studies, showing how RTTs to the far end have increased drastically during the congestion events, before discussing the implications of our observations for both research and network operations. Our findings suggest the need for ISPs to carefully monitor the provision of their peering links, so as to avoid or quickly mitigate the occurrence of such phenomenon (since an IXP only monitors ports sizes/traffic or ensures upgrades upon ISPs requests). Regulators may also define the maximum level of packet loss in those links to provide some protection to communications routed through local IXPs. Although our findings regarding the causes of congestion at IXPs may apply to IXPs in other regions, we have preferred not to attempt to generalize them beyond what we could directly observe and validate with the operators.

The above-listed results are uniquely obtained based on active measurements. Having publicly accessible information that could progressively pinpoint challenges and expose opportunities across the continent (for ISPs and CPs worldwide), was necessary to monitor progress and gaps in the African peering and interconnection landscape. In collaboration with the ISOC, we thus undertook to design a system which uses BGP routing data collected through passive measurements to profile the IXP substrate in a given Internet region and constantly monitor its growth. We implemented it for the Internet frontier, mindful of the strong push for local IXPs setups in the region [6], thus obtaining the "African" Route-collectors Data Analyzer (ARDA) platform. We have highlighted in Section 3.3.1 the key algorithms used to analyze pre-collected BGP data before including analysis results and use-cases of the relevance of the designed system for the Internet community.

In fact, ARDA examines which networks are directly connected to a local IXP, which networks are indirectly connected through that IXP and how far these (both direct and indirectly) connected networks span, in terms of country of origin. This open-source system, which is in production since April 21, 2017 [194], involves 63.1 % of African IXPs as of September 18, 2017. Networks participating at IXPs in Africa are encouraged to peer with the existing route-collectors at those IXPs to improve the accuracy of the platform. Further, all local IXPs are encouraged to deploy a route-collector of at least one type (PCH or RouteViews), and their members to provide BGP feeds to those facilities. This is particularly important because positive metrics values over time and a constant growth exhibited by ARDA regarding IP prefixes, peering ASNs, or origin ASNs may constitute a strong incentive for new members operating worldwide to join the corresponding IXPs and will surely influence upcoming investments decisions, notably from Internet developmental institutions.

6.1.2. African web ecosystem

After highlighting in our aforementioned studies the evolution of the Internet infrastructure in Africa, notably by exhibiting the proofs of the launch of new IXPs and their positive impacts on communications performance among local networks, we have investigated the African web ecosystem. We did so, aware that improving the underlying connectivity network is even more

useful with appropriately provisioned services to exploit it. In Section 4, we have thus measured the availability and utilization of web infrastructure serving end-users in Africa, whereas others have explored web infrastructure in developed regions [37, 76, 117, 123, 157, 210, 276, 315].

To achieve this, we have applied a comprehensive measurement methodology to collect data from a variety of sources. Our analysis of traffic data collected from a large European IXP has underlined the need for a better traffic localization in the African region. Using a significantly improved geolocation technique, we have then focused on Google to reveal that its content infrastructure in Africa is, indeed, expanding. We have, however, found that much of its web content is still served from the US and Europe, despite being the most popular website in many African countries. Next, we have repeated the same analysis across a number of other popular regional websites to find that even national African websites prefer to host their content abroad. To explore the reasons for this, we have evaluated some of the major bottlenecks faced by Content Providers (CPs) in Africa. Amongst other things, we have found a lack of peering between the networks hosting our probes, which prevent the sharing of CPs cache servers, as well as poorly configured DNS resolvers. We have therefore made a few suggestions for alleviating the issues observed (Section 4.3.4).

6.1.3. Topology and infrastructure

As a following step (Section 5), we have logically identified the interconnection challenges (Section 5.1.1). After that, we have looked towards the future of the African Internet while learning from its past (Section 5.1.2), *i.e.*, the reasons behind both failures and achievements of the Internet community in the region. Our goal in Section 5.2 has then been to propose 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 detailed in Sections 3.2.1 and Chapter 4, or in previous studies [49, 81, 85, 89, 117], 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 [72, 203, 211, 268]) are not always realizable due to the prevailing external factors. In addition, we have shown in Section 5.2.2 how naive approaches, which do not take into account prevailing socio-economic realities of the region are infeasible in practice.

In this study, we have first introduced an innovative framework that acknowledges the existence of geographical, political, and socio-economic realities, which 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 a hierarchical IXP substrate: ISP source – local IXP (– regional IXP hub – local IXP) – ISP destination.

Next, we have evaluated the proposed layout and quantified the benefits using extensive simulations with C-BGP. Our results have shown how 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 has demonstrated that it is possible to obtain shorter AS paths and better performance, if local ISPs intensify peering and CPs were to deploy caches at the designated regional hubs. By doing so, we have highlighted 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 have identified, in Section 5.2.12.2.2, three options to realize the proposed IXP infrastructure, 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 [120, 158, 196, 266], we have 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.

All in all, there are some key points, which resort from this research as essential steps in the process of revealing and reshaping the Internet ecosystem in developing regions to help meet the challenge of making the Internet accessible, fast, or affordable for end-users and beneficial for stakeholders:

1. The need to build an open, large, and trustful network measurement infrastructure for transparency in the data collection/publication and for involving all stakeholders, as *we cannot improve what we cannot measure*.

2. A focused and longitudinal data collection with comprehensive measurement methods on which are performed rigorous and reproducible data analysis techniques (some of which are proposed in this thesis) for :

   (a) Inspecting the interdomain routing topology and determining ISPs practises that need to be encouraged or corrected.

   (b) Investigating the nature, prevalence, causes, and impacts of congestion on peering links at local IXPs, since its prevalence may prevent those local Internet markets from growing.

   (c) Monitoring the evolution of local IXPs, as they are essential for an improved local interconnection and therefore more traffic localization.

(d) Exploring the web ecosystem to determine insights for significantly improving future ISPs and CPs deployments.

3. The suggestion, based on the holistic knowledge acquired on the Internet in the region during the aforementioned studies, of a participative and innovative framework for enriched connectivity and increased CPs presence, taking into account external factors that affect connectivity in the region to ensure that the solution is realizable.

6.2. Future Work

Each of our studies can frequently be repeated by automatized systems to inform the Internet community in real-time, similarly to [87]. A key study that will result from our research is a detailed comparison of the interdomain routing in Africa to that in Latin America and the Caribbean (LAC) [29,101,170,313], with the objective of gaining further insights into their common characteristics or deficiencies w.r.t. Internet connectivity and performance. This ongoing study, which is based on routing data collected through passive measurements, aims at suggesting to stakeholders in both regions some ways for cooperation to solve the issues identified and quantifying the expected impacts. Along these lines, the African Route collectors Data Analyzer (ARDA) [87] is planned to be extended to the LACNIC region.

Our future work includes to keep monitoring the African interdomain routing and the evolution of web infrastructure in the region. The former can be achieved by populating the databases of the application [256] with up-to-date measurements data collected through full-mesh measurements run by RIPE Atlas probes or any other vantage point deployed within African networks. The latter can be achieved by building a platform which automatizes the run of the diverse measurements/analysis carried out during our study of the web ecosystem, and frequently launches them to give up-to-date information to the Internet community.

Moving to our inspection of the causes of congestion at IXPs, we plan to continue deploying additional Ark probes at networks and IXPs operating in developed and developing regions, including Africa, to increase our coverage of the African sub-regions that have not received much attention so far. Meanwhile, we intend to keep analyzing collected TSLP data to delve into the dynamics and causes of congestion at IXP infrastructures and compare the results with those presented in this thesis. Further, it will be interesting to correlate our observations from TSLP measurements with data from IXP operators. Towards this end, we are working on strengthening our relationship with operators in the African region to make such a study feasible in the future.

Regarding our proposed IXP interconnection framework, we are aware that there may be further socio-economic factors beyond the ones we could capture, which influence connectivity in the African region. We plan to engage further with stakeholders to discover those parameters and capture them in our framework. Modeling the internal topology of ASes with a regional scope/area, rather than using a single router is also planned as future work. Next, we intend to reach out to local ISPs 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 have eliminated 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.

Appendix A

Curriculum of the AXIS Workshops

We present below the curriculum of the English version of the ISOC's AXIS [6,141] technical workshops. We taught these courses punctuated by several practical labs, as an assistant instructor, to participants from English-speaking countries such as Liberia (LR), Ethiopia (ET). They were taught in French, as a lead instructor, to participants from French-speaking countries Burkina Faso (BF), Benin (BJ), Niger (NE), Mauritania (MR), and Congo-Brazzaville (CG).

Table A.1: Curriculum of the AXIS workshops entitled "technical aspects of setting up, operating and administering IXPs"

HOURS	**ACTIVITIES**
	DAY 1: IP Resources and Routing Basics
08h00 – 08h30	Registration
08h30 – 10h30	**Introduction to IP (Internet Protocol) and Number Resources** ▪ Introduction to IPv4 and IPv6 ▪ Internet Number Resources
10h30 – 11h00	Networking, Coffee Break, and Discussions
11h00 – 13h00	**Introduction to Routing** ▪ Routing Basics
13h00 – 14h00	Networking, Lunch, and Discussions
14h00 – 16h00	**Introduction to Internal Gateway Protocols (IGPs)** ▪ Introduction to OSPF v2 ▪ OSPF Deployment for ISPs

HOURS	ACTIVITIES
16h00 – 16h30	Networking, Coffee Break, and Discussions
16h30 – 18h00	**OSPF Lab Exercises** ▪ Basic OSPF (Module 01a. IPv4 + OSPFv2)
	DAY 2: Introduction to BGP
08h00 – 08h30	Registration
08h00 – 10h30	**Introduction to Exterior Gateway Protocol (EGPs)** ▪ Introduction to BGP
10h30 – 11h00	Networking, Coffee Break, and Discussions
11h00 – 13h00	**BGP Hands-on Lab Exercise** ▪ iBGP Lab Exercise (Module 1c. IPv4 + OSPFv2 + iBGP)
13h00 – 14h00	Networking, Lunch, and Discussions
14h00 – 16h00	**BGP Attributes and Scaling Techniques** ▪ BGP Scaling Techniques ▪ BGP Attributes
16h00 – 16h30	Networking, Coffee break, and Discussions
16h30 – 18h00	**BGP Hands-on Lab Exercise** ▪ eBGP Lab Exercise (Module 6a. IPv4+OSPFv2 + iBGP + eBGP)
	DAY 3: BGP Policy Control and Multihoming
08h30 – 08h30	Registration
08h30 – 10h30	**BGP Policy and Best Practices** ▪ Implementing BGP Policy Controls ▪ BGP Best Practices for ISPs
10h30 – 11h00	Networking, Coffee Break, and Discussions

HOURS	ACTIVITIES
11h00 – 13h00	**BGP Policy Control Hands-on Lab Exercise** ▪ BGP Filtering Exercise (Module 7 IPv4 + OSPF + iBGP + eBGP + Prefix-Lists and ASN Path)
13h00 – 14h00	Networking, Lunch, and Discussions
14h00 – 16h00	**BGP Policy Control Hands-on Lab Exercise** ▪ BGP Filtering Exercise (Module 7 IPv4 + OSPF + iBGP + eBGP + BGP Communities)
16h00 – 16h30	Networking, Coffee break, and Discussions
16h30 – 18h00	**Introduction to Multi-homing** ▪ Simple Multi-homing ▪ Advanced Multi-homing Techniques
DAY 4: Introduction to IXPs	
08h30 – 08h30	Registration
08h30 – 10h30	**Introduction to Multi-homing** ▪ Multi-homing Hands-on Lab Exercise • BGP Local Preference Exercise (Module 8. IPv4 + OSPF + iBGP + eBGP + PrefixList + Local-Pref)
10h00 – 10h30	Networking, Coffee break, and Discussions
10h30 – 11h00	**Introduction to Multi-homing** ▪ Multi-homing Hands-on Lab Exercise • BGP Local Preference Exercise (Module 8. IPv4 + OSPF + iBGP + eBGP + PrefixList + ASPath-Prepend) Scalable Network Design ▪ ISP Network Design

HOURS	ACTIVITIES
13h00 – 14h00	Networking, Lunch, and Discussions
14h30 – 16h00	**Introduction to IXPs** ▪ Euro-IX Video ▪ Value of Peering ▪ IXP Network Design
16h00 – 16h30	Networking, Coffee break, and discussions
16h30 – 18h00	**IXP Hands-on Lab Exercise** ▪ Advanced IXP Lab Exercise (Module 16a OSPF+iBGP+eBGP+IXP+Transit)
DAY 5: IXP and Traffic Monitoring	
08h30 – 08h30	Registration
08h30 – 10h30	**IXP Hands-on Lab Exercise** ▪ Advanced IXP Lab Exercise (Module 16a OSPF+iBGP+eBGP+IXP+Transit)
10h30 – 11h00	Networking, Coffee break, and Discussions
11h00 – 13h00	**IXP Hands-on Lab Exercise** ▪ Advanced IXP Lab Exercise (Module 16a OSPF+iBGP+eBGP+IXP+Transit)
13h00 – 14h00	Networking, Lunch, and Discussions
14h00 – 16h00	**Network Monitoring and Value Added Services** ▪ Traffic Monitoring and Flow Analysis ▪ Selecting an IXP ▪ PeeringDB and Role of Peering Coordinator
16h00 – 16h30	Networking, Coffee break, and Discussions
16h30 – 18h00	Way Forward and Certificate Ceremony

Appendix B

Survey of the African IXPs Operators

This survey available at [291] was conducted during my internship at CAIDA from January to February 2016. The goal was to collect detailed information on existing IXPs, their member ASNs, and the setup of their infrastructures to suggest an innovative an realistic framework for a distributed IXP infrastructure in Africa and thus reshaping the African Internet, goal achieved in [79, 90]. It targeted the IXPs existing in the African region during that period. It was conducted in both French and English, depending on the official language of the country hosting the IXP. On February 28, 2017, at the end of the survey, we found out that 37 of them were active and this number remained steady until July 2017. In total, 59.4 % of IXPs answered.

SURVEY: Collecting African IXP Colocation Data for research purposes (following on Joint Study which Identifies Infrastructure Development as Top Priority for ICT in Africa)

We are currently performing a study in line with the conclusion 3 of the published ISOC report [145] (*i.e.*, Infrastructure development is the top-most priority of ICT African Policymakers today). To give ourselves the means to analyze (for research purposes) the impact that interconnecting African IXPs would have on the Internet ecosystem, we have been collecting colocation data at these IXPs from IXP websites, PeeringDB, Telegeography's Internet Exchange Map, and PCH. Since that information may not be up-to-date, we have tried during the last two months to validate it by sending the survey below to IXP NOC administrators:

1. Could you please give us a list of your members with their corresponding AS numbers?
2. Could you please add/validate your IPv4 and IPv6 peering LAN as well as your Administration LAN?
3. Could you please add the IXP AS numbers if there are any?
4. *a*) Is there bilateral peering between network operators at your IXP?

   b) Is there multilateral peering through route servers provided by the IXP (this implies the route servers have an ASN)?

c) (optional) can peers influence advertisements through the route servers to other peers (by communities)?

d) Must peers use the route servers (Mandatory multi-lateral peering)?

e) if bilateral peering is done, multilateral peering is available and not mandatory, then please give the approximate shares/percentages of peers doing *(i)* only bilateral peering, *(ii)* both bilateral and multilateral peering, or *(iii)* only multi-lateral peering.

5. Do you think your IXP would be open to connecting with other IXPs in order to increase peering in your sub-region?

20 IXPs (52%) replied to our emails out of 38. We thank all IXP NOC admins who answered. However, we believe this is insufficient and are willing to reach 80 – 90% of IXPs data validation and hence a trustful list of the IXPs/ IXP members.

Here are the 18 IXPs that did not respond: RIMIX (MR), SiXP (SD), MUIXP (MU), ZINX (ZW), NAPAFRICA JB (ZA), NAPAFRICA DB (ZA), NAPAFRICA CT (ZA), BINX (BW), ANGOLA-IX (AO), MBABANE-IX (SZ), CIVIX (CI), LIBERIA-IX (LR), BFIX (BF), SIXP-GM (GM), TUNIXP (TN), MOZIX (MZ), RENATERIX (RE), WHINDOEK-IX (NA).

In case the Network Operating Center (NOC) administrator/members of any of those IXPs are members of this mailing list and have not received our mail/have not replied to our survey, it will be helpful, it they could please do so. The survey can be filled in 5 – 10 mn (For question 1, you can give an URL to your looking glass or an updated website if there is any). If you are an ISP, member of these IXPs you can also reply in the thread (so that the information is publicly available and useful for everyone). Finally, if you have the contact of any of those NOC admins, feel free to send it to us.

We have not found any information about the following IXPs. We do not know if they exist and up to now are not planning to include them in the study: SEYCHELLES-IX (SC), MGIX (MG), LIXP (LS), and GA-IXP (GA).

Any comment about them is welcome as well.

[**February 26, 2016 (Date of the sending of this mail to The African IXP association (Af-IX) mailing list)**]

References

[1] Google Africa Blog. https://africa.googleblog.com/, 2017.

[2] Emile Aben. Measuring Countries and IXPs with RIPE Atlas. https://labs.ripe.net/Members/emileaben/measuring-ixps-with-ripe-atlas, March 2015.

[3] Emile Aben. IXP Country Jedi. http://sg-pub.ripe.net/emile/ixp-country-jedi/latest/, August 2017.

[4] Emile Aben, Dimitris Mavrommatis, Massimiliano Stucchi, Andreas Strikos, and Byron Ellacott. IXP Country Jedi: Probe Mesh Measurements. https://github.com/emileaben/ixp-country-jedi, March 2015.

[5] African Union (AU). Study on Harmonisation of Telecommunication, Information and Communication Technologies Policies and Regulation in Africa. *Draft Report*, March 2008.

[6] African Union (AU). African Internet eXchange System (AXIS). www.au.int/web/en/axis, August 2017.

[7] African Union Commission and New Zealand Ministry. African Union Handbook 2015. Technical report, African Union (AU), 2015.

[8] African Union Commission and New Zealand Ministry. African Union Handbook 2016. Technical report, African Union (AU), 2016.

[9] African Union Commission and New Zealand Ministry. African Union Handbook 2017. Technical report, African Union (AU), January 2017.

[10] 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/, April 2017.

[11] AfriNIC. AfriNIC Database. ftp://ftp.afrinic.net/, August 2017.

[12] Agence d'Information d'Africque Centrale. Réseau fibre optique : Brazzaville et Kinshasa bientôt interconnectées. http://adiac-congo.com/content/reseau-fibre-optique-brazzaville-et-kinshasa-bientot-interconnectees-31551, April 2015.

[13] Bernhard Ager, Nikolaos Chatzis, Anja Feldmann, Nadi Sarrar, Steve Uhlig, and Walter Willinger. Anatomy of a Large European IXP. In *Proceedings of the ACM SIGCOMM Conference on Applications, Technologies, Architectures, and Protocols for Computer Communication*, pages 163–174. SIGCOMM, 2012.

[14] Bernhard Ager, Wolfgang Mühlbauer, Georgios Smaragdakis, and Steve Uhlig. Web Content Cartography. In *Proceedings of the 2011 ACM SIGCOMM conference on Internet measurement conference*, pages 585–600. ACM, 2011.

[15] Akamai. Internet Connection Speeds and Adoption Rates by Geography. https://www.akamai.com/uk/en/about/our-thinking/state-of-the-internet-report/state-of-the-internet-connectivity-visualization.jsp, November 2017.

[16] Akamai. State of the Internet. https://www.akamai.com/us/en/about/our-thinking/state-of-the-internet-report/, November 2017.

[17] Abossé Akue-Kpakpo. Study on International Internet Connectivity in Sub-Saharan Africa. *Telecommunication Development Bureau*, March 2013.

[18] Alexa. Alexa Websites. http://www.alexa.com/topsites/, August 2017.

[19] UbuntuNet Alliance. Ubuntunet alliance. https://ubuntunet.net/, 2017.

[20] APNIC. APNIC Database. ftp://ftp.apnic.net/, August 2017.

[21] APNIC. Visible ASNs: Customer Populations (Est.). http://stats.labs.apnic.net/cgi-bin/aspop?c=, August 2017.

[22] AquaLab. DASU – A Platform for Measurement Experimentation and Broadband Characterization. http://www.aqualab.cs.northwestern.edu/projects/115-dasu-isp-characterization-from-the-network-edge, 2017.

[23] ARIN. ARIN Database. ftp://ftp.arin.net/, August 2017.

[24] Brice Augustin, Xavier Cuvellier, Benjamin Orgogozo, Fabien Viger, Timur Friedman, Matthieu Latapy, Clémence Magnien, and Renata Teixera. Avoiding Traceroute Anomalies with Paris Traceroute. In *Proceedings of the ACM SIGCOMM Internet Measurement Conference (IMC)*, pages 153–158. ACM, October 2006.

[25] Brice Augustin, Balachander Krishnamurthy, and Walter Willinger. IXPs: Mapped? In *Proceedings of the ACM SIGCOMM Internet Measurement Conference (IMC)*, pages 336–349. ACM, 2009.

[26] Vaibhav Bajpai and Jürgen Schönwälder. A Survey on Internet Performance Measurement Platforms and Related Standardization Efforts. *IEEE Communications Surveys and Tutorials*, 17(3), 2015.

[27] Philip Bates. Submarine cables in Sub-Saharan Africa: Terrestrial Networks Need to Keep up. Technical report, Analysys mason, 2014.

[28] Benin-IX. Benin Internet Exchange Point. https://benin-ix.org.bj/, August 2017.

[29] Sofía Silva Berenguer, Esteban Carisimo, J Ignacio Alvarez-Hamelin, and Francisco Valera Pintor. Hidden Internet Topologies info: Truth or Myth? In *LANCOMM@ SIGCOMM*, pages 4–6, 2016.

[30] Jérôme Bezzina. Interconnection Challenges in a Converging Environment. *The World Bank*, 2005.

[31] Zachary Bischof, Fabián Bustamante, and Rade Stanojevic. Need, Want, Can Afford - Broadband Markets and the Behavior of Users. In *Proceedings of the ACM Internet Measurement Conference (IMC)*, pages 73–86. ACM, 2014.

[32] Project BISmark. Project bismark– broadband internet service benchmark. http://projectbismark.net/#, September 2017.

[33] Olivier Bonaventure. *Computer Networking: Principles, Protocols, and Practice*. The Saylor Foundation, 2011.

[34] Stephane Bortzmeyer. RIPE Atlas Probes and the User-Defined Measurements. https://ripe67.ripe.net/presentations/153-ripe-atlas-udm-api-1.pdf, October 2013.

[35] Marc Bruyere. TOUSIX First OpenFlow European IXP. https://www.youtube.com/watch?v=VLgJcz-opws&feature=youtu.be, March 2016.

[36] Randy Bush. RIPE Atlas Probes. https://archive.psg.com/130614.AtlasProbes.pdf, June 2013.

[37] Matt Calder, Xun Fan, Zi Hu, Ethan Katz-Bassett, John Heidemann, and Ramesh Govindan. Mapping the Expansion of Google's Serving Infrastructure. In *Proceedings of the ACM Internet Measurement Conference (IMC)*, pages 313–326. ACM, 2013.

[38] Ignacio Castro, Juan Camilo Cardona, Sergey Gorinsky, and Pierre Francois. Remote Peering: More Peering Without Internet Flattening. In *Proceedings of the 10th ACM International on Conference on emerging Networking Experiments and Technologies*, pages 185–198. ACM, 2014.

[39] Center for Applied Internet Data Analysis (CAIDA). BGPstream. http://bgpstream.caida.org, October 2015.

[40] Center for Applied Internet Data Analysis (CAIDA). Archipelago (Ark) Measurement Infrastructure. http://www.caida.org/projects/ark/, August 2017.

[41] Center for Applied Internet Data Analysis (CAIDA). AS Relationships. http://www.caida.org/data/as-relationships, August 2017.

[42] Center for Applied Internet Data Analysis (CAIDA). Automated Autonomous System (AS) Ranking. Research Project. http://as-rank.caida.org, August 2017.

[43] Center for Applied Internet Data Analysis (CAIDA). Border Mapping (bdrmap) Dataset. http://www.caida.org/data/active/bdrmap_dataset.xml, 2017.

[44] Center for Applied Internet Data Analysis (CAIDA). CAIDA AS Relationships Data. Research Project. http://data.caida.org/datasets/as-relationships/, August 2017.

[45] Center for Applied Internet Data Analysis (CAIDA). Internet Topology Datasets Collected on the Archipelago (Ark) Infrastructure. http://www.caida.org/projects/ark/topo_datasets.xml, August 2017.

[46] Central Intelligence Agency (CIA). The World Factbook. https://www.cia.gov/library/publications/the-world-factbook/docs/guidetowfbook.html, August 2017.

[47] Balakrishnan Chandrasekaran, Georgios Smaragdakis, Arthur Berger, Matthew Luckie, and Keung-Chi Ng. A Server-to-Server View of the Internet. In *Proceedings of the 11th ACM Conference on Emerging Networking Experiments and Technologies*, page 40. ACM, 2015.

[48] Nikolaos Chatzis, Georgios Smaragdakis, Anja Feldmann, and Walter Willinger. There is More to IXPs than Meets the Eye. *ACM SIGCOMM Computer Communication Review*, (5):19–28, 2013.

[49] Josiah Chavula, Nick Feamster, Antoine Bagula, and Hussein Suleman. Quantifying the Effects of Circuitous Routes on the Latency of Intra-Africa Internet Traffic: A Study of Research and Education Networks. In *AFRICOMM*, pages 64–73. Springer International Publishing, 2014.

[50] Johnson Chen and Ljiljana Trajkovic. Analysis of Internet Topology Data. In *Proceedings of the International Symposium on Circuits and Systems, ISCAS'04*, volume 4, May 2004.

[51] Pu-Shih Daniel Chen, Amber D Lambert, and Kevin R Guidry. Engaging Online Learners: The Impact of Web-based Learning Technology on College Student Engagement. *Computers & Education*, 54(4):1222–1232, 2010.

[52] Shaohua Chen and Martin Ravallion. The Developing World is Poorer than We Thought, But no Less Successful in the Fight Against Poverty. *The Quarterly Journal of Economics*, 125(4):1577–1625, 2010.

[53] Fangfei Cheng, Ramesh K Sitaraman, and Marcelo Torres. End-user Mapping: Next Generation Request Routing for Content Delivery. In *ACM SIGCOMM Computer Communication Review*, volume 45, pages 167–181. ACM, 2015.

[54] Marshini Chetty, Srikanth Sundaresan, Sachit Muckaden, Nick Feamster, and Enrico Calandro. Measuring Broadband Performance in South Africa. In *Proceedings of the 4th Annual Symposium on Computing for Development*, page 1. ACM, 2013.

[55] Kc Claffy. The 8th Workshop on Active Internet Measurements (AIMS-8) Report. *ACM SIGCOMM Computer Communication Review*, 46:23–29, October 2016.

[56] CNN. Cameroon Goes Offline After Anglophone Revolt. http://edition.cnn.com/2017/02/03/africa/internet-shutdown-cameroon/index.html, February 2017.

[57] Computing Research & Education (CORE). CORE Conference Portal. http://portal.core.edu.au/conf-ranks/, September 2017.

[58] Les Cottrell. Routing in Africa. https://confluence.slac.stanford.edu/display/IEPM/Routing+in+Africa, October 2013.

[59] Pelsser Cristel, Cittadini Luca, Vissicchio Stefano, and Randy Bush. From Paris to Tokyo: On the Suitability of Ping to Measure Latency. In *Proceedings of the ACM Internet Measurement Conference (IMC)*, pages 427–432. ACM, 2013.

[60] Mark Crovella and Balachander Krishnamurthy. *Internet Measurement: Infrastructure, traffic and applications*. John Wiley & Sons, Inc., 2006.

[61] Jon Crowcroft, Adam Wolisz, and Arjuna Sathiaseelan. Towards an Affordable Internet Access for Everyone: The Quest for Enabling Universal Service Commitment (Dagstuhl Seminar 14471). *Dagstuhl Reports*, 4(11), 2015.

[62] Alberto Dainotti, Ethan Katz-Bassett, and Xenofontas Dimitropolous. The BGP Hackathon 2016 Report. ASSOC COMPUTING MACHINERY 2 PENN PLAZA, STE 701, NEW YORK, NY 10121-0701 USA, 2016.

[63] Angus Deaton. Measuring Poverty in a Growing World (or Measuring Growth in a Poor World). *The Review of Economics and Statistics*, 87(1):1–19, 2005.

[64] Timothy Denton. Hurricane Electric's Observations on Internet Exchange Point Management. https://conference.apnic.net/data/41/apricot-he-presentation-1_1456014430.pdf, February 2016.

[65] Amogh Dhamdhere and Constantine 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.

[66] Digital Element. Netacuity. http://www.digital-element.net/ip_intelligence/ip_intelligence.html, August 2017.

[67] Xenofontas Dimitropoulos, Dmitri Krioukov, George Riley, and Kc Claffy. Classifying the Types of ASes in the Internet. In *SIGCOMM*, 2005.

[68] Chris Dixon. *Rural Development in the Third World*. Routledge, 2015.

[69] Rohit Dube. A Comparison of Scaling Techniques for BGP. *ACM SIGCOMM Computer Communication Review*, 29(3):44–46, 1999.

[70] Ramakrishnan Durairajan, P. Barford, J. Sommers, and Walter Willinger. InterTubes: A Study of the US Long-haul Fiber-optic Infrastructure. In *ACM SIGCOMM Computer Communication Review*, volume 45, pages 565–578. ACM, 2015.

[71] 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, May 2014.

[72] 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, October 2015.

[73] Ekinops. Liquid Telecom Deploys New Optical Network in Africa Using EKINOPS Long-Haul DWDM Technology. https://www.ekinops.net/newswire/customers/liquid-telecom-deploys-new-optical-network-in-africa-using-ekinops-long-haul-dwdm-technology, July 2013.

[74] Yehia Elkhatib, Gareth Tyson, and Michael Welzl. Can SPDY Really Make the Web Faster? In *Networking Conference, 2014 IFIP*, pages 1–9. IEEE, 2014.

[75] Kevin R Fall and W Richard Stevens. *TCP/IP illustrated, volume 1: The protocols*. Addison-Wesley, 2011.

[76] Xun Fan, Ethan Katz-Bassett, and John Heidemann. Assessing Affinity Between Users and CDNs sites. In *International Workshop on Traffic Monitoring and Analysis*, pages 95–110. Springer, 2015.

[77] Rodérick Fanou. On the State of Interdomain Routing in Africa. Master's thesis, Universidad Carlos III de Madrid, Spain, 2014.

[78] Rodérick Fanou, Amogh Dhamdhere, and Francisco Valera. Investigating the Causes of Congestion on the African IXP substrate. In *Proceedings of the ACM SIGCOMM Internet Measurement Conference (IMC)*. ACM, November 2017.

[79] Rodérick Fanou, Valera Francisco, Pierre Francois, and Amogh Dhamdhere. Reshaping the african internet: From scattered islands to a connected continent. *Computer Communications*, September 2017.

[80] Rodérick Fanou and Pierre Francois. Drawing the Map of the West African Internet. http://netcom.it.uc3m.es/whats-new/news/2014/drawing-map-west-african-internet, February 2014.

[81] Rodérick Fanou, Pierre Francois, and Emile Aben. On the Diversity of Interdomain Routing in Africa. In *International Conference on Passive and Active Network Measurement (PAM)*, March 2015.

[82] Rodérick Fanou, Pierre Francois, and Emile Aben. On the Diversity of Interdomain Routing in Africa (RIPE Labs). https://labs.ripe.net/Members/fanou_roderick/on-the-diversity-of-interdomain-routing-in-africa, May 2015.

[83] Rodérick Fanou, Pierre Francois, and Emile Aben. On the Interdomain Topology of Africa (Invited Poster). In *The 5th PhD School on Traffic Monitoring and Analysis (TMA) and The 7th International Workshop on TMA*, 2015.

[84] Rodérick Fanou, Pierre Francois, Emile Aben, Michuki Mwangi, Nishal Goburdhan, and Francisco Valera. Four Years Tracking Unreaveled Topology Changes in the African Interdomain: Technical Report. https://techrep_amc_journal:haHeudcis@fourier.networks.imdea.org/external/techrep_amc_journal/index/index.html, June 2017.

[85] Rodérick Fanou, Pierre Francois, Emile Aben, Michuki Mwangi, Nishal Goburdhan, and Francisco Valera. Four Years Tracking Unrevealed Topological Changes in the African Interdomain. *Computer Communications*, 106:117–135, July 2017.

[86] Rodérick Fanou, Vìctor Sánchez-Agüero, Francisco Valera, Michuki Mwangi, and Jane Coffin. The ISOC Compass to Support Peering Growth in the African Region: a Route-collectors Data Analyzer. http://isoc-ny.org/afpif2016/slides/11_African_Route_Collectors_Data_Analyzer_v8_Roderick.pdf, August 2016.

[87] Rodérick Fanou, Vìctor Sánchez-Agüero, Francisco Valera, Michuki Mwangi, and Coffin Jane. African Route-collector Data Analyzer (ARDA). https://arda.af-ix.net/, August 2017.

[88] Rodérick Fanou, Gareth Tyson, Eder Leao Fernandes, Pierre Francois, Francisco Valera, and Arjuna Sathiaseelan. Exploring and Analysing the African Web Ecosystem: Technical Report. https://techrepwebinf:bRCA9hFZ@fourier.networks.imdea.org/external/techrep_web_infrastructure/index/, June 2017.

[89] Rodérick Fanou, Gareth Tyson, Pierre Francois, and Arjuna Sathiaseelan. Pushing the Frontier: Exploring the African Web Ecosystem. In *Proceedings of the 25th International Conference on World Wide Web (WWW)*, April 2016.

[90] Rodérick Fanou, Francisco Valera, Pierre Francois, and Amogh Dhamdhere. Reshaping the African Internet: From Scattered Islands to a Connected Continent (Technical Report). https://fourier.networks.imdea.org/external/techrep_reshaping/index, June 2017.

[91] R. Farahbakhsh, A. Cuevas, A. M. Ortiz, X. Han, and N Crespi. How Far is Facebook from Me? Facebook Network Infrastructure Analysis. *IEEE Communications Magazine*, 53:134–142, 2015.

[92] Peyman Faratin, David Clark, Steven Bauer, William Lehr, Patrick Gilmore, and Arthur Berger. The Growing Complexity of Internet Interconnection. 2008.

[93] Arnaud Fenioux. Why and How to Interconnect IXP. https://es.slideshare.net/InternetSociety/why-and-how-to-interconnect-ixp, September 2015.

[94] Agustin Formoso, Josiah Chavula, Amreesh Phokeer, Arjuna Sathiaseelan, and Gareth Tyson. Dissecting the African Internet: An Intra-Continental Study (White paper). http://www.eecs.qmul.ac.uk/~tysong/files/africa-internet.pdf, September 2017.

[95] Augustin Kwasi Fosu. Growth, Inequality and Poverty in Sub-Saharan Africa: Recent Progress in a Global Context. *Oxford Development Studies*, 43(1):44–59, 2015.

[96] Raspberry Pi Foundation. Raspberry pi. https://www.raspberrypi.org/, September 2017.

[97] France Diplomatie. France Diplomatie: Country Files. http://www.diplomatie.gouv.fr/fr/conseils-aux-voyageurs/conseils-par-pays/, August 2017.

[98] FranceIX. FranceIX: Interconnection with Other IXPs. http://www.franceix.net/en/solutions/interconnection/, August 2017.

[99] Benjamin Frank, Ingmar Poese, Georgios Smaragdakis, Anja Feldmann, Bruce M Maggs, Steve Uhlig, Vinay Aggarwal, and Fabian Schneider. Collaboration Opportunities for Content Delivery and Network Infrastructures. *Recent Advances in Networking*, 1:305–377, 2013.

[100] Hernan Galperin. Connectivity in Latin America and the Caribbean: The role of Internet Exchange Points (IXPs). *Internet Society, November*, 2013.

[101] Hernán Galperín. Localizing internet infrastructure: Cooperative peering in latin america. *Telematics and Informatics*, 33:631–640, 2016.

[102] Lixin Gao. On Inferring Autonomous System Relationships in the Internet. *IEEE/ACM Transactions on Networking (ToN)*, 9(6):733–745, December 2001.

[103] Daniel Genin and Jolene Splett. Where in the Internet is congestion? http://arxiv.org/abs/1307.3696, July 2013.

[104] Ghana Internet Exchange Association (GIXA). GIXA Website. www.gixa.org.gh/, August 2017.

[105] Manaf Gharaibeh, Anant Shah, Bradley Huffaker, Hang Zhang, Roya Ensafi, and Christos Papadopoulos. A Look at Router Geolocation in Public and Commercial Databases. In *Proceedings of the ACM Internet Measurement Conference (IMC)*, 2017.

[106] J.S. Gilmore, N.F. Huysamen, P. Cronje, M.C. de Klerk, and A.E. Krzesinski. Mapping the African Internet. In *Proceedings Southern African Telecommunication Networks and Applications Conference (SATNAC), Mauritius*, September 2007.

[107] Vasileios Giotsas, Matthew Luckie, Bradley Huffaker, et al. Inferring Complex AS Relationships. In *Proceedings of the ACM Internet Measurement Conference (IMC)*, pages 23–30. ACM, 2014.

[108] Nishal Goburdhan. Madagascar IX (MGIX) Looking Glass. http://wiki.mgix.mg/display/PUB/Peers, April 2016.

[109] Google. Google Maps. https://www.google.com/maps, June 2017.

[110] Google. The True Size of. http://thetruesize.com/, August 2017.

[111] Google, Convergence Partners, International Finance Corporation, and Mitsui co-invest in CSquared. CSquared, a Better Way to Connect. http://www.csquared.com/, August 2017.

[112] Google Developers. Google Maps APIs: Distance Matrix API. https://developers.google.com/maps/documentation/distance-matrix/start, June 2017.

[113] Enrico Gregori, Alessandro Improta, Luciano Lenzini, and Chiara Orsini. The Impact of IXPs on the AS-level Topology Structure of the Internet. *Computer Communications*, 34(1):68–82, 2011.

[114] GrenoblIX. GrenoblIX: Infrastructure. http://www.grenoblix.net/en/infrastructure, August 2017.

[115] GSMA Association. The Mobile Economy 2017. Technical report, GSMA, 2017.

[116] Bamba Gueye, Artur Ziviani, Mark Crovella, and Serge Fdida. Constraint-based geolocation of internet hosts. *IEEE/ACM Transactions On Networking*, 14(6):1219–1232, 2006.

[117] Arpit Gupta, Matt Calder, Nick Feamster, Marshini Chetty, Enrico Calandro, and Ethan Katz-Bassett. Peering at the Internet's Frontier: A First Look at ISP Interconnectivity in Africa. In *International Conference on Passive and Active Network Measurement (PAM)*, pages 204–213. Springer, March 2014.

[118] Endika Gutiérrez and Austin Keeley. World universities. https://github.com/endSly/world-universities-csv, October 2017.

[119] Hamed Haddadi, Miguel Rio, Gianluca Iannaccone, Andrew Moore, and Richard Mortier. Network Topologies: Inference, Modeling, and Generation. In *IEEE Communications Surveys and Tutorials*, volume 10, pages 48–69. IEEE, 2008.

[120] Hamilton Research Ltd 2017. Africa Bandwidth Maps. http://www.africabandwidthmaps.com/?p=5348, May 2017.

[121] John Hawkinson and Tony Bates. Guidelines for Creation, Selection, and Registration of an Autonomous System (AS), 1996.

[122] Rick Hofstede, Pavel Čeleda, Brian Trammell, Idilio Drago, Ramin Sadre, Anna Sperotto, and Aiko Pras. Flow Monitoring Explained: From Packet Capture to Data Analysis With NetFlow and IPFIX. volume 16, pages 2037–2064. IEEE, 2014.

[123] Qi Huang, Ken Birman, Robbert van Renesse, Wyatt Lloyd, Sanjeev Kumar, and Harry C Li. An Analysis of Facebook Photo Caching. In *Proceedings of the Twenty-Fourth ACM Symposium on Operating Systems Principles*, pages 167–181. ACM, 2013.

[124] Bradley Huffaker, Marina Fomenkov, and K Claffy. Geocompare: a comparison of public and commercial geolocation databases. *Proc. NMMC*, pages 1–12, 2011.

[125] Bradley Huffaker, Marina Fomenkov, and K Claffy. Internet topology data comparison. In *Citeseer*, 2012.

[126] Bradley Huffaker, Marina Fomenkov, and KC Claffy. Internet Topology Data Comparison. *Citeseer*, 2012.

[127] Bradley Huffaker, Marina Fomenkov, Daniel J Plummer, David Moore, and K Claffy. Distance metrics in the Internet. In *EEE International, Telecommunications Symposium, (ITS2002)*, 2002.

[128] Hurricane Electric. Hurricane Electric Internet Services: BGP Toolkit Home. http://bgp.he.net/, 2017.

[129] G Huston. Interconnection, Peering and Settlements—Part II. The Internet Protocol Journal 2 (2), 1999.

[130] Geoff Huston. Interconnection, peering, and settlements. In *proc. INET (Vol. 9)*, volume 9, 1999.

[131] IMDEA Networks Institute. Roderick Fanou New PhD Student. https://www.networks.imdea.org/whats-new/news/2013/roderick-fanou-new-phd-student, October 2013.

[132] IMDEA Networks Institute and NETCOM Research Group. Roderick Fanou New PhD Student. http://netcom.it.uc3m.es/news/2013/roderick-fanou-new-phd-student, October 2013.

[133] IMDEA Networks Institute. How is the Internet Connected in Africa? https://www.networks.imdea.org/whats-new/news/2017/how-internet-connected-africa, August 2017.

[134] International Development Research Center (IDRC) and International Telecommunication Union (ITU). Via Africa: Creating local and regional IXPs to save money and bandwidth. Technical report, 2005.

[135] International Telecommunication Union (ITU). ICT Facts & Figures 2016. https://www.itu.int/en/ITU-D/Statistics/Documents/facts/ICTFactsFigures2013-e.pdf, 2013.

[136] International Telecommunication Union (ITU). ICT Facts & Figures 2016. https://www.itu.int/en/ITU-D/Statistics/Documents/facts/ICTFactsFigures2016.pdf, 2016.

[137] International Telecommunication Union (ITU). Internet Data Portal: ITU Key ICT Indicators for Developed and Developing Countries and the World (Total and Penetration Rates) 2005–2016. https://idp.nz/Global-Rankings/ITU-Key-ICT-Indicators/6mef-ytq6/, June 2016.

[138] International Telecommunication Union (ITU). Measuring the Information Society Report 2016. Technical report, ITU, 2016.

[139] International Telecommunication Union (ITU). Internet World Stats. https://www.itu.int/en/ITU-D/Statistics/Pages/stat/default.aspx, August 2017.

[140] Internet Service Providers' Association (ISPA). Internet Exchange. http://ispa.org.za/inx/, August 2017.

[141] Internet Society (ISOC). African Union (AU) Selects the Internet Society to Support Establishment of Internet Exchange Points across Africa. http://www.internetsociety.org/news/african-unionau-selects-internet-society-support-establishment-internet-exchange-points-across, August 2012.

[142] Internet Society (ISOC). Africa Peering and Interconnection Forum (AfPIF) 2014: Livestream. https://livestream.com/internetsociety/AfPIF2014day3/videos/60571996, August 2014.

[143] Internet Society (ISOC). Africa Peering and Interconnection Forum (AfPIF) 2015: Livestream. https://livestream.com/internetsociety/afpif2015/statuses/103846879, August 2015.

[144] Internet Society (ISOC). Africa Peering and Interconnection Forum (AfPIF) 2016: Livestream. https://livestream.com/accounts/686369/events/6040448/videos/134496198, August 2016.

[145] Internet Society (ISOC). Joint Study Identifies Infrastructure Development as Top Priority for ICT in Africa. https://www.internetsociety.org/news/press-releases/2016/joint-study-identifies-infrastructure-development-as-top-priority-for-ict-in-africa/, February 2016.

[146] Internet Society (ISOC). Global Internet Maps: Global Internet Maps. http://www.internetsociety.org/map/global-internet-report/, August 2017.

[147] Internet Society (ISOC). The Africa Peering and Interconnection Forum (AfPIF). http://www.afpif.org/, 2017.

[148] Internet Speed in Africa. Internet Speed in Africa. https://www.internetspeedinafrica.org/, August 2017.

[149] Internet World Stats. Internet World Stats: Usage and Population Statistics. http://www.internetworldstats.com/stats.htm, August 2017.

[150] ITWeb. Liquid Telecom Deploys New Optical Network in Africa Using Ekinops Long-Haul DWDM Technology. http://www.itweb.co.za/index.php?id=65384, July 2013.

[151] IXPN. Internet Exchange Point of Nigeria. http://ixp.net.ng, August 2017.

[152] Mike Jensen. African Internet Status Report. http://h-net.msu.edu/cgi-bin/logbrowse.pl?trx=vx&list=h-africa&month=0011&week=a&msg=DkgnIo6w/NaHj0hXMIJBrA&user=&pw=, November 2000.

[153] Karlin Josh, Forrest Stephanie, and Rexford Jennifer. Nation-State Routing: Censorship, Wiretapping, and BGP. *CoRR*, abs/0903.3218 (http://arxiv.org/abs/0903.3218), 2009.

[154] K-NET. K-NET. http://www.knetgh.com, August 2017.

[155] Ethan Katz-Bassett, Harsha V Madhyastha, Vijay Kumar Adhikari, Colin Scott, Justine Sherry, Peter Van Wesep, Thomas E Anderson, and Arvind Krishnamurthy. Reverse Traceroute. In *NSDI*, volume 10, pages 219–234, 2010.

[156] Michael Kende and Charles Hurpy. Assessment of the Impact of Internet Exchange Points (IXPs) - Empirical Study of Kenya and Nigeria. *Internet Society (ISOC)*, (59), April 2012.

[157] Michael Kende and Karen Rose. Promoting Local Content Hosting to Develop the Internet Ecosystem. *ISOC Report*, 2015.

[158] Edris 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/, May 2011.

[159] Kai Krause. The True Size of Africa. http://kai.sub.blue/en/africa.html, October 2010.

[160] LACNIC. LACNIC Database. ftp://ftp.lacnic.net/, August 2017.

[161] Anukool Lakhina, John W Byers, Mark Crovella, and Ibrahim Matta. On the geographic location of internet resources. *IEEE Journal on Selected Areas in Communications*, 21(6):934–948, 2003.

[162] Linode. Diagnosing Network Issues with MTR. https://www.linode.com/docs/networking/diagnostics/diagnosing-network-issues-with-mtr, December 2014.

[163] Liquid Telecom. A Network Like no Other. http://liquidtelecom.com/about-us/network-map, August 2017.

[164] Aemen Lodhi, Natalie Larson, Amogh Dhamdhere, Constantine Dovrolis, and K Claffy. Using PeeringDB to Understand the Peering Ecosystem. *ACM SIGCOMM Computer Communication Review*, 2014.

[165] Matthew Luckie. Scamper: a Scalable and Extensible Packet Prober for Active Measurement of the Internet. In *Proceedings of the ACM SIGCOMM Internet measurement Conference (IMC)*, pages 239–245. ACM, 2010.

[166] Matthew Luckie and K Claffy. A Second Look at Detecting Third-Party Addresses in Traceroute Traces with the IP Timestamp Option. In *International Conference on Passive and Active Network Measurement (PAM)*, pages 46–55. Springer, 2014.

[167] Matthew Luckie, Amogh Dhamdhere, David Clark, Bradley Huffaker, and K Claffy. Challenges in Inferring Internet Interdomain congestion. In *Proceedings of the ACM Internet Measurement Conference (IMC)*, pages 15–22. ACM, 2014.

[168] Matthew Luckie, Amogh Dhamdhere, Bradley Huffaker, David Clark, and K Claffy. bdrmap: Inference of borders between IP networks. In *Proceedings of the ACM Internet Measurement Conference (IMC)*, pages 381–396, 2016.

[169] Matthew Luckie, Bradley Huffaker, Amogh Dhamdhere, Vasileios Giotsas, et al. AS Relationships, Customer Cones, and Validation. In ACM, editor, *Proceedings of the ACM Internet Measurement Conference (IMC)*, 2013.

[170] Andra Lutu. Tangled: Analysis of AS-level Interconnection in LAC Region. *2014*.

[171] LyonIX. LyonIX: Infrastructure. http://www.lyonix.net/en/infrastructure-2, 2017.

[172] Beukes-Amiss M. Three-quarters of African Markets Have Telecoms Competition – Even then the Playing Field is not Always Level. http://www.oafrica.com/business/three-quarters-of-african-markets-have-telecoms-competition-even-then-the-playing-field-is-not-always-level/, June 2014.

[173] Richard TB Ma, Dah Ming Chiu, John Lui, Vishal Misra, and Dan Rubenstein. Internet Economics: The use of Shapley value for ISP settlement. *IEEE/ACM Transactions on Networking (TON)*, 18(3):775–787, 2010.

[174] Richard TB Ma, Dah Ming Chiu, John CS Lui, Vishal Misra, and Dan Rubenstein. On cooperative settlement between content, transit, and eyeball internet service providers. *IEEE/ACM Transactions on networking (TON)*, 19(3):802–815, 2011.

[175] Priya Mahadevan, Dmitri Krioukov, Marina Fomenkov, Bradley Huffaker, Xenofontas Dimitropoulos, Amin Vahdat, et al. Lessons from Three Views of the Internet Topology. *arXiv preprint cs/0508033*, 2005.

[176] Ratul Mahajan, Neil Spring, David Wetherall, and Thomas Anderson. Measuring ISP Topologies with Rocketfuel. In *IEEE/ACM Transactions on Networking (TON)*, volume 12, pages 931–943. IEEE, 2004.

[177] Greg Mahlknecht. Greg's Cable Map. http://www.cablemap.info/, August 2017.

[178] Patrick Maigron. Regional Internet Registries Statistics. http://www-public.tem-tsp.eu/~maigron/RIR_Stats/, August 2017.

[179] MainOne. MainOne: Our Network. http://www.mainone.net/our-network/coverage/, August 2017.

[180] MainOne Cable System. MainOne Cable. https://www.mainone.net/our-network/coverage/, August 2017.

[181] James Manyika, Armando Cabral, Lohini Moodley, Yeboah-Amankwah, Suraj Moraje, Michael Chui, Jerry Anthonyrajah, and Acha Leke. Lions Go digital: The Internet's Transformative Potential in Africa. *McKinsey & Company*, (17), November 2013.

[182] Z Morley Mao, Lili Qiu, Jia Wang, and Yin Zhang. On AS-Level Path Inference. In *ACM SIGMETRICS Performance Evaluation Review*, volume 33, pages 339–349, 2005.

[183] Zhuoqing Morley Mao, Jennifer Rexford, Jia Wang, and Randy H Katz. Towards an Accurate AS-level Traceroute tool. In *Proceedings of Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications*, pages 365–378. ACM, 2003.

[184] Pietro Marchetta, Walter de Donato, and Antonio Pescapé. Detecting Third-Party Addresses in Traceroute Traces with IP Timestamp Option. In *International Conference on Passive and Active Network Measurement (PAM)*, pages 21–30. Springer, 2013.

[185] Miriam Marciel, Foivos Michelinakis, Rodérick Fanou, and Pedro J Muñoz-Merino. Enhancements to Google Course Builder: Assessments Visualisation, YouTube Events Collector and Dummy Data Generator. 2013.

[186] Miriam Marciel, Foivos Michelinakis, Rodérick Fanou, and Pedro J Muñoz-Merino. Learning Analytics on GCB (Code). https://github.com/roderickfanou/Learning_analytics_on_GCB, 2013.

[187] MaxMind. GeoIP. https://dev.maxmind.com/geoip/geoip2/geoip2-city-country-csv-databases/, August 2017.

[188] MaxMind. GeoIP accuracy. https://www.maxmind.com/es/geoip2-city-database-accuracy, August 2017.

[189] David Mayer. University of Oregon Route Views Archive Project. routeviews.org, August 2017.

[190] Mwangi Michuki. Gambia IXP, an Oasis in the Desert. http://www.internetsociety.org/blog/africa-bureau/2014/08/gambia-ixp-oasis-desert, August 2014.

[191] MLab. MLab. https://www.measurementlab.net/, September 2017.

[192] Munin. Main repository for munin master. https://github.com/munin-monitoring/munin, 2017.

[193] Michuki Mwangi. Overview of African IXPs. https://www.isoc.org/isoc/conferences/inet/08/docs/inet2008_mwangi.pdf, June 2008.

[194] Michuki Mwangi and Rodérick Fanou. ARDA 1.0: A Pulse Meter for Africa's Peering and Interconnection Landscape. https://www.internetsociety.org/blog/2017/04/arda-1-0-a-pulse-meter-for-africas-peering-and-interconnection-landscape/, April 2017.

[195] Mohamadreza Najiminaini, Laxmi Subedi, and Ljiljana Trajkovic. Analysis of Internet Topologies: a Historical View. In *IEEE International Symposium on Circuits and Systems (ISCAS'09)*, 2009.

[196] Mthuli Ncube and Charles Leyeka Lufumpa. *Infrastructure in Africa: Lessons for Future Development*. Policy Press, 2017.

[197] Network Startup Resource Center (NSRC). Africa: Connectivity information. http://nsrc.org/AFRICA/, August 2017.

[198] Network Startup Resource Center (NSRC). African Undersea and Terrestrial Fibre Optic Cables. https://afterfibre.nsrc.org/, August 2017.

[199] Mark EJ Newman. The Mathematics of Networks. *The New Palgrave Encyclopedia of Economics*, 2(2008):1–12, 2008.

[200] Arnold 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, February 2012.

[201] Rehan Noordally, Xavier Nicolay, Pascal Anelli, Richard Lorion, and Tahiry Razafindralambo. Poor Peering: a Reflexion About a RIXP. *arXiv preprint arXiv:1709.09842*, 2017.

[202] Rehan Noordally, Xavier Nicolay, Pascal Anelli, Richard Lorion, and Pierre Ugo Tournoux. Analysis of Internet Latency: the Reunion Island Case. In *Proceedings of the Asian Internet Engineering Conference*, pages 49–56. ACM, 2016.

[203] Northern African Community (NAC). 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, June 2014.

[204] William B. Norton. *The Internet Peering Playbook: Connecting to the Core*. 2014.

[205] Number Resource Organization (NRO). List of Country Codes and RIRs Ordered by RIR. https://www.nro.net/about-the-nro/list-of-country-codes-and-rirs/, August 2017.

[206] Number Resource Organization (NRO). Stewardship of the IANA Functions. https://www.nro.net/nro-and-internet-governance/iana-oversight/, August 2017.

[207] Towela Nyirenda-Jere and Tesfaye Biru. Internet Development and Internet Governance in Africa. *ISOC Report*, 2015.

[208] Patrick Okui. International Internet Bandwidth and Pricing trends in Africa (Telegeography). https://www.slideshare.net/InternetSociety/international-bandwidth-and-pricing-trends-in-subsahara-africa, August 2016.

[209] Chiara Orsini, Alistair King, Danilo Giordano, Vasileios Giotsas, and Alberto Dainotti. BGPStream: A Software Framework for Live and Historical BGP Data Analysis. In *Proceedings of the ACM Internet Measurement Conference (IMC)*, volume 429 –444. ACM, November 2016.

[210] John S Otto, Mario A Sánchez, John P Rula, and Fabián E Bustamante. Content Delivery and The Natural Evolution of DNS: Remote DNS Trends, Performance Issues and Alternative Solutions. In *Proceedings of the 2012 ACM conference on Internet measurement conference*, pages 523–536. ACM, 2012.

[211] Andrew Owens. Africa: Regional Insights. https://ripe72.ripe.net/presentations/21-Ripe_2016v5_final.pdf, May 2016.

[212] Banji Oyelaran-Oyeyinka and Catherine Nyaki Adeya. Internet access in Africa: empirical evidence from Kenya and Nigeria. *Elsevier Telematics and Informatics*, 21(1):67–81, February 2004.

[213] Packet Clearing House (PCH). Daily Routing Snapshots. https://www.pch.net/resources/Routing_Data/, August 2017.

[214] Packet Clearing House (PCH). Internet Exchange Directory: Download datasets. https://prefix.pch.net/applications/ixpdir/menu_download.php, 2017.

[215] Packet Clearing House (PCH). Internet eXchange Point Directory Reports. https://prefix.pch.net/applications/ixpdir/summary/, August 2017.

[216] Packet Clearing House (PCH). Internet Exchange Point Growth. https://prefix.pch.net/applications/ixpdir/summary/growth/, August 2017.

[217] Packet Clearing House (PCH). IXP growth per region. https://prefix.pch.net/applications/ixpdir/summary/growth-region/, August 2017.

[218] Packet Clearing House (PCH). PCH IXP directory. http://prefix.pch.net/images/applications/ixpdir/ip_asn_mapping.txt, August 2017.

[219] Packet Clearing House (PCH). Research. https://www.pch.net/about/research, August 2017.

[220] PeeringDB. http://www.peeringdb.com/private/exchange_list.php, August 2017.

[221] Björn Pehrson and Margaret Ngwira. Optical Fibre for Education and Research Networks in Eastern and Southern Africa. http://www.sarua.org/files/publications/Sarua_Fibre_Report_2006.pdf, March 2006.

[222] Veljko Pejovic, David Johnson, Mariya Zheleva, Elizabeth Belding, Lisa Parks, and Gertjan Van Stam. The Bandwidth Divide: Obstacles to Efficient Broadband Adoption in Rural Sub-Saharan Africa. *International Journal of Communication*, 6:2467–2491, 2012.

[223] PingER. PingER Project. http://www-iepm.slac.stanford.edu/pinger/, August 2017.

[224] Planetlab. Planetlab. https://www.planet-lab.org/, 2007-2017.

[225] Ingmar Poese, Steve Uhlig, Mohamed Ali Kaafar, Benoit Donnet, and Bamba Gueye. IP Geolocation Databases: Unreliable? *ACM SIGCOMM Computer Communication Review*, 4(2):53–56, 2011.

[226] F Poletti, NV Wheeler, MN Petrovich, N Baddela, E Numkam Fokoua, JR Hayes, DR Gray, Z Li, R Slavík, and DJ Richardson. Towards high-capacity fibre-optic communications at the speed of light in vacuum. *Nature Photonics*, 7(4):279–284, 2013.

[227] PriMetrica. TeleGeography Internet Exchange Map. http://www.internetexchangemap.com/, August 2017.

[228] Puppet. Download Open Source Puppet. https://puppet.com/download-open-source-puppet, 2017.

[229] Puppet. Dozens of Projects, Used by Thousands of Organizations. https://puppet.com/products/open-source-projects, 2017.

[230] Puppet. Puppet: Server Automation Framework and Application. https://github.com/puppetlabs/puppet, 2017.

[231] 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.

[232] Bruno Quoitin, Cristel Pelsser, Olivier Bonaventure, and Steve Uhlig. A performance evaluation of BGP-based traffic engineering. *International journal of network management*, 15(3):177–191, 2005.

[233] Bruno Quoitin and S Tandiel. C-BGP user's guide. *CSE Departement, UCL, Belgique*, 2004.

[234] Bruno Quoitin and Steve Uhlig. Modeling the Routing of an Autonomous System with C-BGP. *IEEE Network Magazine*, 19(6):12–19, 2005.

[235] Rajiv Ramaswami, Kumar Sivarajan, and Galen Sasaki. *Optical Networks: a Practical Perspective*. Morgan Kaufmann, 2009.

[236] Vicky Ramirez. 8 Maps That Will Change the Way You Look at Africa. https://www.one.org/international/blog/8-maps-that-will-change-the-way-you-look-at-africa/, April 2014.

[237] Martin Ravallion. *The Economics of Poverty: History, Measurement, and Policy*. Oxford University Press, 2015.

[238] Yakov Rekhter, Tony Li, and Susan Hares. A Border Gateway Protocol 4 (BGP-4). Technical report, 2005.

[239] Reporters Without Borders. Media obstructed during Chad's presidential election. https://rsf.org/en/news/media-obstructed-during-chads-presidential-election, April 2016.

[240] 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, March 2016.

[241] Jennifer Rexford, Jia Wang, Zhen Xiao, and Yin Zhang. BGP Routing Stability of Popular Destinations. In *Proceedings of the 2nd ACM SIGCOMM Workshop on Internet measurement*, pages 197–202. ACM, Nov. 2002.

[242] Philipp Richter, Mark Allman, Randy Bush, and Vern Paxson. A Primer on IPv4 Scarcity. *ACM SIGCOMM Computer Communication Review*, 45(2):21–31, 2015.

[243] Philipp Richter, Georgios Smaragdakis, Anja Feldmann, Nikolaos Chatzis, Jan Boettger, and Walter Willinger. Peering at Peerings: On the Role of IXP Route Servers. In *Proceedings of the 2014 Conference on Internet Measurement Conference*, pages 31–44. ACM, 2014.

[244] RINEX. RINEX: Rwanda Internet eXchange. http://www.rinex.org.rw/, 2015.

[245] RIPE NCC. OpenIPMap Database. https://labs.ripe.net/Members/emileaben/infrastructure-geolocation-plan-of-action, December 2013.

[246] RIPE NCC. RIPE70: Daily Meeting Report. https://ripe70.ripe.net/programme/report/, May 2015.

[247] RIPE NCC. The RIPE Academic Cooperation Initiative (RACI). https://ripe70.ripe.net/programme/raci/, May 2015.

[248] RIPE NCC. Global RIPE Atlas Network Coverage. https://atlas.ripe.net/results/maps/network-coverage/, August 2017.

[249] RIPE NCC. RIPE Atlas - Raw Data Structure Documentation. https://atlas.ripe.net/docs/data_struct/, August 2017.

[250] RIPE NCC. RIPE Atlas Probe V2. https://atlas.ripe.net/docs/probe-v2/, September 2017.

[251] RIPE NCC. RIPE Atlas Resources: Code. https://atlas.ripe.net/resources/code/, August 2017.

[252] RIPE NCC. RIPE NCC Database. ftp://ftp.ripe.net/, August 2017.

[253] RIPE NCC. RIPE RIS. https://www.ripe.net/analyse/internet-measurements/routing-information-service-ris/, August 2017.

[254] RIPE NCC. RIPE Stats. https://stat.ripe.net/, 2017.

[255] Samknows. Global Internet Performance Platform. https://samknows.com/global-platform, September 2017.

[256] Víctor Sánchez-Agüero, Rodérick Fanou, Pierre Francois, and Francisco Valera. African Measurements Campaigns (AMC). http://amc.netcom.it.uc3m.es/, October 2017.

[257] Bijal Sanghani. 2012 Report on European IXPs. Technical report, European Internet Exchange Association, February 2013.

[258] Robert Schumann and Michael Kende. Lifting Barriers to Internet Development in Africa: Suggestions for Improving Connectivity. *Analysys Mason Limited, London*, May 2013.

[259] Seacom. Seacom Fiber Network. http://seacom.mu/network-map/, June 2017.

[260] Anuj Sehgal and Jürgen Schönwälder. Getting Familiar with the C-BGP Simulator. http://cnds.eecs.jacobs-university.de/courses/anl-2010/C-BGP%20Tutorial.pdf, September 2010.

[261] Serekunda Internet Exchange Point (SIXP). http://www.sixp.gm/, August 2017.

[262] J. Sherry, E. Katz-Bassett, M. Pimenova, H. V. Madhyastha, T. Anderson, and A. Krishnamurthy. Resolving IP aliases With Prespecified Timestamps. In *Proceedings of the ACM SIGCOMM Internet Measurement Conference (IMC)*, pages 172–178. ACM, 2010.

[263] Philip Smith. BGP Routing Table Analysis. Technical report, PFS Internet Development Pty Ltd, August 2017.

[264] Steve Song. AfTerFibre: Mapping Terrestrial Fibre Optic Cable Projects in Africa. https://manypossibilities.net/2011/06/afterfibre-mapping-terrestrial-fibre-optic-cable-projects-in-africa/, June 2011.

[265] Steve Song. AfTerFiber Mapping African Terrestrial Fibre Optic Infrastructure. https://www.flickr.com/photos/ssong/albums/72157627195113720, January 2015.

[266] Steve Song. Africa Telecoms Infrastructure in 2015. https://manypossibilities.net/2016/01/africa-telecoms-infrastructure-in-2015/, January 2016.

[267] Steve Song. African Undersea Cables. https://manypossibilities.net/african-undersea-cables/, August 2017.

[268] South African Development Community (SADC). Conclusions of the AU-SADC Regional Internet Exchange Point (RIXP) and Regional Internet Carrier (RIC) Workshop 03-07 February 2014. Technical report, February 2014.

[269] Speedchecker. Speedchecker. http://www.speedchecker.xyz/, September 2017.

[270] Kara Sprague, F Grijpink, J Manyika, L Moodley, B Chappuis, K Pattabiraman, and J Bughin. Offline and Falling Behind: Barriers to Internet Adoption. *McKinsey & Company, Technical Report*, September 2014.

[271] Richard Steenbergen. Traceroute. http://cluepon.net/ras/traceroute.pdf, 2011.

[272] John W Stewart III. *BGP4: Inter-domain Routing in the Internet*. Addison-Wesley Longman Publishing Co., Inc., 1998.

[273] Volker Stocker, Georgios Smaragdakis, William Lehr, and Steven Bauer. The growing complexity of content delivery networks: Challenges and implications for the Internet ecosystem. *Elsevier Telecommunications Policy*, 2017.

[274] Florian Streibelt, Jan Böttger, Nikolaos Chatzis, Georgios Smaragdakis, and Anja Feldmann. Exploring EDNS-client-subnet Adopters in Your Free Time. In *Proceedings of the 2013 conference on Internet measurement conference*, pages 305–312. ACM, 2013.

[275] William Stucke. Regional IXPs: the Need for Regional Interconnection in Africa. *Information & communications technologies*, 2006.

[276] Ao-Jan Su, David R Choffnes, Aleksandar Kuzmanovic, and Fabián E Bustamante. Drafting Behind Akamai (Travelocity-Based Detouring). In *ACM SIGCOMM Computer Communication Review*, volume 36, pages 435–446. ACM, 2006.

[277] Submarine Telecoms Forum. Submarine Telecoms Industry Report 2013. Technical Report 2, Terabit Consulting, March 2013.

[278] Submarine Telecoms Forum. Submarine Telecoms Industry Report 2014. Technical Report 3, Terabit Consulting, 2014.

[279] Submarine Telecoms Forum. Submarine Telecoms Industry Report 2016 (5th Anniversary Edition). Technical report, Terabit Consulting, October 2016.

[280] Submarine Telecoms Forum, Inc. Submarine Telecoms Industry Report 2012. Technical Report 1, Terabit Consulting, July 2012.

[281] Lakshminarayanan Subramanian, Sharad Agarwal, Jennifer Rexford, and Randy H. Katz. Characterizing the Internet Hierarchy from Multiple Vantage Points. In *INFOCOM 2002. Twenty-First Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings. IEEE*, volume 2, pages 618–627. IEEE, 2002.

[282] Srikanth Sundaresan, Amogh Dhamdhere, Mark Allman, and K Claffy. TCP Congestion Signatures. In *Proceedings of the ACM Internet Measurement Conference (IMC)*, 2017.

[283] Srikanth Sundaresan, Danny Lee, Xiaohong Deng, Yun Feng, and Amogh Dhamdhere. Challenges in Inferring Internet Congestion using Throughput Measurements. In *Proceedings of the ACM Internet Measurement Conference (IMC)*, 2017.

[284] Tanzania Internet Service Providers Association (TISPA). Tanzania Internet eXchange – TIX. https://www.tix.or.tz/, August 2017.

[285] Wayne A Taylor. Change-Point Analysis: A powerful New Tool for Detecting Changes. http://www.variation.com/cpa/tech/changepoint.html, 2000.

[286] Team Cymru. Team Cymru Services. https://www.team-cymru.com/, August 2017.

[287] Telecommunications Service Providers of Kenya (TESPOK). Kenya Internet Exchange Point (KIXP). https://www.tespok.co.ke/, August 2017.

[288] Telegeography. Submarine Cable Map 2017. http://submarine-cable-map-2017.telegeography.com/, September 2017.

[289] Telegeography. Telegeography Submarine Cable Map. https://www.submarinecablemap.com/, September 2017.

[290] Dave Thaler. Evolution of the IP Model. 2011.

[291] The African IXP Association (Af-IX). Survey of the African IXPs operators. http://af-ix.net/pipermail/af-ix-discuss_af-ix.net/2016-February/000006.html, February 2016.

[292] The African IXP Association (Af-IX). List of Active Internet eXchange Points in Africa. http://www.af-ix.net/ixps-list, August 2017.

[293] The History Guy. Current Wars in Africa. http://www.historyguy.com/wars_of_africa_current.html, November 2016.

[294] The World Bank. The World Bank: Data By Topic. http://data.worldbank.org/topic, August 2017.

[295] Caroline J Tolbert and Ramona S McNeal. Unraveling the Effects of the Internet on Political Participation? *Political research quarterly*, 56(2):175–185, 2003.

[296] Paul Traina, Danny McPherson, and John Scudder. Autonomous system confederations for BGP. Technical report, 2007.

[297] United Nations Economic Commission for Africa. Regional Economic Communities. http://www.uneca.org/oria/pages/history-background-africas-regional-integration-efforts, June 2017.

[298] Jacobson Van. Traceroute. ftp://ftp.ee.lbl.gov/traceroute.tar.gz, February 1989.

[299] WACREN. West and central african research and education network english français twitter iconfacebook iconflickr iconrss icon main menu west and central african research and education network (wacren). http://www.wacren.net/, September 2017.

[300] Matthias Wählisch, Thomas Schmidt, Markus de Brün, and Thomas Häberlen. Exposing a Nation-centric View on the German Internet–a Change in Perspective on AS-level. In *International Conference on Passive and Active Network Measurement (PAM)*, pages 200–210. Springer, 2012.

[301] James Wan. Why Google Maps Gets Africa wrong? https://www.theguardian.com/world/2014/apr/02/google-maps-gets-africa-wrong, April 2014.

[302] Wikipedia. List of active rebel groups. https://en.wikipedia.org/wiki/List_of_active_rebel_groups, August 2017.

[303] Wikipedia. List of conflicts in Africa. https://en.wikipedia.org/wiki/List_of_conflicts_in_Africa, August 2017.

[304] Wikipedia. List of non-state terrorist incidents. https://en.wikipedia.org/wiki/List_of_terrorist_incidents, August 2017.

[305] Wikipedia. National research and education network. https://en.wikipedia.org/wiki/National_research_and_education_network#, September 2017.

[306] Wikipedia. Regions of the African Union. https://en.wikipedia.org/wiki/Regions_of_the_African_Union, August 2017.

[307] Mark Winther. Tier 1 ISPs: What They Are and Why They Are Important. *IDC White Paper, NTT Communications*, May 2006.

[308] Mühlbauer Wolfgang, Feldmann Anja, Olaf Maennel, Matthew Roughan, and Steve Uhlig. Building an AS-Topology Model that Captures Route Diversity. In *SIGCOMM*, 2006.

[309] Woodrow Wilson International Center for Scholars - Africa Program. African Regional and Sub-Regional Organizations: Assessing Their Contributions to Economic Integration and Conflict Management. Technical report, October 2008.

[310] World Bank. Africa. https://www.worlddata.info/africa/index.php, June 2017.

[311] World Bank. Average Income Around the World. https://www.worlddata.info/average-income.php, June 2017.

[312] Worldometers. Countries in the World by Population. http://www.worldometers.info/world-population/population-by-country/, August 2017.

[313] Marcelo Yannuzzi, Xavi Masip-Bruin, Eduardo Grampín, Roque Gagliano, Alberto Castro, and Martín Germán. Managing Interdomain Traffic in Latin America: A New Perspective Based on LISP. *IEEE Communications Magazine*, 47(7), 2009.

[314] Marcelo Yannuzzi, Xavier Masip-Bruin, and Olivier Bonaventure. Open Issues in Interdomain Routing: a Survey. In *IEEE network*, pages 49–56, 2005.

[315] Yasir Zaki, Jay Chen, Thomas Pötsch, and Talal Ahmad Lakshminarayanan Subramanian. Dissecting Web Latency in Ghana. In *Proceedings of the ACM Internet Measurement Conference (IMC)*, pages 241–248. ACM, 2014.

[316] M. Zennaro, E. Canessa, K. R. Sreenivasan, A. A. Rehmatullah, and R. L. Cottrell. Scientific Measure of Africa's Connectivity. *Information Technologies & International Development*, (1):55–64, 2006.