

On the Benefits of Using a Large IXP as an Internet Vantage Point

Nikolaos Chatzis
TU Berlin
nikolaos@inet.tu-berlin.de

Georgios Smaragdakis
T-Labs/TU Berlin
georgios@inet.tu-berlin.de

Jan Böttger
TU Berlin
jan@inet.tu-berlin.de

Thomas Krenc
TU Berlin
tkrenc@inet.tu-berlin.de

Anja Feldmann
TU Berlin
anja@inet.tu-berlin.de

ABSTRACT

In the context of measuring the Internet, a long-standing question has been whether there exist well-localized physical entities in today's network where traffic from a representative cross-section of the constituents of the Internet can be observed at a fine-enough granularity to paint an accurate and informative picture of how these constituents shape and impact much of the structure and evolution of today's Internet and the actual traffic it carries.

In this paper, we first answer this question in the affirmative by mining 17 weeks of continuous sFlow data from one of the largest European IXPs. Examining these weekly snapshots, we discover a vantage point with excellent visibility into the Internet, seeing week-in and week-out traffic from all 42K+ routed ASes, almost all 450K+ routed prefixes, from close to 1.5M servers, and around a quarter billion IPs from all around the globe. Second, to show the potential of such vantage points, we analyze the server-related portion of the traffic at this IXP, identify the server IPs and cluster them according to the organizations responsible for delivering the content. In the process, we observe a clear trend among many of the critical Internet players towards network heterogenization; that is, either hosting servers of third-party networks in their own infrastructures or pursuing massive deployments of their own servers in strategically chosen third-party networks. While the latter is a well-known business strategy of companies such as Akamai, Google, and Netflix, we show in this paper the extent of network heterogenization in today's Internet and illustrate how it enriches the traditional, largely traffic-agnostic AS-level view of the Internet.

Categories and Subject Descriptors

Computer Systems Organization, Computer Communication Networks, Network Operations [**Network monitoring**]:

Keywords

Internet Exchange Points, Internet Topology, Traffic Characterization, Content Delivery

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
IMC'13, October 23–25, 2013, Barcelona, Spain.

Copyright is held by the owner/author(s). Publication rights licensed to ACM.
ACM 978-1-4503-1953-9/13/10 ...\$15.00.
<http://dx.doi.org/10.1145/2504730.2504746>.

Acknowledgments

This work would not have been possible without the help, engagement, and commitment of Walter Willinger.

We want to express our gratitude towards the IXP for their generous cooperation and support. Moreover, we thank the anonymous reviewers for their useful feedback. This work was supported in part by the EU projects BigFoot (FP7-ICT-317858) and CHANGE (FP7-ICT-257422), EIT Knowledge and Innovation Communities program, and an IKY-DAAD award (54718944).

1. INTRODUCTION

An ever-growing demand for Web-based traffic [30, 41] (e.g., HD video and other streaming media, e-commerce services), together with the proliferation of new Internet-enabled devices and the emergence of new content distribution models and cloud infrastructure providers are radically transforming the nature of content delivery in today's Internet. These features are also having a profound impact on how some of the main Internet players (e.g., ISPs, CDNs, Web hosting companies, and content providers) operate in such a dynamic environment and do business in an increasingly competitive marketplace. Unfortunately, carefully tracking these developments to obtain an accurate picture of how this critical cast of players shapes and impacts much of the Internet and its traffic has become an increasingly daunting task. In the past, attempts at painting such a picture had limited success because they typically relied on piecing together incomplete and often inaccurate information from many different sources of varying quality [40, 45, 50] or using various types of hard-to-get (i.e., proprietary) datasets [41] or hard-to-justify estimates of difficult-to-measure (e.g., inter-AS traffic matrix) quantities [27].

A main reason for our current inability to accurately track a constantly changing Internet has been the lack of global vantage points where traffic from a sufficiently large portion of the Internet can be observed at a granularity that is sufficiently fine-grained to discern the make up of today's Internet traffic and the interactions of the responsible parties. This raises the questions whether or not such vantage points do indeed exist in today's Internet, and if so, what exactly they enable us to discern about the Internet as a whole as well as its individual constituents?

Main contribution (part 1): The first finding reported in this paper is that some of the largest Internet Exchange Points (IXPs) in Europe (i.e., AMS-IX in Amsterdam, DE-CIX in Frankfurt, LINX in London) do indeed serve as the kind of vantage points we are looking for. In particular, our detailed analysis of 17 contiguous weeks of complete sFlow measurements from one of the largest European IXPs reveals that in addition to being a critical part of the

European portion of the Internet, this IXP also plays a global role, “seeing” traffic from a large fraction of the Internet. To highlight the kind of visibility into the Internet as a whole that is possible at these vantage points, we show that week-in and week-out, our IXP “sees” traffic from all 42K+ routed ASes, almost all 450K+ routed prefixes, from about 1.5M servers, and around a quarter billion IP addresses from all the countries around the globe. Importantly, the fact that these largest European IXPs all operate in a very similar fashion, handle comparable traffic volumes, and have similar profiles with respect to the members and range of service offerings allows us to conclude that either one of them fits the role of being a global Internet vantage point.

However, like all vantage points used in practice, visibility from these large European IXPs into the Internet is not perfect. Thus, knowing what we cannot discern about what aspects of the Internet from mining the traffic seen at one of these vantage points is as important as knowing what we can discern. To deal with this important issue, we consider a variety of different IXP-external measurements that not only complement the IXPs own data, but also enable us to either check or validate our findings that are based on the IXPs data alone or help us determine what aspects are impossible to discern from this vantage point, and why.

Main contribution (part 2): Our second main finding illustrates the new opportunities or benefits for Internet measurements that arise from having access to vantage points with such good visibility into the Internet. By squarely focusing on the analysis of the Web server-related portion of the traffic “seen” by our IXP and applying an original methodology for identifying server-based infrastructures, classifying them by ownership (e. g., company, organization), and associating traffic with them, we show that many of the major commercial Internet players are actively contributing to a clearly discernible trend towards network heterogenization. That is, many of these players’ networks are undergoing major changes as a result of business decisions that either favor the hosting of servers from third-party networks within their own network infrastructures or the deployment of their own servers, often in massive numbers, in purposefully-selected third-party networks.

On the one hand, this finding simply confirms what is well-documented in press releases of some of the main commercial Internet players and technology blogs, but remains a largely under-reported issue in the networking research literature. For example, it is well-known that CDN companies such as Akamai are accelerating the deployment of their own servers in eyeball networks by forming strategic content delivery alliances with large ISPs [1, 6]. Similarly, it is also widely reported that large content providers such as Netflix are installing their own brand of single-purpose CDNs (e. g., Netflix’s Open Connect) inside various regional and local ISPs for the sole purpose of enabling those ISPs to deliver Netflix video data directly to their end users [5]. Other key companies such as Google, Amazon, or Facebook have followed suit and further add to the complexities that result from this proliferation of intertwined network infrastructures and traffic.

On the other hand, by providing a methodology for discovering an organization’s servers, whether they are deployed within the organization’s own AS (or ASes) or inside some third-party network’s infrastructure, we enable and advocate a measurement-driven approach to inferring and assessing the *extent* to which individual network infrastructures or the traffic that they carry are heterogeneous. The results of our study show that network heterogenization is wide-spread and not just confined to well-known players such as Akamai or Google. We also study the impact that this finding has on the usage of peering links at IXPs, in particular, and AS-links in the larger Internet, in general. Specifically, we

show that not only is AS-link usage becoming more heterogeneous, but the task of attributing traffic to the party responsible for it faces enormous challenges in view of the complexities that result from the observed proliferation of intertwined network infrastructures and traffic. In effect, these findings argue that future attempts at accurately and meaningfully studying the business strategies of the various Internet constituents (e. g., ISPs, content providers, CDNs, IXPs, networks without an ASN, resellers, etc.) and the business relationships among them have to move beyond the traditional and largely traffic-agnostic AS-level view of the Internet. They must account for the complexities that result from today’s Internet realities and are more concerned with how traffic flows across the network than with how the network is connected at the AS level.

Head-on comparison with [13]: The large European IXP considered in this paper also featured prominently in the recent work by Ager et al. [13]. However, while that study focused squarely on the discovery of a surprisingly rich peering fabric among the member ASes of that IXP, in this paper, we are mainly concerned with mining the traffic seen at this IXP to determine the IXP’s visibility into the Internet. Put differently, while [13] exploited the IXP measurements to obtain an accurate picture of the “inside” of this IXP (i. e., its member ASes, their peerings, and the IXP-specific traffic matrix), this paper mines recent traffic data to obtain a view of the “outside” of the IXP; that is, the larger Internet beyond the boundary formed by the members of the IXP. In terms of results, while [13] highlighted the severe level of incompleteness of the commonly-studied AS maps of the Internet, this paper establishes and provides concrete evidence for why and in what sense this traditional AS-level view— although still useful for exploring and understanding various connectivity- or reachability-related aspects— is largely inept for accounting for critical elements of the networks that make up today’s Internet. Thus, representing two largely complementary efforts, the combined findings of [13] and of this paper take the observations of the study by Labovitz et al. [41] to the next level. In the process, we identify and outline an alternative and largely orthogonal perspective to the traditional AS-level view that centers around organizations or companies and their server-based infrastructures that are spread across many networks and countries and defy traditional network and geographic boundaries.

Road-map: The remainder of the paper is organized as follows. In Section 2, we describe the available IXP measurements and report on how we extract the relevant traffic data for our study. We focus in Section 3 on our first finding that the IXP defines a vantage point that offers unprecedented visibility into the Internet, and we support this finding with concrete results based on a week’s worth of traffic. Our results are complemented in Section 4 by a longitudinal analysis, and we discuss what measurements taken at this single vantage point for a 17-week long period reveal about the Internet. In Section 5, we describe our methodology for grouping an organization’s servers, irrespective of their location within the Internet, and report on our second main finding; that is, evidence for and extent of network heterogenization. We continue in Section 6 with a discussion of related and future work and conclude with a summary of our main observations in Section 7.

2. IXP AS RICH DATA SOURCE

In this section, we describe the IXP measurements that are at our disposal for this study and sketch and illustrate the basic methodology we use to identify the traffic components relevant for our work. We also list and comment on the different IXP-external datasets that




Figure 1: Traffic filtering steps

we rely on throughout this paper to show what we can and cannot discern from the IXP-internal measurements alone.¹

2.1 Available IXP-internal datasets

The work reported in this paper is based on traffic measurements collected between August 27 (beginning of week 35) and December 23 (end of week 51) of 2012 at one of the largest IXPs in Europe. At the beginning of the measurement period in week 35, this IXP had 443 member ASes that exchanged on average some 11.9 PB of traffic per day over the IXP’s public peering infrastructure (i. e., a layer-2 switching fabric distributed over a number of data centers within the city where the IXP is located). During the measurement period, the IXP added between 1-2 members per week. Specifically, the measurements we rely on consist of 17 consecutive weeks of uninterrupted anonymized sFlow records that contain Ethernet frame samples that were collected using a random sampling of 1 out of 16K. sFlow captures the first 128 bytes of each sampled frame. This implies that in the case of IPv4 packets the available information consists of the full IP and transport layer headers and 74 and 86 bytes of TCP and UDP payload, respectively. For further details about the IXP infrastructure itself as well as the collected sFlow measurements (e. g., absence of sampling bias), we refer to [13]. In the following, we use our week 45 data to illustrate our method. The other weekly snapshots produce very similar results and are discussed in more detail in Section 4.

2.2 Methods for dissecting the IXP’s traffic

2.2.1 Peering traffic

Figure 1 details the filtering steps that we applied to the raw sFlow records collected at this IXP to obtain what we refer to as the “peering traffic” component of the overall traffic. As shown in Figure 1, after removing from the overall traffic, in succession, all non-IPv4 traffic (i. e., native IPv6 and other protocols; roughly 0.4% of the total traffic, most of which is native IPv6), all traffic that is either not member-to-member or stays local (e. g., IXP management traffic; about 0.6%), all member-to-member IPv4 traffic that is not TCP or UDP (i. e., ICMP and other transport protocols; less than 0.5%), this peering traffic makes up more than 98.5% of the total traffic. As an interesting by-product, we observe that 82% of the peering traffic is TCP and 18% is UDP.

2.2.2 Web server-related traffic

We next identify the portion of the peering traffic that can be unambiguously identified as Web server-related traffic. Our motivation is that Web servers are generally considered to be the engines

of e-commerce, which in turn argues that Web server-related traffic is, in general, a good proxy for the commercial portion of Internet traffic. Accordingly, we focus on HTTP and HTTPS and describe the filtering steps for extracting their traffic.

To identify HTTP traffic, we rely primarily on commonly-used string-matching techniques applied to the content of the 128 bytes of each sampled frame. We use two different patterns. The first pattern matches the initial line of request and response packets and looks for HTTP method words (e. g., GET, HEAD, POST) and the words HTTP/1.{0,1}. The second pattern applies to header lines in any packet of a connection and relies on commonly used HTTP header field words as documented in the relevant RFCs and W3C specifications (e. g., Host, Server, Access-Control-Allow-Methods). Using these techniques enables us to identify which of the IP endpoints act as servers and which ones act as clients. When applied to our week 45 data, we identify about 1.3 million server IPs together with roughly 40 million client IPs. Checking the port numbers, we verify that more than 80% of the server IPs use the expected TCP ports, i. e., 80 and 8080. Some 5% of them also use 1935 (RTMP) as well as 443 (HTTPS). Note that by relying on string-matching, we miss those servers for which our sFlow records do not contain sufficient information; we also might mis-classify as clients some of those servers that “talk” with other servers and for which only their client-related activity is captured in our data.

With respect to HTTPS traffic, since we cannot use pattern matching directly due to encryption, we use a mixed passive and active measurement approach. In a first step, we use traffic on TCP port 443 to identify a candidate set of IPs of HTTPS servers. Here, we clearly miss HTTPS servers that do not use port 443, but we consider them not to be commercially relevant. However, given that TCP port 443 is commonly used to circumvent firewalls and proxy rules for other kinds of traffic (e. g., SSH servers or VPNs running TCP port 443), in a second step we rule out non-HTTPS related use by relying on active measurements. For this purpose, we crawl each IP in our candidate set for an X.509 certificate chain and check the validity of the returned X.509 certificates. For those IPs that pass the checks of the certificate, we extract the names for which the X.509 certificate is valid and the purpose for which it was issued. In particular, we check the following properties in each retrieved X.509 certificate: (a) *certificate subject*, (b) *alternative names*, (c) *key usage* (purpose), (d) *certificate chain*, (e) *validity time*, and (f) *stability over time*. If a certificate does not pass any of the tests, we do not consider it in the analysis.

We keep only the IPs that have a certificate subject and alternative names with valid domains and also valid country-code second-level domains (ccSLD) according to the definition in [35]. Next, we check if the key usage explicitly indicates a Web server role. In the certificate chain we check if the delivered certificates do really refer to each other in the right order they are listed up to the root certificate, which must be contained in the current Linux/Ubuntu white-list. Next, we verify the validity time of each certificate in the chain by comparing it to the timestamp the certificate fetching was performed. Lastly, we perform the active measurements several times and check for changes because IPs in cloud deployments can change their role very quickly and frequently. Ignoring validity time, we require that all the certificates fetched from a single IP have the same properties. In the case of our week 45 data, starting with a candidate set of approximately 1.5M IPs, some 500K respond to repeated active measurements, of which 250K are in the end identified as HTTPS server IPs.

When combined, these filtering steps yield approximately 1.5M different Web server IPs (including the 250K HTTPS server IPs). In total, these HTTP and HTTPS server IPs are responsible for or

¹For an overview of the importance of IXPs for today’s Internet, we refer to [28].

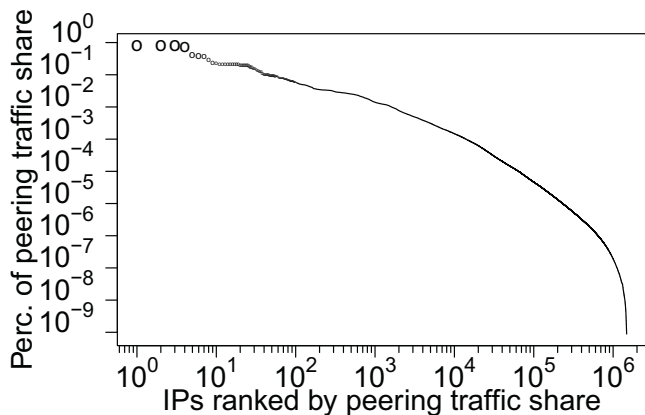


Figure 2: Traffic per server IP sorted by traffic share

“see” more than 70% of the peering traffic portion of the total traffic. Some 350K of these IP addresses appear in both sets and are examples of multi-purpose servers; that is, servers with one IP address that see activity on multiple ports. Multi-purpose servers are popular with commercial Internet players (e.g., Akamai which uses TCP port 80 (HTTP) and TCP port 1935 (RTMP)), and their presence in our data partially explains why we see a larger percentage of Web server-related traffic than what is typically reported in the literature [33, 36, 43], but is often based on a strictly port-based traffic classification [30, 41].

Among the identified HTTP and HTTPS server IPs, we find some 200K IPs that act both as servers and as clients. These are responsible for some 10% of the server-related traffic. Upon closer inspection of the top contributors in this category, we find that they typically belong to major CDNs (e.g., EdgeCast, Limelight) or network operators (e.g., Eweca). Thus, the large traffic share of these servers is not surprising and reflects typical machine-to-machine traffic associated with operating, for example, a CDN. Another class of IPs in this category are proxies or clients that are managed via a server interface (or vice versa).

In the context of this paper, it is important to clarify the notion of a server IP. Throughout this paper, a server IP is defined as a publicly routed IP address of a server. As such, it can represent many different real-world scenarios, including a single (multi-purpose) server, a rack of multiple servers, or a front-end server acting as a gateway to possibly thousands of back-end servers (e.g., an entire data center). In fact, Figure 2 shows the traffic share of each server IP seen in the week 45 data. It highlights the presence of individual server IPs that are responsible for more than 0.5% of all server-related traffic! Indeed, the top 34 server IPs are responsible for more than 6% of the overall server traffic. These server IPs cannot be single machines. Upon closer examination, they are identified as belonging to a cast of Internet players that includes CDNs, large content providers, streamers, virtual backbone providers, and resellers, and thus represent front-end servers to large data centers and/or anycast services. Henceforth, we use the term server to refer to a server IP as defined above.

2.3 Available IXP-external datasets

When appropriate and feasible, we augment our IXP-based findings with active and passive measurements that do not involve the IXP in any form or shape and are all collected in parallel to our IXP data collection. Such complementary information allows us to verify, check, or refine the IXP-based findings.

One example of a complementary IXP-external dataset is a proprietary dataset from a large European Tier-1 ISP consisting of packet-level traffic traces.² With the help of the network intrusion detection system Bro [47] we produce the HTTP and DNS logs, extract the Web server-related traffic and the corresponding server IPs from the logs, and rely on the resulting data in Section 3.

For another example, we use the list of the top 280K DNS recursive resolvers—as seen by one of the largest commercial CDNs—as a starting set to find a highly distributed set of DNS resolvers that are available for active measurements such as doing reverse DNS lookups or performing active DNS queries. From this initial list of DNS servers, we eliminate those that cannot be used for active measurements (i.e., those that are not open, delegate DNS resolutions to other resolvers, or provide incorrect answers) and end up with a final list of about 25K DNS resolvers in some 12K ASes that are used for active measurements in Section 3.3.

Other examples of IXP-external data we use in this work include the publicly available lists of the top-1M or top-1K popular Web sites that can be downloaded from www.alexa.com. We obtained these lists for each of the weeks for which we have IXP data. We also utilized blogs and technical information found on the official Web sites of the various technology companies and Internet players. In addition, we make extensive use of publicly available BGP-based data that is collected on an ongoing basis by RouteViews [8], RIPE RIS [7], Team Cymru [9], etc.

2.4 IP server meta-data

Our efforts in Section 5 rely on certain meta-data that we collect for server IPs and that is obtained from DNS information, URIs, and X.509 certificates from HTTPS servers.

Regarding DNS information, obvious meta-information is the hostname(s) of a server IP. This information is useful because large organizations [38] often follow industrial standards in their naming schema’s for servers that they operate or host in their own networks. Another useful piece of meta-data is the Start of Authority (SOA) resource record which relates to the administrative authority and can be resolved iteratively. This way one can often find a common root for organizations that do not use a unified naming schema. Note that the SOA record is often present, even when there is no hostname record available or an ARPA address is returned in the reverse lookup of a server IP.

Next, the URI as well as the authority associated with the hostname give us hints regarding the organization that is responsible for the content. For example, for the URI `youtube.com`, one finds the SOA resource record `google.com` and thus can associate Youtube with Google.

Lastly, the X.509 certificates reveal several useful pieces of meta-data. First, they list the base set of URIs that can be served by the corresponding server IP. Second, some server IPs have certificates with multiple names that can be used to find additional URIs. This is typically the case for hosting companies that host multiple sites on a single server IP. In addition, it is used by CDNs that serve multiple different domains with the same physical infrastructure. Moreover, the names found in the certificates can be mapped to SOA resource records as well.

Overall, we are able to extract DNS information for 71.7%, at least one URI for 23.8%, and X.509 certificate information for 17.7% of the 1.5M server IPs that we see in our week 45 data. For 81.9% of all the server IPs, we have at least one of the three pieces of information. For example, for streamers, one typically has no

²For this trace we anonymized the client information before applying the analysis with the network intrusion detection system Bro. We always use a prefix preserving function when anonymizing IPs.

		week 45	educated guesses of ground-truth
Peering Traffic	IPs	232,460,635	unknown < 2^{32}
	#ASes	42,825	approx. 43K
	Subnets	445,051	450K+
	countries	242	250
Server Traffic	IPs	1,488,286	unknown
	#ASes	19,824	unknown
	Subnets	75,841	unknown
	Countries	200	250

Table 1: IXP summary statistics—week 45

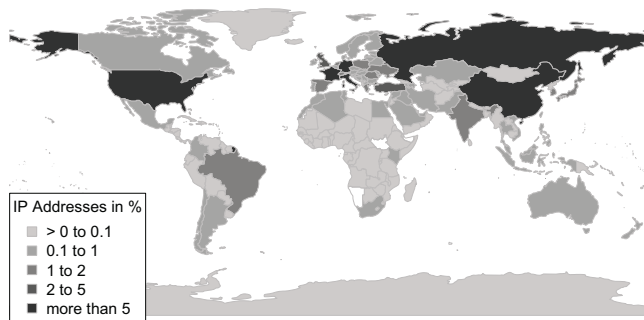


Figure 3: Percentage of IPs per country—week 45

assigned URI, but information from DNS. Before using this rich meta-data in Section 5, we clean it by removing non-valid URIs, SOA resource records of the Regional Internet Registries (RIRs) such as `ripe.net`, etc. This cleaning effort reduces the pool of server IPs by less than 3%.

3. LOCAL YET GLOBAL

The main purpose of this section is to show that our IXP represents an intriguing vantage point, with excellent visibility into the Internet as a whole. This finding of the IXP’s important global role complements earlier observations that have focused on the important local role that this large European IXP plays for the greater geographic region where it is located [13], and we further elaborate here on its dual role as a local and as a global player. Importantly, we also discuss what we can and cannot discern about the Internet as a whole or its individual constituents based on measurements taken at this vantage point. The reported numbers are for the week 45 measurements when the IXP had 452 members that exchanged some 14 PB of traffic per day and are complemented by a longitudinal analysis in Section 4.

3.1 On the global role of the IXP

By providing a well-defined set of steps and requirements for establishing peering links between member networks, IXPs clearly satisfy the main reason for why they exist in the first place – keeping local traffic local. To assess the visibility into the global Internet that comes with using a large European IXP as a vantage point, we focus on the peering traffic component (see Section 2.2.1) and summarize in Table 1 the pertinent results.

First, in this single geographically well-localized facility, we observe during a one-week period approximately a quarter billion

unique IPv4 addresses (recall that the portion of native IPv6 traffic seen at this IXP is negligible). While the total number of publicly routed IPv4 addresses in the Internet in any given week is unknown, it is some portion of the approximately three and a half billion allocated IPv4 addresses, which suggests that this IXP “sees” a significant fraction of the ground truth. The global role of this IXP is further illuminated by geo-locating all 230M+ IP addresses at the country-level granularity [49] and observing that this IXP “sees” traffic from every country of the world, except for places such as Western Sahara, Christmas Islands, or Cocos (Keeling) Islands. This ability to see the global Internet from this single vantage point is visualized in Figure 3, where the different countries’ shades of gray indicate which percentage of IPs a given country contributes to the IPs seen at this IXP.

Second, when mapping the encountered 230M+ IP addresses to more network-specific entities such as subnets or prefixes and ASes, we confirm the IXP’s ability to “see” the Internet. More precisely, in terms of subnets/prefixes, this IXP “sees” traffic from 445K subnets; that is, from essentially all actively routed prefixes. Determining the precise number of actively routed prefixes in the Internet in any given week remains an imprecise science as it depends on the publicly available BGP data that are traditionally used in this context (e.g., RouteViews, RIPE). The reported numbers vary between 450K-500K and are only slightly larger than the 445K subnets we see in this one week. With respect to ASes, the results are very similar – the IXP “sees” traffic from some 42.8K actively routed ASes, where the ground truth for the number of actively routed ASes in the Internet in any given week is around 43K [3] and varies slightly with the used BGP dataset.

Lastly, to examine the visibility that this IXP has into the more commercial-oriented Internet, we next use the Web server-related component of the IXP’s peering traffic (see Section 2.2.2). Table 1 shows that this IXP “sees” server-related traffic from some 1.5M IPs that can be unambiguously identified as Web server IPs. Unfortunately, we are not aware of any numbers that can be reliably considered as ground truth of all server IPs in the Internet in any given week. Even worse, available white papers or reports that purportedly provide this information are typically very cavalier about their definition of what they consider as “Web server” and hence cannot be taken at face value [40, 50].

To indirectly assess how the roughly 1.5M Web server IPs seen at this IXP stack up against the unknown number of Web server IPs Internet-wide, we use an essentially orthogonal dataset, namely the HTTP and DNS logs from a large European Tier-1 ISP that does not exchange traffic over the public switching infrastructure of our IXP. Applying the method as described in Section 2, we extract the Web server IPs from this ISP dataset and find that of the total number of server IPs that are “seen” by this ISP, only some 45K are not seen at the IXP. Importantly, for the server IPs seen both at the IXP and the ISP, those we identified as server IPs using the IXP-internal data are confirmed to be indeed server IPs when relying on the more detailed ISP dataset. In any case, mapping the 1.5M server IPs from the IXP to prefixes, ASes, and countries shows that this IXP “sees” server-traffic from some 17% of all actively routed prefixes, from about 50% of all actively routed ASes, and from about 80% of all the countries in the world.

3.2 On the IXP’s dual role

Visuals such as Figure 3 illustrate that by using this IXP as a vantage point, we are able to see peering traffic from every country and corner of the world or from almost every AS and prefix that is publicly routed. However, such figures do not show whether or not certain countries or corners and ASes or prefixes are better visible than

rank	All IPs Country	Server IPs Country	All IPs Network	Server IPs Network	
IPs	1	US	DE	Chinanet	Akamai
	2	DE	US	Vodafone/DE	1&1
	3	CN	RU	Free SAS	OVH
	4	RU	FR	Turk Telekom	Softlayer
	5	IT	GB	Telecom Italia	ThePlanet
	6	FR	CN	Liberty Global	Chinanet
	7	GB	NL	Vodafone/IT	HostEurope
	8	TR	CZ	Comnet	Strato
	9	UA	IT	Virgin Media	Webazilla
	10	JP	UA	Telefonica/DE	Plusserver
Traffic	1	DE	US	Akamai	Akamai
	2	US	DE	Google	Google
	3	RU	NL	Hetzner	Hetzner
	4	FR	RU	OVH	Vkontakte
	5	GB	GB	Vkontakte	Leaseweb
	6	CN	EU	Kabel Deu.	Limelight
	7	NL	FR	Leaseweb	OVH
	8	CZ	RO	Vodafone/DE	EdgeCast
	9	IT	UA	Unitymedia	Link11
	10	UA	CZ	Kyivstar	Kartina

Table 2: Top 10 contributors—week 45

		Member AS $A(L)$	Distance 1 $A(M)$	Distance > 1 $A(G)$
Peering Traffic	IPs	42.3%	45.0%	12.7%
	Prefixes	10.1%	34.1%	55.8%
	ASes	1.0%	48.9%	50.1%
	Traffic	67.3%	28.4%	4.3%
Server Traffic	IPs	52.9%	41.2%	5.9%
	Prefixes	17.2%	61.9%	20.9%
	ASes	2.2%	61.5%	36.3%
	Traffic	82.6%	17.35%	0.05%

Table 3: IXP as local yet global player—week 45

others in the sense that they are responsible for more traffic that is exchanged over the public switching fabric of the IXP. In particular, we would like to know whether the importance of the local role that this IXP plays for the larger geographic region within which it is situated is more or less recovered when considering the peering or server-related traffic that the IPs or server IPs are responsible for, respectively. To this end, we show in Table 2 the top-10 countries in terms of percentage of IP addresses (and associated traffic) and percentage of server IPs (and associated traffic). In addition, we show the top-10 networks. While the role of the IXP for the European region becomes more dominant when we change from peering to server-related traffic, there are still prominent signs of the IXP’s global role, even with respect to the commercial Internet, and they reflect the relative importance of this IXP for countries such as USA, Russia, and China or ASes such as 20940 (Akamai), 15169 (Google), and 47541 (VKontakte).

For a somewhat simplified illustration of the IXP’s dual role as a local as well as global player, we divide the set of all actively routed ASes into three disjoint sets, $A(L)$, $A(M)$, and $A(G)$. $A(L)$ consists of the member ASes of the IXP; $A(M)$ consists of all ASes that are distance 1 (measured in AS-hops) from a member AS; and

$A(G)$ is the complement of $A(L) \cup A(M)$ and contains those ASes that are distance 2 or more from the member ASes. Intuitively, the set $A(L)$ captures the importance of the local role of the IXP, whereas the set $A(G)$ is more a reflection of the IXP’s global role, with $A(M)$ covering some middle ground. Table 3 shows the breakdown of the IPs, prefixes, and ASes for peering traffic and Web server-related traffic, respectively, for the three sets. It basically confirms our above observation that there is a general trend towards the set $A(L)$ as we move from IPs and the peering traffic they are responsible for to server IPs and their traffic. Note, while the relative importance of the IXP’s local role over its global role with respect to the commercial Internet (i.e., server-related traffic) makes economic sense and is well-captured by this cartoon picture, in reality, there is potentially significant overlap between the sets $A(L)$, $A(M)$, and $A(G)$, e.g., due to remote peerings, IXP resellers, and non-European networks joining the IXP for purely economic reasons. But this is unlikely to invalidate our basic findings concerning the IXP’s dual role.

3.3 On the IXP’s “blind spots”

While the IXP “sees” traffic from much of the Internet, taking measurements exclusively at this single vantage point can tell us only so much about the network as a whole or its individual constituents. Hence, knowing what we can discern about the network with what sort of accuracy is as important as understanding what we cannot discern about it, and why.

We show in Section 3.1 how the use of an essentially orthogonal IXP-external dataset (i.e., the HTTP and DNS logs from the large European Tier-1 ISP) enables us to indirectly assess how the approximately 1.5M server IPs seen at the IXP in a given week compare to the unknown number of server IPs network-wide. In the following, we discuss additional examples where the use of IXP-external data, either in the form publicly available measurements, active or passive measurements, or proprietary information, enables us to check, validate, or refine what we can say with certainty when relying solely on IXP measurements.

To examine in more detail how the approximately 1.5M server IPs seen at the IXP in a given week compare to all server IPs in the Internet, we now use a more extensive combination of IXP-external measurements. To start, using the list of the top-1M Web sites available from `www.alexa.com` and based on the URIs retrieved from the limited payload part of the sampled frames at the IXP, we recover about 20% of all the second-level domains on Alexa’s top-1M list of sites; this percentage increases to 63% if we consider only the top-10K list and to 80% for the top-1K. Note that many hostnames on these lists are dynamic and/or ephemeral. Next, to assess how many additional server IPs we can identify using the approximately 80% of domains we cannot recover using the URIs seen at the IXP, we rely on active measurements in the form of DNS queries to those uncovered domains using our set of 25K DNS resolvers across 12K ASes (see Section 2.3). From this pool of resolvers, we assign 100 randomly-selected resolvers to each URI. This results in approximately 600K server IPs, of which more than 360K are already seen at the IXP and identified as servers.

To provide insight into the remaining 240K server IPs that are not seen as a server at the IXP, we classify them into four distinct categories. First, there are servers of CDNs that are hosted inside an AS and serve exclusively clients in that AS (“private clusters”). These servers reply only to resolvers of that AS for content that is delivered by the global footprint of those CDNs. Traffic to these servers should not be observable at the IXP as it should stay internal to the AS. Second, there are servers of CDNs or cloud/hosting providers that are located geographically far away from the IXP. If

these networks have a global footprint and distribute content in a region-aware manner, it is unlikely that these server IPs are seen at the IXP. The third group includes servers that some ASes operate for the sole purpose of handling invalid URIs. Finally, the last category contains those servers of small organizations and/or universities in parts of the world that are geographically far away from the IXP. These IPs are typically not visible at the IXP. In terms of importance, the first two categories account for more than 40% of the 240K servers not seen at the IXP.

For a concrete example for illustrating “what we know we don’t know”, we consider Akamai. In our week-long IXP dataset, we observe some 28K server IPs for Akamai in 278 ASes (for details, see Section 5). However, Akamai publicly states that it operates some 100K servers in more than 1K ASes [15]. The reasons why we cannot see this ground truth relying solely on our IXP-internal data are twofold and mentioned above. First Akamai is known to operate “private clusters” in many third-party networks which are generally not visible outside those ASes and therefore cannot be detected at the IXP. Second, we cannot expect to uncover Akamai’s footprint in regions that are geographically far away from the IXP, mainly because Akamai uses sophisticated mechanisms to localize traffic [42, 46]. Akamai’s large footprint makes discovering all of its servers difficult, but by performing our own diligently chosen IXP-external active measurements [39] that utilize the URIs collected in the IXP and the open resolvers discussed in Section 2.3, we were able to discover about 100K servers in 700 ASes. Thus, even for a challenging case like Akamai, knowing what our IXP-internal data can and cannot tell us about its many servers and understanding the underlying reasons is feasible.

Regarding our assumption that server-related traffic is a good proxy for the commercial portion of Internet traffic, there are clearly components of this commercial traffic that are not visible at the IXP. For example, the recently introduced hybrid CDNs (e. g., Akamai’s NetSession [12]) serve content by servers as well as by end users that have already downloaded part of the content. Since the connections between users are not based on a HTTP/HTTPS server-client architecture but are P2P-based, we may not see them at the IXP. However, while the traffic of these hybrid CDNs is increasing (e. g., the service is mainly used for large files such as software downloads), the overall volume is still very low [31].

Lastly, by the very definition of an IXP, any traffic that does not pass through the IXP via its public-facing switching infrastructure remains entirely invisible to us. For example, this includes all traffic that traverses the IXP over private peering links. IXPs keep the private peering infrastructure separate from its public peering platform, and we are not aware of any kind of estimates of the amount of private peering traffic handled by the IXPs.

Summary: By analyzing in detail the traffic that traverses the physical infrastructure of one of the largest IXPs in Europe, we provide evidence that the large European IXPs such as AMS-IX, DE-CIX, and LINX represent global Internet vantage points that “see” week-in and week-out traffic from hundreds of millions of IPs, from almost all routed prefixes and from all routed ASes, and from essentially every country around the world. We also illustrate these IXPs’ dual role as a global and a local player within the Internet’s ecosystem and caution that despite their outstanding visibility into the Internet, their use as global Internet vantage points comes with caveats (e. g., having “blind spots”).

4. STABLE YET CHANGING

In this section, we report on a longitudinal analysis that covers 17 consecutive weeks and describes what using our large IXP as a vantage point through time enables us to say about the network

as whole, about some of its constituents, and about the traffic that these constituents are responsible for.


4.1 Stability in the face of constant growth

Publicly available data shows that during 2012, this IXP has experienced significant growth, increasing the number of member ASes by 75 and seeing the average daily traffic volume grow by 0.1%. In terms of absolute numbers, we see in week 35 a total of 443 IXP member ASes sending an average daily traffic volume of 11.9 PB over the IXP’s public-facing switching infrastructure. By week 51, the member count stood at 457, and the average traffic volume went up to 14.5 PB/day. For what follows, it is important to note that these newly added member ASes are typically regional and local ISPs or organizations and small companies outside of central Europe for which membership at this IXP makes economic sense. To contrast, all the major content providers, CDNs, Web hosting companies, eyeball ASes, and Tier-1 ISPs have been members at this IXP for some time, but may have seen upgrades to higher port speeds since the time they joined.


Given our interest in the commercial Internet and knowing (see Section 2) that the server-related traffic is more than 70% of the peering traffic seen at the IXP, we focus in the rest of this paper on the server-related portion of the IXP traffic. The initial set of findings from our longitudinal analysis paints an intriguingly stable picture of the commercial Internet as seen from our vantage point. In particular, analyzing in detail each of the 17 weekly snapshots shows that during every week, we see server-related traffic at this IXP from about 20K (i. e., about half of all) actively routed ASes, some 75K or approximately 15% of all actively routed prefixes, and from a pool of server IPs whose absolute size changes only so slightly but tends to increase in the long term.

This last property is illustrated in Figure 4(a) when focusing only on the absolute heights of the different bars that represent the total number of server IPs seen in a given week. When considering the full version of this figure, including the within-bar details, Figure 4(a) visualizes the weekly churn that is inherent in the server IPs seen at the IXPs. To explain, the first bar in Figure 4(a) shows the approximately 1.4M unique server IPs that we see in week 35. The next bar shows that same quantity for week 36, but splits it into two pieces. While the lower (white) piece reflects the portion of all week 36 server IPs that were already seen during week 35, the upper (black) piece represents the set of server IPs that were seen for the first time during week 36. Starting with week 37, we show for each week $n \in \{37, 38, \dots, 51\}$ snapshot a bar that has three pieces stacked on top of one another. While the first (bottom, white) piece represents the server IPs that were seen at the IXP in each one of the week k snapshot ($k = 35, 36, \dots, n$), the second (grey-shaded) piece shows the server IPs that were seen at the IXP in at least one previous week k snapshot ($k = 35, 36, \dots, n - 1$), but not in all; the third (top black) piece represents all server IPs that were seen at the IXP for the first time in week n .


A key take-away from Figure 4(a) is that there is a sizable pool of server IPs that is seen at the IXP during each and every week throughout the 17-week long measurement period. In fact, this stable portion of server IPs that is seen at the IXP week-in and week-out is about 30% as can be seen by looking at the bottom (white) portion of the week 51 bar. Instead of requiring for a server IP to be seen in each and every week, we also consider a more relaxed notion of stability called recurrence. This recurrent pool of server IPs consists of all server IPs that, by week 51, have been seen at the IXP during at least one previous week (but not in each and every previous week), is represented by the grey-shaded portion of the week 51 bar, and consists of about 60% of all server IPs seen in



(a) Churn of server IPs



(b) Churn of server IPs per region



(c) Churn of ASes with servers

Figure 4: Churn of server IPs and ASes that host servers—weeks 35-51

week 51. Note that the number of server IPs seen for the first time in week n (top black portion) decreases over time and makes up just about 10% of all server IPs seen in week 51.

To look in more detail at the stable and recurrent pools of server IPs and examine their churn or evolution during the 17-week long measurement period, we rely on the GeoLite Country database [44] to geo-locate the server IPs to the country level and group them by geographic “region” as follows: DE, US, RU, CN, RoW (rest of world). Figure 4(b) is similar to Figure 4(a), but shows for each week the portions of IPs for each of these five regions and visualizes the make-up of these server IPs in the same way as we did in Figure 4(a). Note that the shown region-specific stable portions in week 51 add up to the 30% number observed in Figure 4(a), and similarly for the region-specific recurrent portions in week 51 (their sum adds up to the roughly 60% portion of the recurrent pool shown in Figure 4(a)). Interestingly, while the stable pool for DE is consistently about half of the overall stable pool of server IPs seen at the IXP, that pool is vanishing small for CN, slightly larger for RU. This is yet another indication of the important role that this IXP plays for the European part of the Internet.

An even more intriguing aspect of stability is seen when we consider the server-related traffic that the server IPs that we see at the IXP are responsible for. For one, we find that the stable pool of server IPs is consistently contributing more than 60% of the server-related traffic. That is, of the server IPs that this IXP “sees” on a weekly basis, more than 30% of them are not only seen week after week, but they are also responsible for most of the server-related traffic that traverses the IXP each week. When considering the weekly recurrent pools of server IP (grey-shaded segments in Figure 4(a)), their traffic portions keep increasing, but only to less than 30% of all server traffic. To examine the make-up of the server-related traffic attributed to the stable and recurrent pools of server IPs, respectively, Figure 5 shows for each week n three bars, each with five segments corresponding to the five regions considered earlier. The first bar is for the server-related traffic portion of all peering traffic that all server IPs see at the IXP in week n ; the second bar reflects the server-related traffic portion in week n attributed to the recurrent pool of server IPs in that week, while the third bar shows the server-related traffic portion in week n that the stable pool of server IPs is responsible for. From Figure 5, we see that while the stable and recurrent pools of server IPs from China are basically invisible at the IXP in terms of their traffic, both US and Russia have the property that the stable pool of server IPs is responsible for much all the server-related traffic seen from those regions at the IXP.

In addition to examining the churn and evolution of the server IPs seen at the IXP, it is also instructive to study the temporal behavior of the subnets and ASes that the encountered server IPs map into. To illustrate, we only consider the ASes and show in Figure 4(c) for ASes what we depicted in Figure 4(b) for server IPs. The key difference between server IPs and ASes is that the stable pool of ASes represented by the white portion of the week 51 bar is about 70% compared to the 30% for the stable pool of server IPs. Thus, a majority of ASes with server IPs is seen at the IXP during each and every week, and the number of ASes that are seen for the first time becomes miniscule over time. In summary, the stable pool of server IPs (about 1/3 of all server IPs seen at the IXP) gives rise to a stable pool of ASes (about 2/3 of all ASes seen at the IXP and have server IPs) and is responsible for much of the server-related traffic seen at the IXP.

4.2 Changes in face of significant stability

One benefit of observing a significant amount of stability with respect to the server-related portion of the overall peering traffic seen at the IXP is that any weekly snapshot provides more or less the same information. At the same time, subsequent weekly snap-




Figure 5: Churn of server traffic by region—weeks 35-51

shots that differ noticeably may be an indication of some change. Next, we briefly discuss a few examples of such changes that we can discern about the Internet as a whole and some of its individual constituents when we have the luxury to observe and measure the network at this IXP for a number of consecutive weeks.

The first example is motivated primarily by our ability described in Section 2.2.2 to specifically look for and identify HTTPS server IPs, but also by anecdotal evidence or company blogs [21, 22] that suggest that due to widespread security and privacy concerns, the use of HTTPS is steadily increasing. To examine this purported increase, we extract for each weekly snapshot all HTTPS server IPs and the traffic that they contribute. When comparing for each week the number of HTTPS server IPs relative to all server IPs seen in that week and the weekly traffic associated with HTTPS server IPs relative to all peering traffic, we indeed observe a small, yet steady increase, which confirms that the Internet landscape is gradually changing as far as the use of HTTPS is concerned.

For a different kind of example for using our IXP vantage point, we are interested in tracking the recently announced expansion of Netflix using Amazon’s EC2 cloud service [10] into a number of Scandinavian countries [23]. To this end, we relied on publicly available data to obtain Amazon EC2’s data center locations [16] and the corresponding IP ranges [17]. We then mined our 17 weeks worth of IXP data and observed for weeks 49, 50, and 51 a pronounced increase in the number of server IPs at Amazon EC2’s Ireland location, the only data center of Amazon EC2 in Europe. This was accompanied by a significant (but still small in absolute terms) increase in Amazon EC2’s traffic. All this suggests that it may be interesting to watch this traffic in the future, especially if the observed changes are in any way related to Netflix becoming available in Northern Europe towards the end of 2012.

Yet another example concerns the detection of regional or national events at this IXP. For example, considering in more detail week 44, which shows up as a clear dip in, say, Figure 4(a), we notice that this week coincides with Hurricane Sandy that had a major impact on the US East Coast region. To examine its impact,

we use the IXP vantage point to discern this natural disaster from traffic that we see at the IXP from a particular Internet constituent, a major cloud provider. Using publicly available information about the cloud platform’s data centers and corresponding IP ranges, we look in our data for the corresponding server IPs and find a total of about 14K. A detailed breakdown by data center location for weeks 43-45 shows a drastic reduction in the number of server IPs seen at the IXP from the US East Coast region, indicating that the platform of this major cloud provider faced serious problems in week 44, with traffic dropping close to zero. These problems made the news, and the example highlights how a geographical distant event such as a hurricane can be discerned from traffic measurements taken at this geographically distant IXP.

Lastly, we also mention that an IXP is an ideal location to monitor new players such as “resellers”. Resellers are IXP member ASes, and their main business is to provide and facilitate access to the IXP for smaller companies that are typically far away geographically from the IXP. For IXPs, the emergence of resellers is beneficial as they extend the reach of the IXP into geographically distant regions and thereby the potential membership base. For example, for a particular reseller at our IXP, we observed a doubling of the server IPs from 50K to 100K in four months, suggesting that this reseller has been quite successful in attracting new networks with significant server-based infrastructures as its customers.

Summary: When concentrating on the Web server-related portion of the IXP traffic and performing a longitudinal analysis over a 17-week long period, we observe significant stability – of all the server IPs for which traffic is observed at the IXP during this 17-week period, some 30% are seen at the IXP week-in and week-out and are responsible for around 60% of all server-related traffic in each and every week. At the same time, the traffic seen at the IXP does exhibit differences from one week to the next, and we illustrate with some examples what different types of changes enable us to say about the network as a whole or some of its individual constituents. In this sense, the traffic seen at these global Internet vantage points can be used as complementary source of information for recent ef-

forts analyzing Internet events such as large-scale outages due to censorship [32], natural disasters [29], etc.

5. BEYOND THE AS-LEVEL VIEW

To illustrate the benefits and new opportunities for Internet measurements that arise from being able to use our IXP as a vantage point with very good visibility into the Internet, we describe in this section an approach for identifying server-based network infrastructures and classifying them by ownership. In the process, we report on a clear trend towards more heterogeneous networks and network interconnections, provide concrete examples, and discuss why and how this observed network heterogenization requires moving beyond the traditional AS-level view of the Internet.

5.1 Alternative grouping of server IPs

To this point, we have followed a very traditional approach for looking at our IXP data in the sense that we measured the IXP’s visibility into the Internet in terms of the number of actively routed ASes or subnets seen at the IXP. However, there exist Internet players (e.g., CDN77, a recently launched low-cost no-commitment CDN; Rapidshare, a one-click hosting service; or certain meta-hosters that utilize multiple hosters) that are not ASes in the sense that they do not have an assigned ASN. Thus, as far as the traditional AS-level view of the Internet is concerned, these players are invisible, and the traffic that they are responsible for goes unnoticed, or worse misattributed to other Internet players. Yet, being commercial entities, these companies actively advertise their services, and in the process often publish the locations and IP addresses of their servers. This then suggests an alternative approach to assessing the IXP’s ability to “see” the Internet as a whole—group servers according to the organization or company that has the administrative control over the servers and is responsible for distributing the content. While this approach is easy and works to perfection for companies like CDN77 that publish all their server IPs, the question is what to do if the server IPs are not known.

Accordingly, our primary goal is to start with the server IPs seen at the IXP and cluster them so that the servers in one and the same cluster are provably under the administrative control of the same organization or company. To this end, we rely in parts on methods described by Plonka et al. [48] for traffic and host profiling, Bermudez et al. [20] for discerning content and services, and Ager et al. [14] for inferring hosting infrastructures from the content they deliver. We also take advantage of different sets of meta-data obtained from assorted active measurement efforts or available by other means as discussed in Section 2.4. Recall that this meta-data includes for every server IP seen in the IXP data the corresponding URIs, the DNS information from active measurements, and, where available, the list of X.509 certificates retrieved via active measurements. In the rest of this section the reported numbers are for week 45.

The clustering proceeds in three steps. First, we focus on those server IPs for which we have a SOA resource record and consider a first category of clusters that have the property that all server IPs assigned to a given cluster have the IP and the content managed by the same authority. We identify those clusters by grouping all server IPs where the SOA of the hostname and the authority of the URI lead to the same entry. Prominent organizations that fall into this first category are Amazon and big players like Akamai and Google when they are located in their own ASes or when they are in third-party ASes but have assigned names to their own servers. 78.7% of all our server IPs are clustered in this first step.

In a second step, we consider clusters with the property that for the server IPs in a given cluster, most of the server IPs and most of the content are managed by the same authority. This can happen

if the SOA is outsourced (e.g., to a third-party DNS provider) and is common property among hosters and domains served by virtual servers. In these cases, to group server IPs, we rely on a majority vote among the SOA resource records, where the majority vote is by (i) the number of IPs and (ii) the size of the network footprint. This heuristic enables us to group some server IPs together with organizations inferred in the previous step and also applies to meta-hosters such as Hostica. 17.4% of all our server IPs are clustered in this second step. Lastly, for the remaining 3.9% of server IPs that have been seen in our IXP data and have not yet been clustered, we only have partial SOA information. This situation is quite common for parts of the server infrastructure of some large content providers and CDNs such as Akamai that have servers deployed deep inside ISPs. In this case, we apply the same heuristic as in the second step, but only rely on the available subset of information.

To validate our clustering that results from this three-step process, we manually compare the results by (1) checking against the coverage of the public IP ranges that some organizations advertise (see Section 4.2), (2) utilizing the information of certificates that point to applications and services, and (3) actively downloading either the front page (e.g., in the case of Google, it is always the search engine front page) or requested content that is delivered by a CDN (e.g., in the case of Akamai, any content is delivered by any of its servers [53]). Our method manages to correctly identify and group the servers of organizations with a small false-positive rate of less than 3%. Moreover, we observe that the false-positive rate decreases with increasing size of the network footprint. However, there are false-negatives in the sense that our methodology misses some servers due to the “blind spots” discussed in Section 3.3.

5.2 New reality (I): ASes are heterogeneous

Equipped with an approach for grouping server IPs by organizations, we examine next to what extent this grouping is orthogonal to the prevailing AS-level view of the Internet. The issues are succinctly illustrated in Figure 6(a) where we augment the traditional AS-level view (i.e., a number of different ASes exchanging traffic over (public) peering links at an IXP) with new features in the form of AS-internal details (i.e., the third-party servers that the ASes host). Note that while the traditional view that makes a tacit homogeneity assumption by abstracting away any AS-internal details may have been an adequate model for understanding some aspects of the Internet and the traffic it carries at some point in the past, things have changed, and we assert that the cartoon picture in Figure 6(a) captures more accurately the current Internet reality; that is, a trend towards distributed network infrastructures that are deployed and operated by today’s commercial Internet players.

To quantify how much closer the cartoon Figure 6(a) is to reality than the traditional AS-level view, we apply our clustering approach to the 1.5M server IPs seen in our week 45 IXP data and obtain some 21K clusters, henceforth referred to organizations or companies. Among them are the well-known big players like Akamai with 28K active server IPs, Google with 11.5K server IPs, and several large hosters, each with more than 50K server IPs (e.g., AS92572 with 90K+ server IPs; AS56740 and AS50099, both with more than 50K server IPs). Indeed, of the 21K identified organizations, a total of 143 organizations are associated with more than 1000 server IPs and more than 6K organizations have more than 10 servers IPs. For the latter, Figure 6(b) shows a scatter plot of the number of server IPs per organization vs. the number of ASes that they cover. More precisely, every dot in the plot is an organization, and for a given organization, we show the number of its server IPs (x-axis) and the number of ASes that host servers from






Figure 6: Heterogeneity of organizations and ASes

that organization (y-axis).³ We observe that operating a highly diverse infrastructure is commonplace in today’s Internet and is not limited to only the Internet’s biggest players, but reporting on the

³While in a few isolated cases, the ASes that host servers from a given organization are part of that organization (e. g., see [24]), hand-checking the 143 organizations with more than 1000 servers confirmed that in almost all cases, these ASes are genuine third-party networks that are run and operated independently from the organization whose servers they host.

bewildering array of scenarios we encountered when examining the extent of the different organizations and the networks they partner with is beyond the scope of this paper.


The realization that many organizations operate a server infrastructure that is spread across many ASes implies that the complementary view must be equally bewildering in terms of diversity or heterogeneity. This view is captured in Figure 6(a) by focusing on, say AS1, and examining how many third-party networks host some of their servers inside that AS. Thus, yet another way to quantify how much closer that cartoon figure is to reality than the traditional AS-level view with its implicit homogeneity assumption concerning the administrative authority of servers hosted within an AS is shown in Figure 6(c). Each dot in this figure represents an AS, and the number of organizations a given AS hosts is given on the y-axis while the number of identified server IPs is shown on the x-axis. As before, the figure only shows organizations with more than 10 servers. We observe that many ASes host a sizable number of server IPs that belong to many organizations; there are more than 500 ASes that host servers from more than five organizations, and more than 200 ASes that support more than 10 organizations.

Indeed, this observation is again fully consistent with public announcements [1, 6] and content providers’ efforts [5] to install their own brand of single-purpose CDNs inside various ISPs. The end effect of such developments is a clear trend towards more heterogeneous eyeball ISP networks by virtue of such ASes hosting more servers from an increasing number of interested third-party networks. In view of similar announcements from key companies such as Google [51, 25, 34], Amazon [2], or Facebook [4], the challenges of studying, leave alone controlling, such increasingly intertwined networks and traffic are quickly becoming daunting. As an example, consider a large Web hosting company (AS36351), for which we identified more than 40K server IPs belonging to a total more than 350 different organizations (highlighted in Figure 6(c) with a square).


5.3 New reality (II): Links are heterogeneous

In Section 5.2, we show that organizations take advantage of network diversity and purposefully spread their infrastructure across multiple networks. This development creates very fluid and often transparent network boundaries, which in turn causes havoc when trying to attribute the right traffic to the right network. The issues are illustrated in the cartoon Figure 7(a). The figure shows the traditional AS-level perspective, whereby Akamai is a member AS (AS1) of this IXP, and so are a generic AS3 and another generic (non-Akamai) AS2, and the Akamai AS peers at this IXP with AS3 which, in turn, peers also with AS2. This traditional AS perspective is enriched with member-specific details that specify that there is an Akamai server behind/inside (non-Akamai) AS2 and behind/inside the Akamai AS. Note that in terms of the traditional AS-level view, the question of how much Akamai traffic is seen at this IXP is clear-cut and can be simply answered by measuring the traffic on the peering link between AS3 and the Akamai AS. However, when accounting for the fact that there is an Akamai server behind/inside the non-Akamai member AS2, answering that same question becomes more involved. It requires measuring the traffic on the (Akamai) peering link between AS3 and the Akamai AS as well as accounting for the Akamai traffic on the (non-Akamai) peering link between AS3 and (non-Akamai) AS2.


Clearly, for accurately attributing traffic to the responsible parties in today’s network, the trend towards network heterogenization creates problems for the traditional AS-level view of the Internet. To illustrate the extent of these problems, we show in Figure 7(b) what we observe at the IXP for Akamai. Recall that Aka-



(a) Observing traffic from a direct and a non direct link of Akamai



(b) Perc. of Akamai traffic vs. perc. of Akamai traffic via direct link



(c) Perc. of CloudFlare traffic vs. perc. of CloudFlare traffic via direct link

Figure 7: AS link heterogeneity: Traffic via direct member link relative to other member links.

mai (AS20940) is a member of the IXP and peers with some 400 other member ASes. In the traditional view, accounting for Akamai traffic traversing the IXP simply means capturing the traffic on all the peering links between Akamai and those member ASes. Unfortunately, this simple view is no longer reflecting reality when Akamai servers are hosted inside or “behind” (non-Akamai) IXP member ASes. To capture this aspect, Figure 7(b) shows for each IXP member that peers with Akamai (indicated by a dot) the percentage of Akamai traffic on the direct peering link to Akamai (x -axis) vs. the percentage of total Akamai-server traffic for this member AS (y -axis). Under the traditional assumption, all dots would be stacked up at $x=100$, reflecting the fact that to account for Akamai-related traffic, all that is needed is to measure the Akamai peering links. However, with Akamai servers being massively deployed in third-party networks, including many of the other mem-

ber ASes of the IXP, we observe that some members get all their Akamai-related traffic from ASes other than the (member) Akamai AS ($x=0$), even when that traffic is sizable ($y>>0$). Moreover, the scattering of dots across Figure 7(b) succinctly captures the diverse spread of traffic across the direct peering link vs. the other member links. In terms of numbers, Akamai sends 11.1% of its traffic not via its peering links with the member AS. Put differently, traffic from more than 15K out of the 28K Akamai servers that we identified in our IXP data is seen at the IXP via non-IXP member links to Akamai. The same holds true for other major CDNs but also for relatively new players such as CloudFlare. Figure 7(c) shows the same kind of plot as Figure 7(b) for CloudFlare. It demonstrates that despite adhering to very different business models (i. e., Akamai deploys servers inside ISPs vs. CloudFlare operates its own data centers), the two CDNs have similar usage patterns as far as their peering links are concerned.

Looking beyond Akamai, we observe that different services from the same organization use their servers differently resulting in different usage patterns of the peering links. For example, for Amazon CloudFront, Amazon’s “CDN part”, almost all traffic is sent via the IXP’s Amazon links. However, for Amazon EC2, the “cloud part”, a sizable fraction comes via other IXP peering links. We also noticed that for most cases where we see the use of the non-IXP member links, the percentage of traffic in those links increases during peak times. This may be due to reasons such as load balancing, performance improvement, or cost savings. Lastly, how our view of the usage of the IXP’s public peering links is impacted by private peerings that may be in place between member ASes of the IXP remains unexplored.

Summary: To illustrate the kind of benefits that arise from having access to a global Internet vantage point in the form of our large European IXP, we confirm a feature of today’s Internet that is well-known among experts but remains largely under-reported in the networking research literature—a tendency of certain Internet players to either host servers from third-party networks within their own network infrastructures or deploy their own servers in strategically-chosen third-party ASes. More importantly, we present a methodology for discovering an organization’s servers, whether they are deployed within the organization’s own AS (or ASes) or inside some third-party network’s infrastructure, and use it to systematically assess the *extent* of this network heterogenization and study its impact on the usage of peering links at IXPs by these increasingly more heterogeneous member ASes. However, we want to stress that our AS-links usage-related findings are not IXP-specific (i. e., public peering links), but apply to any AS-link in the Internet, pointing towards serious challenges when trying to attribute the right traffic to the right party.

6. DISCUSSION AND CAVEATS

We are not the first to try and uncover the footprints of the infrastructures of commercial Internet players. One group of prior studies targets specific Internet companies (e. g., Akamai [54, 39, 52], Youtube [11, 37, 26], Netflix [10]), or one click hosters [19]). Other work is more concerned with inferring Web hosting infrastructures by relying on content only [14]. Our approach differs from these earlier works. For one, we rely on a unique vantage point in the form of one of the largest European IXPs to supply us with a weekly pool of some 230M IPs from which we diligently extract some 1.5M server IPs. Next, we rely exclusively on publicly available data⁴ to group these servers by organizations that have

⁴Note that our use of the set of DNS resolvers from a large commercial CDN in Section 2.3 is a shortcut. A similar list could also

the administrative authority over them and are responsible for their content. In doing so, we are inspired by earlier studies such as [20, 48]. Lastly, the methodology we develop for grouping servers by their organization is general in the sense that it applies equally well to content providers, CDNs, hosting companies, cloud infrastructure providers, eyeball ASes, or other Internet players.

The difference in perspective between the more traditional AS-level view of the Internet and our perspective that centers around organizations and companies and their heterogeneously deployed server-based infrastructure becomes evident when comparing our approach to the recent work by Cai et al. [24] on mapping ASes to organizations. For one, the starting point for [24] is the traditional AS-level view of the Internet, and two ASes are grouped into two different organizations if neither of the organizations is a subsidiary of the other (i. e., majority-owned by the other). While such a top-down ownership-based grouping of ASes captures one aspect of how ASes are inter-related, it is oblivious to how network infrastructures get used and deployed in today's Internet. In particular, while the method described in [24] may succeed in clustering all Akamai-owned ASes under the umbrella organization Akamai, the publicly known fact that Akamai has more than 100K servers deployed in hundreds of different third-party non-Akamai ASes [46] cannot be accounted for at all by that approach.

Our work relies critically on the sFlow records provided by one of the largest IXPs in Europe, and it can be argued that for many researchers, access to such data cannot be taken for granted. However, it is important to note that some of these largest IXPs in Europe generally welcome collaborations with researchers and are supportive of research efforts that make explicit use of their data (see for instance [18]). Once access to data collected from such unique and powerful vantage points is established, the opportunities for researchers are plentiful.

After presenting evidence for the kind of visibility into the Internet that comes with using one of these largest European IXPs as a vantage point, we highlight in this paper some of the benefits that arises from having access to such a vantage point. However, despite its impressive capabilities, our IXP and the measurements it collects can only tell us so much about the network's "state", and many important issues remain concerning our knowledge about what exactly we can and cannot discern about the Internet as whole and its individual constituents. While we have identified a number of "blind spots", much remains to be done in terms of identifying and collecting IXP-external information that can be brought to the table for either checking, validating, or refining the findings obtained from the use of IXP-internal data only. The question of how to appropriately fuse selective IXP-external data with IXP-internal measurements to obtain a picture of the global network and its individual constituents that is unprecedented in terms of its accuracy, details, and insight looms as an important open research problem.

7. CONCLUSION

This paper contributes to Internet measurements by reporting on the existence of single, well-localized physical locations or vantage points within the Internet infrastructure where one can "see" much of the global Internet. Mining the data collected at one such vantage point reveals a network that teems with heterogeneity whichever way one looks. Given that economic incentives drive many of the main commercial Internet players to either host third-party servers in their own network infrastructures or deploy their own servers, often in massive numbers, in strategically selected (close to the

have been obtained by relying on publicly available data only [52, 54, 39], e. g., via active scanning or from DNS logs.

end users) third-party networks, we expect the observed trend towards increasingly more heterogeneous networks and increasingly diverse usage of IXP peering links, in particular, and AS-links, in general, to accelerate, especially in view of the growing importance of cloud providers. As an interesting consequence of more servers being deployed close to the end users, we also expect that IXPs in the future will "see" less end user-to-server traffic but an increasing amount of server-to-server traffic.

In response to this observed heterogeneity, the paper also contributes to Internet topology research by advancing a new mental model for the Internet's ecosystem that accounts for the observed network heterogenization, points towards measurements that reveal and keep track of this ongoing heterogenization process, and is rich and flexible enough to adapt to a constantly changing Internet environment. Doing so only scratches the surface of a new and rich problem space, and our efforts reported in this paper that focus less on the Internet's connectivity structure and more on how traffic flows over this connectivity structure are just a first step towards exploring that space.

8. REFERENCES

- [1] Akamai and AT&T Forge Global Strategic Alliance to Provide Content Delivery Network Solutions. http://www.akamai.com/html/about/press/releases/2012/press_120612.html.
- [2] Amazon CloudFront - Amazon Web Services. <http://aws.amazon.com/cloudfront/>.
- [3] CIDR Report. <http://www.cidr-report.org/>.
- [4] Like Netflix, Facebook is boosting its edge network. <http://gigaom.com/2012/06/21/like-netflix-facebook-is-planning-its-own-cdn/>.
- [5] Netflix Open Connect. <https://signup.netflix.com/openconnect>.
- [6] Orange and Akamai form Content Delivery Strategic Alliance. http://www.akamai.com/html/about/press/releases/2012/press_112012_1.html.
- [7] RIPE RIS. <http://www.ripe.net/ris/>.
- [8] Route Views Project, University of Oregon. <http://www.routeviews.org/>.
- [9] Team Cymru. <http://www.team-cymru.org/>.
- [10] V. Adhikari, Y. Guo, F. Hao, M. Varvello, V. Hilt, M. Steiner, and Z.-L. Zhang. Unreeling Netflix: Understanding and Improving multi-CDN Movie Delivery. In *IEEE INFOCOM*, 2012.
- [11] V. K. Adhikari, S. Jain, Y. Chen, and Z.-L. Zhang. Vivisecting YouTube: An Active Measurement Study. In *IEEE INFOCOM*, 2012.
- [12] P. Aditya, M. Zhao, Y. Lin, A. Haeberlen, P. Druschel, B. Maggs, and B. Wishon. Reliable Client Accounting for Hybrid Content-Distribution Networks. In *NSDI*, 2012.
- [13] B. Ager, N. Chatzis, A. Feldmann, N. Sarrar, S. Uhlig, and W. Willinger. Anatomy of a Large European IXP. In *ACM SIGCOMM*, 2012.
- [14] B. Ager, W. Mühlbauer, G. Smaragdakis, and S. Uhlig. Web Content Cartography. In *ACM IMC*, 2011.
- [15] Akamai. Facts and Figures: Network Deployment. http://www.akamai.com/html/about/facts_figures.html.
- [16] Amazon. AWS Dashboard. <http://status.aws.amazon.com/>.
- [17] Amazon. EC2 Public IP ranges. <https://forums.aws.amazon.com/ann.jspa?annID=1701>.

- [18] AMS-IX. AMS-IX hosts BGP-Mux research project. <https://www.ams-ix.net/newsitems/82>.
- [19] D. Antoniadis, E. Markatos, and C. Dovrolis. One-click Hosting Services: A File-Sharing Hideout. In *ACM IMC*, 2009.
- [20] I. Bermudez, M. Mellia, M. Munafà, R. Keralapura, and A. Nucci. DNS to the Rescue: Discerning Content and Services in a Tangled Web. In *ACM IMC*, 2012.
- [21] Facebook blog. A Continued Commitment to Security. <http://www.facebook.com/blog/blog.php?post=486790652130>.
- [22] Google Search Blog. Making Search More Secure. <http://googleblog.blogspot.com/2011/10/making-search-more-secure.html>.
- [23] Netflix Nordics blog. Netflix Launches Today in Sweden, Denmark, Norway, Finland. <http://nordicsblog.netflix.com/2012/10/>.
- [24] X. Cai, J. Heidemann, B. Krishnamurthy, and W. Willinger. Towards an AS-to-Organization Map. In *ACM IMC*, 2010.
- [25] M. Calder, X. Fan, Z. Hu, E. Katz-Bassett, J. Heidemann, and R. Govindan. Mapping the Expansion of Google's Serving Infrastructure. In *ACM IMC*, 2013.
- [26] M. Cha, H. Kwak, P. Rodriguez, Y. Y. Ahn, and S. Moon. I Tube, You Tube, Everybody Tubes: Analyzing the World's Largest User Generated Content Video System. In *ACM IMC*, 2008.
- [27] H. Chang, S. Jamin, Z. M. Mao, and W. Willinger. An Empirical Approach to Modeling Inter-AS Traffic Matrices. In *ACM IMC*, 2005.
- [28] N. Chatzis, G. Smaragdakis, and A. Feldmann. On the Importance of Internet eXchange Points for Today's Internet Ecosystem. <http://arxiv-web3.library.cornell.edu/abs/1307.5264v2>.
- [29] K. Cho, C. Pelsser, R. Bush, and Y. Won. The Japan Earthquake: the Impact on Traffic and Routing Observed by a Local ISP. In *In ACM SWID*, 2011.
- [30] Cisco. Visual Networking Index (VNI) and Forecast. http://www.cisco.com/en/US/netsol/ns827/networking_solutions_sub_solution.html.
- [31] Private communication.
- [32] D. Dainotti, C. Squarcella, E. Aben, K.C. Claffy, M. Chiesa, M. Russo, and A. Pescapé. Analysis of Country-Wide Internet Outages Caused by Censorship. In *ACM IMC*, 2011.
- [33] J. Erman, A. Gerber, M. Hajiaghayi, D. Pei, and O. Spatscheck. Network-aware Forward Caching. In *WWW*, 2009.
- [34] T. Flach, N. Dukkipati, A. Terzis, B. Raghavan, N. Cardwell, Y. Cheng, A. Jain, S. Hao, E. Katz-Bassett, and R. Govindan. Reducing Web Latency: the Virtue of Gentle Aggression. In *ACM SIGCOMM*, 2013.
- [35] Mozilla Foundation. Publicsuffix.org. <http://publicsuffix.org/>.
- [36] A. Gerber and R. Doverspike. Traffic Types and Growth in Backbone Networks. In *OFC/NFOEC*, 2011.
- [37] P. Gill, M. F. Arlitt, Z. Li, and A. Mahanti. Youtube Traffic Characterization: A View From the Edge. In *ACM IMC*, 2007.
- [38] Google. What is 1e100.net? <http://support.google.com/bin/answer.py?hl=en&answer=174717>.
- [39] C. Huang, A. Wang, J. Li, and K. Ross. Measuring and Evaluating Large-scale CDNs. In *ACM IMC*, 2008.
- [40] Data Center Knowledge. Who Has the Most Web Servers? <http://www.datacenterknowledge.com/archives/2009/05/14/whos-got-the-most-web-servers>.
- [41] C. Labovitz, S. Iekel-Johnson, D. McPherson, J. Oberheide, and F. Jahanian. Internet Inter-Domain Traffic. In *ACM SIGCOMM*, 2010.
- [42] T. Leighton. Improving Performance on the Internet. *Commun. ACM*, 52(2):44–51, 2009.
- [43] G. Maier, A. Feldmann, V. Paxson, and M. Allman. On Dominant Characteristics of Residential Broadband Internet Traffic. In *ACM IMC*, 2009.
- [44] Maxmind. GeoLite Country. <http://dev.maxmind.com/geoip/legacy/geolite>.
- [45] Netcraft. January 2013 Web Server Survey. <http://news.netcraft.com/archives/2013/01/07/january-2013-web-server-survey-2.html>.
- [46] E. Nygren, R. K. Sitaraman, and J. Sun. The Akamai Network: A Platform for High-performance Internet Applications. *SIGOPS Oper. Syst. Rev.*, 2010.
- [47] V. Paxson. Bro: A System for Detecting Network Intruders in Real-Time. In *Usenix Security Symposium*, 1998.
- [48] D. Plonka and P. Barford. Flexible Traffic and Host Profiling via DNS Rendezvous. In *SATIN*, 2011.
- [49] I. Poese, S. Uhlig, M. A. Kaafar, B. Donnet, and B. Gueye. IP Geolocation Databases: Unreliable? *ACM CCR*, 2011.
- [50] Sandvine. Global Internet Phenomena Report. http://www.sandvine.com/news/global_broadband_trends.asp.
- [51] F. Streibelt, J. Böttger, N. Chatzis, G. Smaragdakis, and A. Feldmann. Exploring EDNS-Client-Subnet Adopters in your Free Time. In *ACM IMC*, 2013.
- [52] A. Su, D. Choffnes, A. Kuzmanovic, and F. Bustamante. Drafting Behind Akamai. In *ACM SIGCOMM*, 2006.
- [53] S. Triukose, Z. Al-Qudah, and M. Rabinovich. Content Delivery Networks: Protection or Threat? In *ESORICS*, 2009.
- [54] S. Triukose, Z. Wen, and M. Rabinovich. Measuring a Commercial Content Delivery Network. In *WWW*, 2011.