Open access • Journal Article • DOI:10.1145/3211852.3211857

Open Connect Everywhere: A Glimpse at the Internet Ecosystem through the Lens of the Netflix CDN — Source link

Timm Böttger, Felix Cuadrado, Gareth Tyson, Ignacio Castro ...+1 more authors

Institutions: Queen Mary University of London

Published on: 27 Apr 2018 - ACM Special Interest Group on Data Communication

Topics: Internet exchange point, Content delivery network and The Internet

Open Connect Everywhere: A Glimpse at the Internet Ecosystem through the Lens of the Netflix CDN

Timm Boettger timm.boettger@qmul.ac.uk
Felix Cuadrado felix.cuadrado@qmul.ac.uk
Gareth Tyson gareth.tyson@qmul.ac.uk
Ignacio Castro i.castro@qmul.ac.uk
Steve Uhlig steve.uhlig@qmul.ac.uk
Queen Mary University of London

ABSTRACT

Netflix has become a worldwide video-streaming platform and the source of a large amount of the Internet traffic. It has achieved this without building its own datacentres, by controlling a network of servers deployed at Internet eXchange Points (IXPs) and within Internet Service Providers (ISPs). Despite its wide success and novel approach to infrastructure deployment, we have little understanding of its deployments and operations. Through extensive measurements, we reveal Netflix's footprint, its traffic patterns and volumes. Our analysis of Netflix's server deployment exposes the diversity of the Internet ecosystem world-wide. Our findings also suggest that the specifics of each regional ecosystem, coupled with the demand of each local market, explain the different deployment strategies.

1. INTRODUCTION

What began as a DVD-rental-by-mail company in 1997 has become one of the largest video streaming companies in the world. Netflix serves nowadays more than 75 million paying customers, is available in almost every country[1], and delivers a large amount of traffic. Indeed, more than a third of the US downstream peak traffic is attributed to Netflix [11].

To meet the challenge of global high-quality video delivery, Netflix has built a dedicated Content Delivery Network (CDN). Netflix's CDN pursues an original strategy that combines the advantages of both centralised (e.g., Limelight) and decentralised models (e.g., Akamai). First, Netflix appears to exploit mature Internet Exchange Points (IXPs), i.e., those with rich enough ecosystems. These IXP locations are complemented by deployments across multiple Internet Service Providers (ISPs). This combination allows Netflix to fine-tune its various regional deployments to best match the local operating needs.

In this paper, we explore the Neflix CDN and observe its deployment approach. We not only gain vantage on Netflix itself, but also on the complex and regionally heterogeneous inter-domain ecosystem. To this end, we have launched a measurement campaign and collected data across the whole Netflix infrastructure. By exploiting the structured naming approach followed by Netflix, we compiled a comprehensive list of Netflix's servers. Our approach avoids the limitations of infrastructure sampling approaches that exploit DNS redirection, e.g., [16, 18, 14]. We also launch active probes towards those servers to study their traffic volume.

Our results provide two different perspectives. First, we obtain a global view, exploring the scale of the Netflix deployment in terms of its servers and traffic volumes. We confirm that Netflix has built an impressive and well-engineered infrastructure that covers every continent. We contend that its deployment provides an additional and insightful reference point for future CDNs. Second, the different deployments across regions exhibit a pattern, based mostly on IXP locations, which are complemented by within ISP deployments. This apparent pattern confirms the central role of IXPs in the Internet ecosystem, and how useful they can be to large CDNs.

In this paper we make the following contributions:

1. We unveil the server deployment of the Netflix CDN, showing its world-wide footprint. We find servers deployed at 233 locations across 6 continents.

2. We expose the region-specific deployment of the Netflix CDN. We explore how Netflix composes its CDN of servers strategically placed at IXPs and within ISPs. We demonstrate how the strategy of Netflix has led to rather different footprints in each region.

3. We also show Netflix's traffic volumes and its dynamics. We study the traffic delivered by the servers in each region, highlighting how the reliance on IXP and ISP servers differs. Our estimates confirm the relative importance of the various regions provided by the sheer server deployment, with the

[1]China and Russia are notable exceptions.

USA still accounting for a vast majority of the traffic, unexpectedly followed by Mexico, UK, Canada, and Brazil.

Our story line unfolds in the following way. We first provide relevant background and describe related work in Section 2. The bulk of the paper related to the Netflix CDN deployment comes next in Section 3. We then study the traffic of Netflix in Section 4, and conclude in Section 5.

2. BACKGROUND & RELATED WORK

2.1 The evolving Internet Ecosystem

The traditional mental model of the Internet ecosystem was shaken by the seminal work from Labovitz et al. [28]. This large study of the Internet inter-domain traffic evidenced the significant impact of large players, such as Google, on the Internet ecosystem. This study provided an estimate on the huge amount of traffic that these players serve to end-users, and the corresponding flattening of the Internet. One of the major players supporting this flattening of the Internet are IXPs, which are known to interconnect the Internet ecosystem at many locations across the world. IXPs facilitate the direct interconnection between content players and ISPs. For example, through the study of one of the largest European IXPs, Ager at al. [15] uncovered a rich and varied ecosystem. The ecosystem present at large IXPs is so large that it fundamentally questions our current knowledge of the AS-level topology.

Despite their importance in the Internet ecosystem, only a few studies have targeted IXPs [17, 15, 19, 20, 21, 22]. The work from Augustin et al. [17] aimed at systematically mapping IXP infrastructures through large-scale active measurements, leading to the first evidence of the large number of IXPs around the world. Ager et al. [15] studied the ecosystem and traffic of one of the largest European IXPs, while Restrepo et al. [19] looked at two smaller European IXPs. Subsequent studies from Chatzis et al. [20, 21, 22] reinforced the critical role played by IXPs in the Internet ecosystem. Our paper, though focused on the Netflix CDN, further reinforces the strategic nature of IXPs for the footprint of a large-scale video delivery platform.

2.2 Large-scale content delivery

To cope with the increasing demand for content, CDNs deploy massively distributed server infrastructures [29]. Two main deployment approaches have been traditionally used by large CDNs: (1) placing servers deep inside the network at many locations, as done by Akamai [23], or (2) concentrating resources in a few selected locations, i.e., datacentres, as done by Google, Microsoft, Amazon, and Limelight. The first approach is naturally fit to web content, due to its latency-sensitive nature. Video, on the other hand, is mostly served through the second approach, due to its storage and bandwidth requirements.

The work from Torres et al. [30] explored Youtube's server selection policies, and showed that content was indeed delivered from a small set of datacentres deployed across multiple continents. Being an Internet giant with a user base of hundreds of millions, Google has inspired researchers to explore the related challenges in delivering services at scale. Google manages a complex cloud infrastructure and delivers various services, mainly from its datacentres interconnected through an internal backbone network [27]. Mapping the infrastructure of such a company is a moving target, as the infrastructure evolves to meet the needs of the users. Calder et al. performed EDNS-based measurements to estimate the footprint of Google's CDN infrastructure [18], and found evidence for a substantial expansion. Google has currently deployed thousands of caches across more than 1000 ASes world-wide. Interestingly, at the time of these studies, the majority of requests were still addressed by the infrastructure deployed inside Google, i.e., its datacentres. As we show in this paper, the server deployment of Netflix shares similarities to this approach.

In contrast to the pure digital nature of Google, Netflix started as a physical video company. The transformation of Netflix into a digital content delivery company has heavily relied on Amazon Web Services, rather than building its own datacentres as done by Google. A single piece of research studied the Netflix content delivery infrastructure, by Adhikari et al. in 2015 [14]. This measurement study explored how Netflix and Hulu redirect users to specific CDN servers. Their findings show that, at the time of those measurements, Netflix relied on three external CDN partners (Akamai, LimeLight, and Level3) for the actual video delivery. However, that approach is no longer in place and therefore past observations about Netflix do not hold any longer. Nowadays Netflix exclusively uses its own Open Connect CDN for video delivery [2]. Furthermore, since the previous Netflix studies, its market has vastly expanded from a handful of countries into a global service. In the next section, we describe the current approach used by the Netflix CDN.

2.3 The Netflix CDN

Netflix exclusively uses its own CDN: Netflix Open Connect [2]. To satisfy bandwidth requirements, two options are publicly offered to network providers [9]: peering at an IXP and deploying servers inside ISPs. Similar to many other network operators, Netflix offers to deliver video through settlement-free peerings at IXPs (referred to as Settlement Free Interconnect (SFI) by Netflix). Netflix follows an open peering policy. The

content is delivered from servers operated by Netflix at IXPs, and each peering partner has to bear its own cost to reach the exchange point.

Netflix also offers to deploy video delivery servers to networks which have a large Netflix customer base. Those servers, called Open Connect Appliance (OCA), are offered by Netflix free of charge. By deploying servers in its network, the ISP reduces the cost of traffic. Netflix offers the ISP fine grained control over which prefixes are served by which servers. Advanced setups, like filling one server from another server within the ISP network or fail-over scenarios are also supported. In the following, we investigate the different roles these two deployment options play.

3. SERVER DEPLOYMENT

In this section we discuss the observed geographical footprint of the Netflix CDN infrastructure. We start by describing our measurement approach, building on assumptions regarding how Netflix built its infrastructure (Section 3.1). This description is followed by a validation of the methodology (Section 3.2). The bulk of this section follows, with an analysis of the server deployment, first at the global world-wide level (Section 3.3), then specifically across the various regions where Netflix is present: North America (Section 3.3.1), Europe (Section 3.3.2), South America (Section 3.3.3), Asia & Oceania (Section 3.3.4), and Africa (Section 3.3.5). We end this section with a discussion on the future of Netflix's infrastructure deployment (Section 3.4).

3.1 Methodology

In a first step, by manually requesting Netflix videos and inspecting the traffic, we inferred how the Netflix platform behaves. We enhanced this process using Selenium [12] to automate browser actions and BrowserMob Proxy [1] to simultaneously capture all HTTP(S) requests. To increase our geographic coverage, we relied on the Hola browser extension [4], which is a free P2P-based VPN plugin, giving access to vantage points in different networks around the world. We tried to sample the Netflix infrastructure from as many different vantage points as possible, given the footprint of Hola. Hola allowed us to use 753 different IPs in 94 ASes. Our measurement study followed the terms of service agreed by Hola users (Hola is a browser plugin for the exact purpose of pretending to be at a different network location). Netflix is officially available in all countries we targeted, and as we requested a short video sequence only, our measurements had a negligible impact on the cooperating users.

Closer investigation of the DNS names of Netflix video servers in our initial measurements revealed a deterministic naming structure. Figure 1 lists representative examples of the server names.

```
ipv4_1-lagg0-c020.1.lhr001.ix.nflxvideo.net
ipv6_1-lagg0-c002.1.lhr005.bt.isp.nflxvideo.net
```

Figure 1: Examples of Netflix server names.

We conjecture that the individual components of each server name are the following:

`ipv4/ipv6`: IP protocol version.

`lagg0`: Type of network connection. We also found samples hinting at different NICs (i.e., cxgbe0, ixl0 or mlx5en0).

`c020`: Counter enumerating servers at a given location.

`lhr001`: Airport code of server location, with counter.

`bt.isp`: Server operated by an ISP, e.g., BT in this case.

`ix`: Server operated by Netflix at an IXP.

`nflxvideo.net`: Common domain for all servers.

The server naming scheme is consistent with network operators' best practices, e.g., using airport codes to denote locations [32].

CDNs typically rely on DNS (CNAME chaining) to direct clients to nearby servers [18, 26], often based on the estimated client location. Our observations of domain names suggest that Netflix uses a different redirection strategy for its CDN. Indeed, Netflix server domain names are too specific, i.e., including details on the physical location of the server. This makes it unlikely that redirection takes place based on these domain names. Additionally, we ran a distributed DNS lookup campaign from Planetlab nodes to verify our assumption that Netflix DNS names always resolve to a single IP address, independently of the client location.

To obtain a more complete picture of the Netflix infrastructure, we systematically exploited the structured nature of the server names. We implemented a crawler, which generated DNS names following the pattern outlined above and attempted to resolve those generated domains names to IP addresses. The crawler was fed with lists of known countries, airport codes and ISPs. We curated these lists manually, consulting publicly available lists of countries and airport codes. For the ISPs, we conducted an extensive search on the Internet, including but not limited to the Netflix ISP Speed Index, Wikipedia, various ISP rankings or recommendation lists and discussions in web forums. In total, our curated lists contained 243 countries, 3,468 airport codes[2] and 1,728 ISP names. In total we tested 16 dif-

[2]The DNS names actually include a typo, the Carrasco International Airport in Uruguay is abbreviated MDV instead of MVD. We used the ISP names found to verify that what was really meant is the airport in Uruguay.

ferent network connection types. We provide details on how many servers were found and how our two approaches (using Hola and the crawler[3]) performed in Table 1.

		Hola	Crawler
Servers found	ISP	157	1,428
	IX	379	3,241
	total	536	4,669
Locations found	ISP	54	217
	IX	26	39
	total	80	256
Countries found	ISP	32	39
	IX	16	17
	total	48	56
ISPs found		61	120

Table 1: Initial Results and Comparison of the Hola vs. Crawler approach.

Table 1 summarises the number of servers, locations, and countries across ISPs and IXPs found by the Hola and crawler-based approaches. The crawler-based approach clearly outperforms the Hola-based approach, finding every server identified through Hola. This happens despite the significant country-level footprint of the Hola infrastructure, and highlights the importance of exploiting the naming structure of the servers.

We are highly confident that our crawling-based approach provides a lower bound of the Netflix CDN, and is representative of the complete infrastructure, especially in terms of countries and locations. In the absence of a definitive ground truth, we believe our results are closest to what would be obtainable through direct measurements.

3.2 Validation

As a first step towards ascertaining the physical location of the servers, we used the MaxMind GeoLite2 [3] database for geolocation. For 84.5% of the servers, MaxMind confirmed the country we expected the server to be in, assuming that the airport code provides a ground truth. We manually investigated those cases where MaxMind maps the IP to a different country than expected. Those errors stem from (1) servers at IXP locations which use a prefix which is mapped to the country where it is registered instead of where it is used (13.3%); (2) cases which appear to relate to ongoing deployments (1.9%); (3) and international network operators with presence in multiple countries (0.16%). In the cases of ongoing deployments the domain names indicate ISPs locations in different countries, whereas the address space still belongs to Netflix.

We further checked our assumption regarding the part of the naming indicating the location of the server inside an ISP or at an IXP, by resolving the server IP address to its AS number using Team Cymru's IP to ASN mapping [7]. For 98.1% of the domain names found, the IP gets mapped to the correct AS, i.e., either to Netflix (inside an IXP) or to the ISP indicated in the domain name. We manually checked the remaining 1.9% and found that all cases appear to relate to ongoing deployment of Netflix servers, as the corresponding address space currently belongs to Netflix.

Finally, we conducted traceroutes from 166 planetlab nodes towards all Netflix servers to gain further confidence on the physical locations of the servers. We used a methodology similar to the one presented by Calder et al. [18]. For each Netflix server we chose (1) the closest, (2) the closest five, (3) the closest ten and (4) closest 25 planetlab nodes based on the measured RTT. We then calculated the geographical distance between the expected location of the Netflix server and the closest (in terms of RTT) planetlab nodes. We observed that nodes which are close from a network perspective (i.e., measured in RTT) are also geographically close, reinforcing our assumption on the validity of the location information embedded in the DNS names.

All in all, these validation steps give us high confidence that the information embedded in the domain names is indeed trustworthy, and can be considered as a ground truth.

3.3 World-wide Server Deployment

In this section, we visualise and discuss the footprint of the world-wide Netflix content delivery infrastructure. We explain the overall picture first, to later delve into the deployments on a per-region basis.

Figure 2(a) shows the locations and deployment sizes for servers operated by Netflix at IXPs. The map shows that Netflix is primarily present in the USA and Europe, consistently with their main customer base [10]. In addition to the USA and Europe, Brazil is also one of the main countries where Netflix servers are deployed at IXPs, and the only country in the Southern American continent, where this is the case. We also observe a sparse server deployment at IXPs in Oceania and Asia, specifically in Japan, Singapore and Hong Kong.

Figure 2(b) is the counterpart of Figure 2(a), but shows the locations and deployment sizes at ISPs. When comparing Figure 2(b) with Figure 2(a), it appears that

[3]Omitting the structured nature of the name, and simply iterating over all possible character sequences is practically infeasible and not desirable. Indeed, assuming an alphabet of 26 characters plus '.', '-', '_' as special characters and a prefix length of 30 characters, enumerating all 29^{30} possible combinations in one year's time would require 2^{36} DNS queries per second.

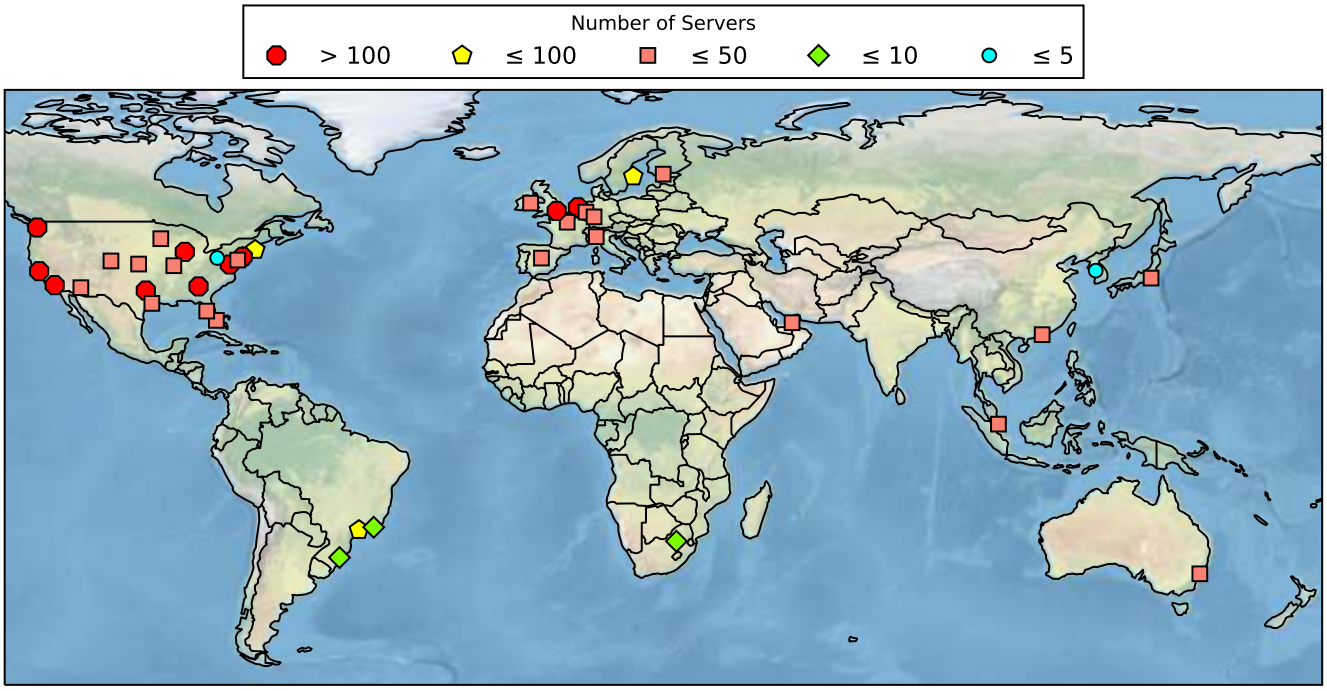

(a) CDN servers operated by Netflix at IXPs.

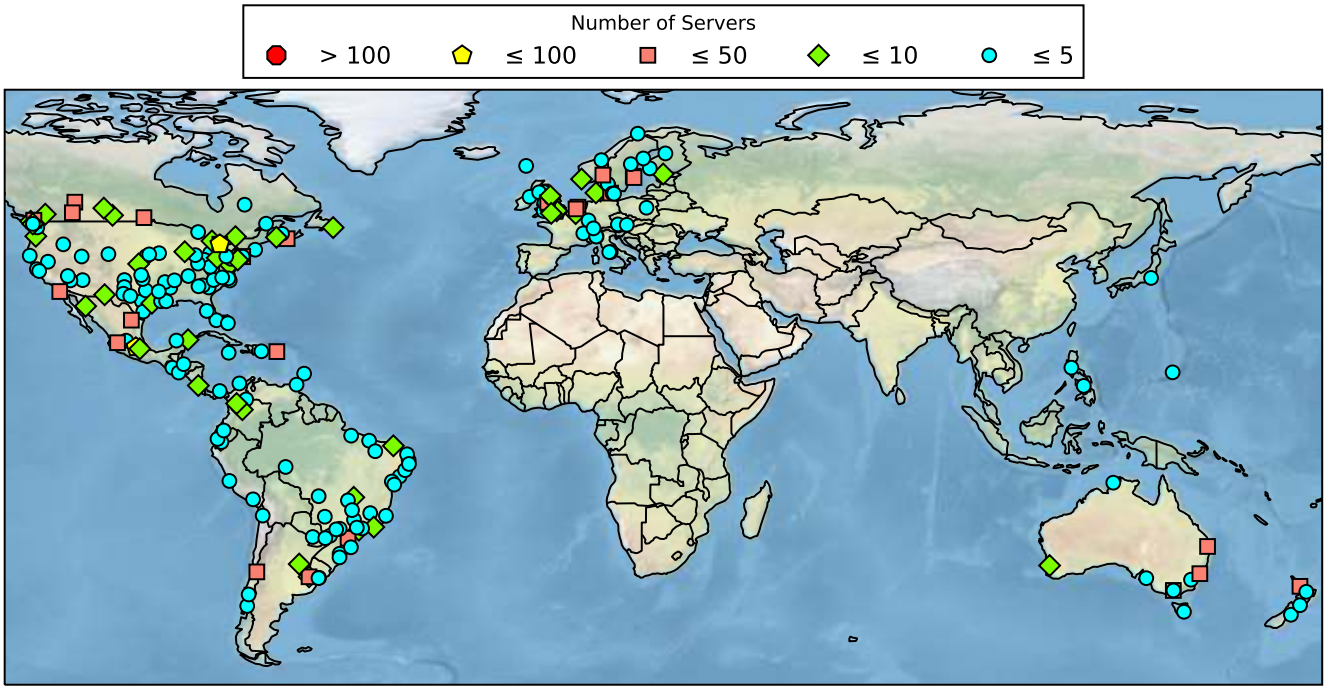

(b) CDN servers deployed within ISPs.

Figure 2: CDN servers deployed by Netflix. The colours and sizes of the markers depict the number of servers deployed at a given location.

the set of servers operated by Netflix at IXPs and those operated within ISPs are to some extent complementary geographically. This seems to be especially the case for regions where Netflix's own IXP deployment is relatively sparse, such as South America, Oceania and the Caribbean Islands. The strength of this complementarity is striking in specific European countries and US states, suggesting that IXP locations are used as a foothold to reach a regional ecosystem, and ISP locations complement the IXP deployment to provide additional coverage when necessary.

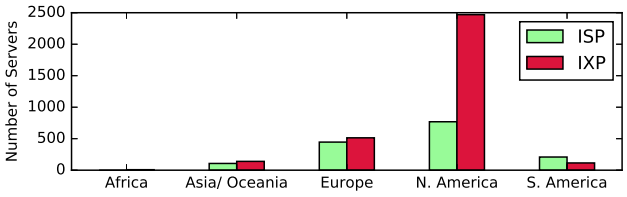

(a) Total number of servers per continent.

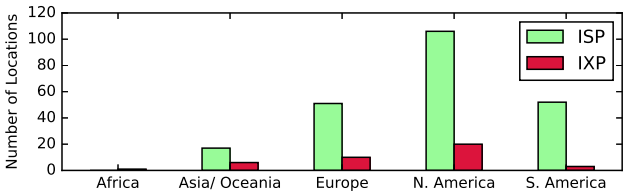

(b) Total number of server locations per continent.

Figure 3: World-wide server deployment.

Figure 3 shows the number of servers and their distribution across locations for each continent. These results again show that most Netflix servers are in the Americas and Europe. Furthermore, the figure emphasises a significant difference in deployment: for IXP servers, each deployment consists of a large number of servers, but at a few selected locations only. ISP servers on the other hand are deployed at many locations, but in contrast every location has only a small number of servers. This again supports our claim about the complementarity of the IXP and ISP server deployments.

Due to how we discover the Netflix deployment, we lack a definitive ground truth to verify precisely what fraction of the whole infrastructure is covered by our measurements. However, we argue that it is very unlikely that we missed significant parts of it. Indeed, Netflix publishes an ISP speed index [8], which ranks the most important ISPs per region according to the perceived end user performance on a monthly basis. We expect, on a per region basis, that this ranking provides information of which networks have Netflix servers. Hence, if an ISP is not significant enough to appear in the Netflix ISP ranking, we expect to see no (or very few) server to be deployed inside this ISP. Therefore, if the Netflix ISP index is highly consistent with our own measurements, this validates our coverage inside ISP networks. Unless stated otherwise, we will use the Speed Index from March 2016 in the following paragraphs.

Table 2 lists the ISPs mentioned in the Netflix Speed Index for the five top markets (according to a 2014 study [10]). It further shows how many of the ISPs we found are deploying servers and gives the number of servers per ISP. We find that a majority of listed ISPs in Canada, the United Kingdom and Brazil deploys Netflix servers. The USA has significantly fewer servers at ISPs relative to the other main markets, given that USA is the major customer base for Netflix, with about 72% of all its subscribers according to [10]. The limited number of servers at ISPs in the USA is a consequence of the major deployment at IXP locations there. This finding is probably among the most surprising ones of our paper.

All in all, we are highly confident that our Netflix sample includes most of the relevant infrastructure parts and especially does not miss any significant location.

3.3.1 *North America*

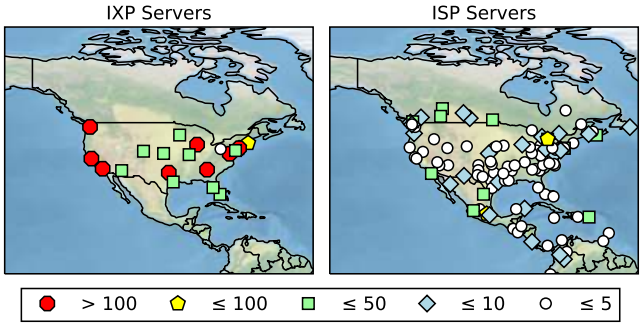

Figure 4: Server deployment in North America, split across IXP and ISP locations. The colours and sizes of the markers again depict the size of the deployments per location.

Being historically the original Netflix market, North America is the region with the densest deployment of Netflix servers world-wide.

Figure 5 shows the distribution of servers and their locations in North America. We observe that an overwhelming majority of the servers in this region is located at IXPs, and exclusively in the USA. From Figure 4, we also see that the 20 IXP locations correspond to the largest cities in the USA, while the locations within ISPs complement the geographic footprint provided by the IXP locations. Figure 4 also shows that the number of servers deployed at each IXP location is much larger than the number of servers deployed at ISP locations.

We find the limited deployment at ISP locations surprising, given that the USA is the main market for Netflix. To discard the possibility of a measurement error,

USA		Canada		United Kingdom		Brazil		Mexico	
AT&T	-	Bell Aliant	22	BT	63	Algar	4	Axtel	10
Bright House	-	Bell Canada	32	EE	-	GVT	43	Totalplay	7
Cablevision	28	Cogeco	14	Plusnet	6	Live TIM	5	izzi	36
CenturyLink	6	Distributel	-	Sky	71	Net Virtua	-	Cablevision Mont.	-
Charter	-	Eastlink	22	TalkTalk	37	Oi Velox	28	Cablemás	-
Comcast	-	MTS	18	Virgin	41	Vivo	-	Telnor	-
Cox	-	Rogers	43					Megacable	2
Frontier	13	SaskTel	12					Telecable	1
Mediacom	-	Shaw	125					Telmex	64
Suddenlink	50	TekSavvy	12						
TWC	-	Telus	54						
Verizon	2	Videotron	-						
Windstream	20								

Table 2: Netflix servers deployed in ISPs within countries with largest number of paying subscribers. ISPs are taken from the Netflix ISP Speed Index, the number next to each ISP denotes the number of servers found within the considered ISP. ISPs listed multiple times in the index (e.g., due to different broadband connection types), are listed only once in this table.

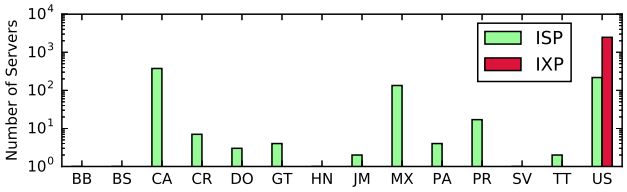

(a) Total number of servers per country.

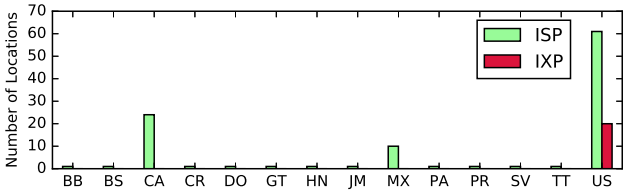

(b) Total number of server locations per country.

Figure 5: Server deployment in North America.

we included all reasonable combinations of these ISP names as input for the DNS crawler. However, even after this extensive search, we could not discover further servers. The absence of deployment inside specific US ISPs (AT&T, Comcast, Time Warner Cable and Verizon), as per Table 2, can be explained by the fact that they publicly refused to deploy Netflix servers and instead insisted on signing paid peering contracts with Netflix [6, 13]. This makes sense given the strong position of these networks in the US market. While the beginning of this dispute reaches back a few years, until today we see no server deployed in those networks, except one specific case[4].

[4]We discovered two Netflix servers in Verizon's network, which does not offer significant advantage in traffic savings for such a large network, but might be part of a trial.

Despite being a neighbour of the USA and one of the biggest Netflix markets, Canada's situation in terms of Netflix deployment is rather different from the one in the USA. In Table 2, we list the ISPs mentioned in Netflix's Speed Index. In contrast to the USA, most ISPs in Canada deploy Netflix servers. We find servers in ten out of twelve Canadian ISPs listed in the Speed Index, the largest Netflix server deployment inside ISPs of any country. As there are no servers operated by Netflix at IXPs in Canada, such wide deployment inside ISPs is therefore expected. Servers are located at 24 locations, mostly distributed across the Southern Canadian border, as shown by Figure 4. This is expected as the majority of Canadians live in that part of the country. We believe that the reasons for the absence of IXP server deployment in Canada are (1) most of the Canadian ISPs can be reached by Netflix through IXPs in the north of the USA, and (2) the fragmented market (in terms of customer coverage) of Canadian IXPs. Every large city in Canada has its own IXP with a limited reach, making the deployment at IXPs there not cost-effective. Therefore, a natural solution for Netflix is to exclusively deploy servers within ISPs to reach the Canadian market.

The Southern neighbour of the USA, Mexico, appears to be in a similar situation to Canada. As per Figure 5, there are quite some servers deployed in Mexico, but all of them within ISPs. Further, we do not see a Netflix server from an IXP located in Mexico. Different from the situation in Canada, where every major city is served by an IXP, in Mexico only Mexico City has two IXPs. These Mexican IXPs are very recent, and therefore their peering ecosystem is likely to be very limited, making the deployment of Netflix servers

there not economically reasonable. Therefore, contrary to Canada where the absence of IXP server deployment likely stems from the fragmented IXP/ISP market, the Mexican situation stems from its immature IXP ecosystem. Further, we notice on Table 2 that some ISPs do not deploy servers. We conjecture that the explanation lies in the fact that Telnor and Telmex are sibling networks, under the same ownership. Further, Cablemas has become part of Izzi in 2015, as did Cablevision in 2014, and therefore these three ISPs are actually one.

Due to their very limited population, other countries in North and Central America show server deployments in ISPs only, with no presence in local IXPs. Server deployments occur in a single location per country, hosting very few servers.

In summary, our findings give evidence that the US IXP ecosystem is strong enough to be used as a stronghold for North America, to reach not only the US market but also Canada, and is complemented by deployment within ISPs to reach Canada and Mexico.

3.3.2 *Europe*

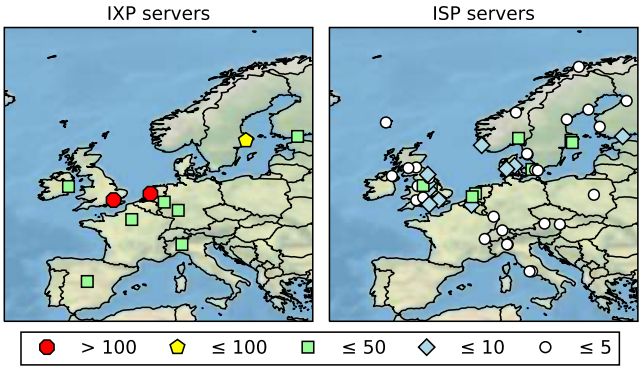

Figure 6: Server deployment in Europe, split across IXP and ISP servers. The colours and sizes of the markers again depict the size of the deployment per location.

The deployment situation in Europe looks quite different from the one in North America. Figure 6 reveals that in Europe, Netflix's deployment is more scattered across countries, and within each country at few locations. The European market is more fragmented due to the language diversity of its customer base, and the different regulations of each country. Netflix approached its European roll-out in multiple stages, addressing a specific set of countries at a time. The Netflix European expansion started in early 2012, with English-speaking countries: the UK and Ireland. The UK has the highest number of servers inside Europe, as shown in Figure 7. The UK deployment relies on a significant number of servers both at IXPs and inside ISPs. The single IXP location (London) is complemented by many ISP locations that provide a much more complete geographical coverage of the country. The second wave of European deployment targeted countries with a strong English proficiency. Nordic countries (Norway, Sweden, Finland, and Denmark) were covered in September 2012, followed by the Netherlands in 2013. We observe a significant Netflix presence in two IXPs, AMS-IX (NL) and Netnod (SE), and a smaller deployment in FINIX (FI). We also notice a diversity of deployments across ISP and IXP locations. While the IXP deployments in Sweden, Finland, and the Netherlands comprise more servers than the total ISP deployments in these countries, the deployment is more balanced than in North America. Norway and Denmark on the other hand operate purely through an ISP server deployment, because both countries' main IXPs are reachable through Netnod anyway.

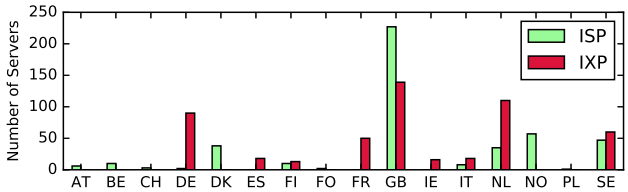

(a) Total number of servers per country.

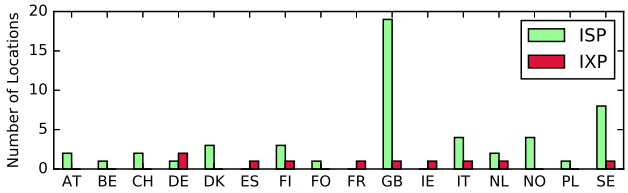

(b) Total number of server locations per country.

Figure 7: Server deployment in Europe.

The next stage of Netflix expansion in 2014 targeted the rest of the Central European market. These countries have a higher language barrier, with customers demanding content in their native languages. The deployment covers two IXPs: DE-CIX (DE) and France-IX (FR). Surprisingly, the ISP deployment consists in very few servers per country in central Europe, and no presence at all in France. These markets seem less mature than those from the first two stages, and we observe that the deployment mostly relies on IXPs. This again supports our claim that Netflix exploits IXP locations whenever possible and usable, which are complemented by ISP deployments. The following stage of Netflix's European expansion, in 2015, included Southern countries, which share many characteristics with Central Europe. Note that Eastern Europe does not currently have any Netflix server deployment, because those countries have only recently been covered by Netflix in its 2016 world-wide expansion. If the demand grows enough in these countries, we expect to observe server deployment

there as well.

3.3.3 South America

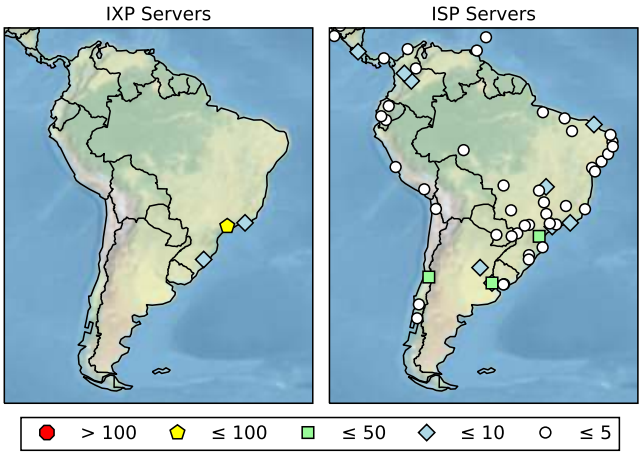

Figure 8: Server deployment in South America.

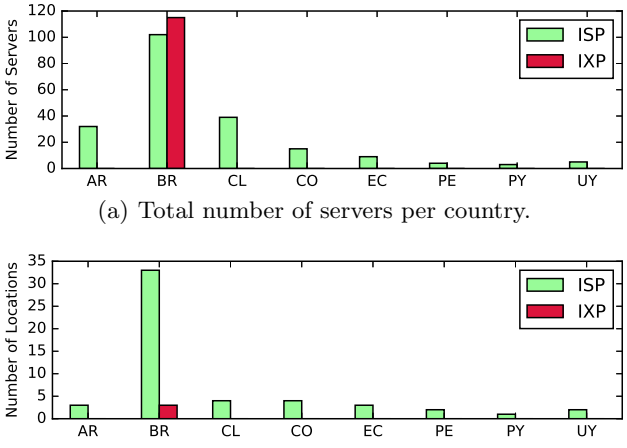

(a) Total number of servers per country.

(b) Total number of server locations per country.

Figure 9: Server deployment in South America.

The map of the server deployment situation in South America is depicted in Figure 8. We notice the absence of IXP servers outside Brazil, as well as the ISP coverage in the Western coastal countries. We also notice the absence of servers in specific countries in the central and northern parts of the region. Figure 9 confirms the visual impression from Figure 8, with Brazil having the largest number of servers deployed in this region, mostly at IXPs, complementing its footprint with more than 30 locations within ISPs. Chile, Argentina, and Colombia follow Brazil both in terms of the number of servers and locations, but exclusively within ISP networks. Despite the existence of IXPs across most South American countries, the limited market for Netflix seems to make it unworthy to deploy inside these IXPs, explaining the exclusive ISP-based deployment outside Brazil.

3.3.4 Asia & Oceania

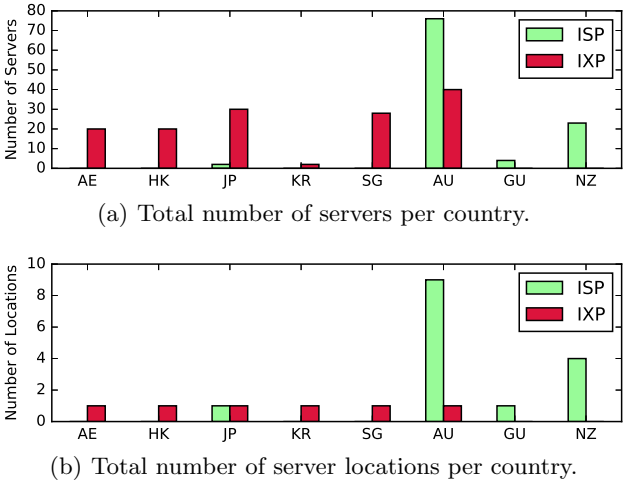

(a) Total number of servers per country.

(b) Total number of server locations per country.

Figure 10: Server deployment in Asia & Oceania.

As can be seen in Figure 10, Asia is mostly served by IXP locations. We find deployments of IXP servers in the United Arab Emirates, Honk Kong, Japan, and in Singapore, at one location per country each. The deployments in the United Arab Emirates and Hong Kong have exactly the same number of servers per location (20). In Japan and Singapore, we observe roughly thirty servers per location. We find two ISP server locations in the Philippines and one in Japan. The reason for not observing IXP servers in the Philippines is that they are currently being deployed, as some of the server names have the "isp" substring while their address space belongs to Netflix, indicating a future IXP deployment. We do not observe servers deployed in China, as Netflix currently is not available there. Overall, the Netflix market in Asia is still very limited, and as the region is geographically large it is more efficiently served through IXPs rather than servers at ISPs.

The server deployment situation in Oceania is shown in Figure 10. In Oceania, we see IXP servers only in Sidney, Australia. These servers have therefore to be complemented by ISP locations in every major city of Australia, Tasmania, and New Zealand. Surprisingly, we also found a few servers deployed in the middle of the Pacific, in Guam. Deploying servers in such a remote location, with a limited population, is explained by the historical nature of Guam as a US (military) stronghold.

3.3.5 Africa

The server deployment in Africa is very limited, representing a tiny market for Netflix. We found eight IXP servers, deployed in South Africa, in Johannesburg. We did not find any other IXP or ISP server in Africa. Given that most African countries are still developing countries, and South Africa is the most developed of

the African countries, this server deployment is sensible. Nevertheless, Africa is, next to Asia, probably the region where the Netflix deployment is most likely to change in the future.

3.4 Discussion

The Netflix infrastructure is clearly engineered and sophisticated, as evidenced by the neatly structured DNS naming scheme, which plays a vital role in redirecting clients to video servers. This structured nature makes it feasible for Netflix to make redirection decisions within their application logic, without the need to rely on DNS. The number of servers deployed follows the relevance of its various markets: the more important a specific market for Netflix, the more servers we observe in the specific region.

Further, we clearly observe that the latest expansion in 130 countries announced by Netflix in January 2016 was only a virtual expansion. Obviously, obtaining the necessary licensing agreements with each local movie industry had to be done for the content to be available in these countries. Interestingly, from the server deployment footprint we do not observe this huge expansion. We observe some additional servers deployed in Internet exchanges in Asia, but many countries where Netflix is officially available still remain without any Netflix infrastructure.

The markets in Africa and Asia are very likely to change and mature in the near future. We expect the Netflix infrastructure to become more dense in these regions, as demand grows and relevant content becomes available. The deployment in Brazil and Australia indicates that these are mature markets. We nevertheless also anticipate further growth there, however not on a scale comparable to Africa or Asia. We do not expect Netflix to increase its hold on the South American markets besides the current main ones (Brazil, Argentina, and Chile), due to the limited size of these markets and the difficulty in reaching end-users in these regions. Still, the main markets for Netflix are clearly North America and Europe. In Europe, we observe as expected a very mature IXP ecosystem, which makes ISP deployment limited or even completely unnecessary when the demand (or the content in the right language) is still limited, e.g., France. Unexpectedly, we found a strong IXP ecosystem in the US, much stronger than anticipated given the existing works from the literature. US IXPs are capable of serving the purpose of a big player like Netflix, and are even used as a base to serve content to Canadian ISPs. This makes ISP deployments in the US unnecessary in many cases, and leads to no IXP deployment in Canada at the moment.

4. TRAFFIC

The reason why Netflix is so interesting is not only its diverse geographical footprint, but also because we expect that it delivers a significant amount of traffic. In this section, we first explain how we obtained traffic measurements (Section 4.1). Then, we provide a validation of the results (Section 4.2). The main part of this section deals with the findings and conclusions derived from the traffic measurements (Section 4.3). First we ascertain the peak and non-peak times across the various locations, before delving into a closer discussion of the actual traffic. We discuss total traffic volumes on a per continent and per country basis and finally compare ISP and IXP deployments on a per server and per location basis.

4.1 Methodology

The identification field (ID) in the IPv4 header was designed to reassemble fragmented IP packets. To reassemble those packets at the destination, this ID must be unique (within certain limitations) for each flow. Most Operating Systems (OS) use a global ID, which is incremented for every packet sent [25] irrespectively of the various flows kept by the end-host. The values of the identification field provide a proxy measurement for the volume of traffic generated by a specific server. However this only holds if the server increments the ID value for each packet sent, instead of using arbitrary or random values, as also allowed per RFC6864 [31].

Fortunately this is the case with Netflix. All Netflix servers run FreeBSD 10 [9]. The two production releases at the time of writing this paper (FreeBSD 10.1 and FreeBSD 10.3), do indeed generate IDs in a predictable way. For every packet sent the ID value is increased by one, independently of the transport layer. In particular, the ID values for both ICMP and TCP packets are derived from the same counter. Accordingly, we estimate the total amount of traffic sent from a server by issuing ICMP ping requests and tracking the evolution of the ID values in the ICMP answers.

We conducted a large scale measurement campaign to assess the temporal behaviour of the ID field of Netflix's servers. From all the servers discovered in the previous section, 4,340 IPv4 addresses (92.9%) responded to our ping requests (see Table 3).

However, overflows of the ID field can happen, specially on busy servers, due to the limited size of the ID field (16 bits). To be able to detect and correct such overruns, we send a ping request every 30ms.

Because the network load induced by simultaneously sampling all the servers would be prohibitive, we sampled each server for one minute in a round-robin fashion using 150 hping3 [5] processes. This led to one measurement per server approximately every 30 minutes. We ran the measurements for a period of 10 days, starting 29/04/16. The traffic load generated by targeting each server twice in an hour is small enough to be negligi-

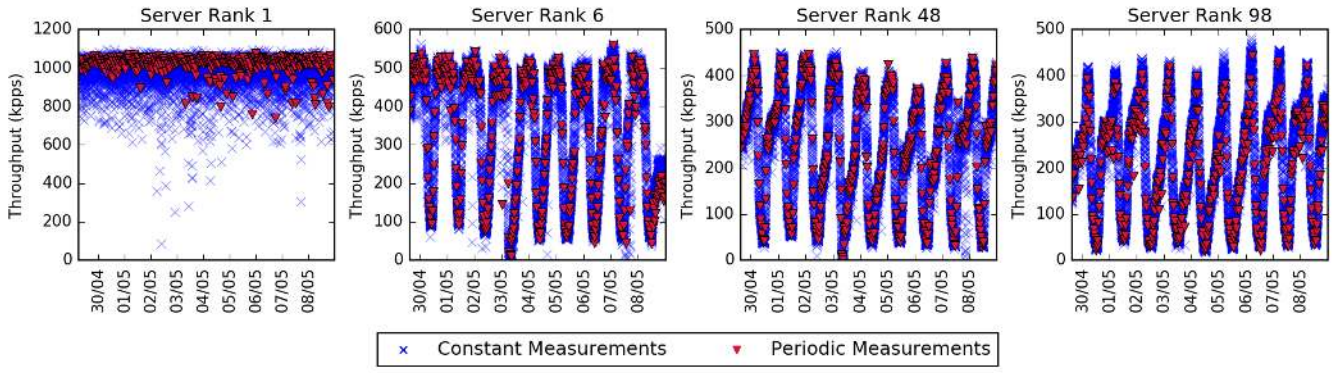

Figure 11: Estimation of the sampling error. Servers are ranked according to maximum throughput.

	reachable			non-reachable		
	IXP	ISP	$\sum$	IXP	ISP	$\sum$
NA	2,356	652	3,008	110	66	176
SA	95	104	199	20	79	99
EU	484	401	885	30	22	52
OC	40	99	139	-	1	1
AS	97	4	101	1	-	1
AF	8	-	8	-	-	-
total	3,080	1,260	4,340	161	168	329

Table 3: Number of servers responding to our hping3 measurements.

ble for Netflix, and hence not disturbing to its regular business activities.

In the next subsection, we present the results and give further evidence that our measurements are correct.

4.2 Validation

We validate our periodic sampling methodology by comparing its results with a constant sampling for a set of selected servers. Figure 11 shows a close match for both, the periodic and constant measurements, thus ensuring that our approach does not miss important temporal dynamics or counter overruns. We ran the constant measurements for four servers with different maximum throughputs.

Figure 11 also shows that for the most active server (Rank 1), our periodic sampling does not reveal a traffic pattern and the measurements oscillate at high values. This indicates that the sampling frequency is too low and we miss counter overflows. In such cases, our results are an underestimation of the real traffic. This issue only occurs for the two most active servers (i.e., Rank 1 and Rank 2). We refrained from targeting those two servers at a significantly higher sampling frequency to avoid stressing the Netflix infrastructure unnecessarily. As these traffic estimates still provide a lower bound on the actual traffic, we decided to keep these results in the evaluation.

For all other servers, we are highly confident that our results are reliable. The expected diurnal patterns exposed in Figure 11, reinforce our confidence that our measurements observe the actual traffic of the servers.

4.3 Footprint

We first study the traffic peak and off-peak times, before characterising Netflix from a traffic perspective. We first motivate the importance of traffic measurements by contrasting IXP with ISP deployments on a per server and per location basis. We then analyse total traffic volumes on a per continent and per country basis.

4.3.1 Peak and Off-peak traffic

Given the nature of the content delivered by Netflix, we expect the hours of peak traffic throughput to be similar across all countries. More specifically, we expect traffic to peak late in the evening, just before midnight. As servers are distributed amongst different time zones, we normalised our measurements to UTC time to expose these peaks. Figure 12 depicts the time and value of the peak packet throughput for each server and day (i.e., one marker per server and day).

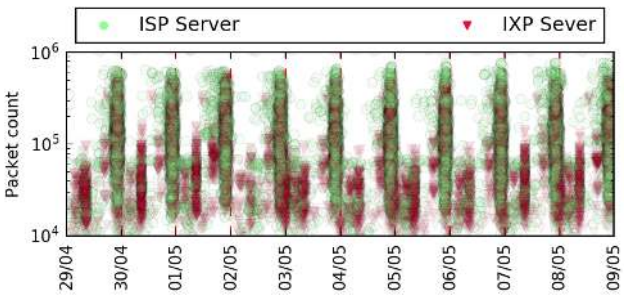

Figure 12: Time of peak traffic throughput per server (normalised to UTC time).

Interestingly, Figure 12 reveals two distinct kinds of traffic. As expected, we observe a majority of servers peaking just before midnight. These late evening peaks further support the validity of our geolocation and traffic estimation methodology. We also observe a much smaller daily peak around 8am and only for IXP servers. This points to server-to-server traffic to update the available content. As a matter of fact, Netflix claims that the default period to upload new video material on the servers is from 2am to 2pm. A peak at 8am is thus right in the middle of this period.

8am is very likely to be off-peak time for ISPs. We conjecture that this off-peak content update strategy is a cost saving mechanism. Since traffic charges typically depend on the peak traffic, this is a cost reducing mechanism for Netflix and an incentive for ISPs to deploy Netflix servers.

4.3.2 *Traffic-Footprints per Location*

We now contrast the traffic-footprints corresponding to different locations. We use the same mechanism to translate observed packet throughput in expected traffic volume as in the previous section.

In Figure 13, we show the average traffic volume per location and the number of servers deployed at each location. On a first glimpse, the figure shows that IXP locations produce significantly more traffic than ISP locations do. Furthermore, the figure also exposes that relying only on the number of servers to infer the total traffic volume of Netflix would lead to very inaccurate picture. The traffic volume per server is not strongly correlated with the size of the deployment at that location, though larger deployment locations will naturally deliver more traffic. We also observe that despite most ISP deployments being comparably small in size, the generated traffic volumes are quite different. IXP deployments show the same behaviour but with traffic levels significantly larger than for ISP deployments.

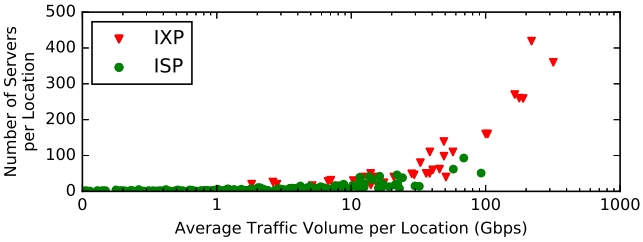

Figure 13: Deployment size per location vs average traffic per location.

Figure 14 shows the cumulative distribution of average traffic per server. We do not observe significant differences between ISP and IXP servers. However, the actual amount varies significantly from server to server, ranging from 10Mbps up to a few Gbps of traffic. Most servers generate traffic volumes in the range between 100 Mbps and 1Gbps. The variations in traffic across servers are strong enough, making accurate estimates of regional traffic volumes difficult only from the sheer server deployment.

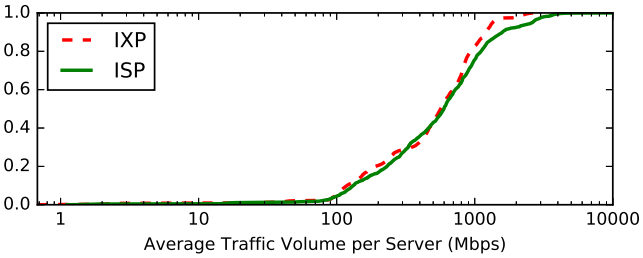

Figure 14: Average traffic generated by a single Netflix server.

4.3.3 *Traffic-Footprints per Region*

From our measurements, we estimate how much traffic is served by the Netflix infrastructure. We assume that most packets use the Ethernet Maximum Transmission Unit (MTU) of 1500 bytes. To make better use of the available link capacity, servers tend to make individual packets as large as possible, particularly for video traffic, when the content is large in size. We therefore estimate the total traffic by multiplying the observed packet rates with this MTU size[5].

In Figure 15, we show the results on a per continent basis. The plots expose a diurnal pattern, a closer investigation again reveals that the peaks per continent are shifted consistently with their geographic location.

North America is the dominating continent in terms of traffic. South America and Oceania are comparable, whereas in Asia and Africa we only observe minor traffic volumes. This is in line with our expectations on the strength of the individual Netflix markets.

To observe the specifics of individual markets, we now break down the results on a country level. Figure 16 shows the average traffic per country and server type. We observe that the USA is by far the largest Netflix market. Surprisingly, Mexico comes in the second place. Indeed, we found evidence that already back in 2014, Mexico was amongst the five largest countries in terms of Netflix paying subscribers [10]. While Netflix produces content in Spanish language, it is rather surprising that Mexico is ahead of the English-speaking UK. The UK, the largest English-speaking market outside the USA, comes in third, followed by Canada. Brazil comes unexpectedly fifth in front of Australia, despite not being an English-speaking country.

[5]While this is an assumption, we are confident that at least the distribution of MTU sizes is similar across all regions, so the relative orderings within our findings are reliable.

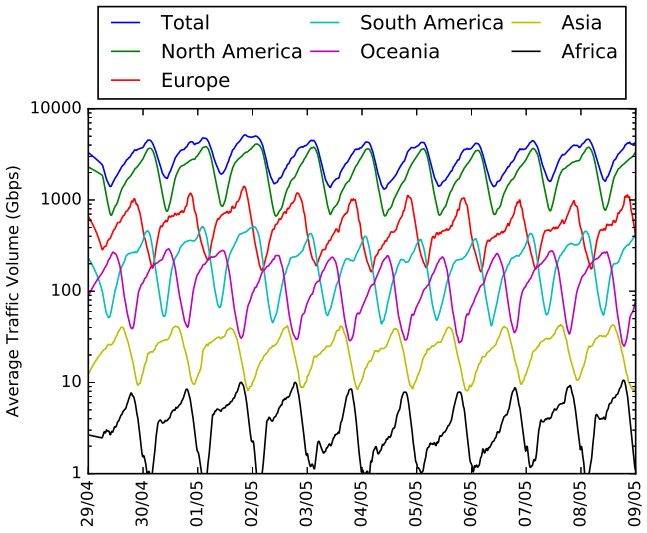

Figure 15: Average traffic per continent over all IXP and ISP servers.

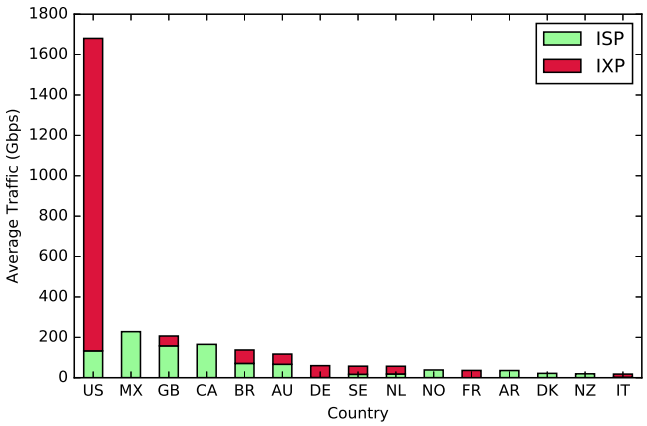

Figure 16: Average traffic generated by Netflix servers per country.

4.3.4 *Discussion*

As already highlighted in the previous sections, our results reinforce how important IXPs are for Netflix to deliver its content. For the USA, the vast majority of Netflix's traffic is served from IXP locations. Outside the USA, IXPs also play a major role to deliver Netflix traffic, e.g., in Brazil, Germany, Netherlands, and France. Netflix provides evidence for the central role played by IXPs in the Internet ecosystem [17, 15]. This also suggests that IXPs constitute strategic locations from where it is possible to deliver significant amounts of traffic to end-users in a cost-effective manner. As Netflix servers are located at the edge of the Internet (either at IXPs or inside ISPs), it also reinforces the observed flattening of the Internet [24] and the impact that big content players have in it [28].

5. SUMMARY

In this paper we uncovered the world-wide server deployment used by Netflix to deliver video content. Our approach relied on the deterministic structure of the DNS names used by Netflix. We enriched the findings on the deployment with estimates of the traffic generated by the servers. For this, we exploited the way in which the servers populate the ID field in the IP header of the generated packets.

We studied the Netflix infrastructure deployment on a per region basis. We discovered different deployment strategies across regions, in terms of the relative importance of IXP and ISP locations. We found that IXP locations do rely on significant number of servers at few locations, while ISP deployments are smaller in size but often at many locations. We also find that the sheer size of the deployments reflect the various markets of Netflix quite well. The traffic figures we obtained support our findings from the deployment sizes.

6. ACKNOWLEDGEMENTS

This work is supported by the European Union's Horizon 2020 research and innovation programme under the ENDEAVOUR project (grant agreement 644960).

7. REFERENCES

[1] Browsermob-proxy by lightbody. `https://bmp.lightbody.net/`.

[2] Completing the Netflix Cloud Migration. `https://media.netflix.com/en/company-blog/completing-the-netflix-cloud-migration/`.

[3] GeoLite2 Free Downloadable Databases Maxmind Developer Site. `https://dev.maxmind.com/geoip/geoip2/geolite2/`.

[4] Hola - Free VPN, Secure Browsing, Unrestricted Access. `https://bmp.lightbody.net/`.

[5] Hping - Active Network Security Tool. `http://hping.org/`.

[6] Internet Tolls And The Case For Strong Net Neutrality. `https://media.netflix.com/en/company-blog/internet-tolls-and-the-case-for-strong-net-neutrality/`.

[7] IP to ASN Mapping - Team Cymru. `http://www.team-cymru.org/IP-ASN-mapping.html`.

[8] Netflix ISP Speed Index. `https://ispspeedindex.netflix.com/`.

[9] Netflix Open Connect. `https://openconnect.netflix.com/`.

[10] Netflix: subscribers by country 2014. `http://www.statista.com/statistics/324050/number-netflix-paying-streaming-subscribers/`.

[11] Sandvine Global Internet Phenomena Latin America & North America - May 2015. https://www.sandvine.com/downloads/general/global-internet-phenomena/2015/global-internet-phenomena-report-latin-america-and-north-america.pdf.

[12] Selenium - Web Browser Automation. http://www.seleniumhq.org/.

[13] Verizon Won't Use Netflix's Hardware to Boost Streaming Speeds. http://time.com/2866004/verizon-netflix/.

[14] V. K. Adhikari, Y. Guo, F. Hao, V. Hilt, Z.-L. Zhang, M. Varvello, and M. Steiner. Measurement Study of Netflix, Hulu, and a Tale of Three CDNs. *IEEE/ACM Transactions on Networking*, 23(6):1984–1997, 2015.

[15] B. Ager, N. Chatzis, A. Feldmann, N. Sarrar, S. Uhlig, and W. Willinger. Anatomy of a Large European IXP. In *Proc. of ACM SIGCOMM*, 2012.

[16] B. Ager, W. Mühlbauer, G. Smaragdakis, and S. Uhlig. Web content cartography. In *Proc. of IMC*, 2011.

[17] B. Augustin, B. Krishnamurthy, and W. Willinger. IXPs: mapped? In *Proc. of IMC*, 2009.

[18] M. Calder, X. Fan, Z. Hu, E. Katz-Bassett, J. Heidemann, and R. Govindan. Mapping the expansion of Google's serving infrastructure. In *Proc. of IMC*, 2013.

[19] J. C. Cardona Restrepo and R. Stanojevic. IXP Traffic: A Macroscopic View. In *Proc. of LANC*, 2012.

[20] N. Chatzis, G. Smaragdakis, J. Böttger, T. Krenc, and A. Feldmann. On the Benefits of Using a Large IXP As an Internet Vantage Point. In *Proc. of IMC*, 2013.

[21] N. Chatzis, G. Smaragdakis, A. Feldmann, and W. Willinger. There is More to IXPs Than Meets the Eye. *CCR*, 43(5), Nov. 2013.

[22] N. Chatzis, G. Smaragdakis, A. Feldmann, and W. Willinger. Quo Vadis Open-IX? *CCR*, 45(1), Jan. 2015.

[23] F. Chen, R. K. Sitaraman, and M. Torres. End-user mapping: Next generation request routing for content delivery. In *Proc. of ACM SIGCOMM*, 2015.

[24] P. Gill, M. Arlitt, Z. Li, and A. Mahanti. The flattening internet topology: Natural evolution, unsightly barnacles or contrived collapse? In *Proc. of PAM*. 2008.

[25] M. Handley, V. Paxson, and C. Kreibich. Network Intrusion Detection: Evasion, Traffic Normalization, and End-to-End Protocol Semantics. In *USENIX Security Symposium*, pages 115–131, 2001.

[26] C. Huang, A. Wang, J. Li, and K. W. Ross. Measuring and evaluating large-scale CDNs. In *ACM IMC*, volume 8, 2008.

[27] S. Jain, A. Kumar, S. Mandal, J. Ong, L. Poutievski, A. Singh, S. Venkata, J. Wanderer, J. Zhou, M. Zhu, J. Zolla, U. Hölzle, S. Stuart, and A. Vahdat. B4: Experience with a Globally-deployed Software Defined Wan. In *Proc. of ACM SIGCOMM*, 2013.

[28] C. Labovitz, S. Iekel-Johnson, D. McPherson, J. Oberheide, and F. Jahanian. Internet inter-domain traffic. *CCR*, 41(4):75–86, 2011.

[29] T. Leighton. Improving Performance on the Internet. *Commun. ACM*, 52(2), Feb. 2009.

[30] R. Torres, A. Finamore, J. R. Kim, M. Mellia, M. M. Munafo, and S. Rao. Dissecting video server selection strategies in the youtube CDN. In *Proc. of IEEE ICDCS*, 2011.

[31] J. Touch. Updated Specification of the IPv4 ID Field. RFC 6864 (Proposed Standard), Feb. 2013.

[32] M. Zhang, Y. Ruan, V. S. Pai, and J. Rexford. How DNS Misnaming Distorts Internet Topology Mapping. In *USENIX ATC*, 2006.