

Web Content Cartography

Bernhard Ager
T-Labs/TU Berlin
bernhard@net.t-labs.tu-berlin.de

Georgios Smaragdakis
T-Labs/TU Berlin
georgios@net.t-labs.tu-berlin.de

Wolfgang Mühlbauer
ETH Zurich
muehlbauer@tik.ee.ethz.ch

Steve Uhlig
T-Labs/TU Berlin
steve@net.t-labs.tu-berlin.de

ABSTRACT

Recent studies show that a significant part of Internet traffic is delivered through Web-based applications. To cope with the increasing demand for Web content, large scale content hosting and delivery infrastructures, such as data-centers and content distribution networks, are continuously being deployed. Being able to identify and classify such hosting infrastructures is helpful not only to content producers, content providers, and ISPs, but also to the research community at large. For example, to quantify the degree of hosting infrastructure deployment in the Internet or the replication of Web content.

In this paper, we introduce *Web Content Cartography*, i. e., the identification and classification of content hosting and delivery infrastructures. We propose a lightweight and fully automated approach to discover hosting infrastructures based only on DNS measurements and BGP routing table snapshots. Our experimental results show that our approach is feasible even with a limited number of well-distributed vantage points. We find that some popular content is served exclusively from specific regions and ASes. Furthermore, our classification enables us to derive content-centric AS rankings that complement existing AS rankings and shed light on recent observations about shifts in inter-domain traffic and the AS topology.

Categories and Subject Descriptors

C.2.5 [Computer-Communication Networks]: Local and Wide-Area Networks—*Internet*

General Terms

Measurement

Keywords

Content delivery, hosting infrastructures, measurement, DNS

The measurement traces are available from
<http://www.inet.tu-berlin.de/?id=cartography>

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. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

IMC'11, November 2–4, 2011, Berlin, Germany.

Copyright 2011 ACM 978-1-4503-1013-0/11/11 ...\$10.00.

1. INTRODUCTION

Today's demand for Web content in the Internet is enormous, reflecting the value Internet users give to content [18]. Recent traffic studies [15, 12, 22, 27] show that Web-based applications are again very popular. To cope with this demand, Web-based applications and Web content producers use scalable and cost-effective hosting and content delivery infrastructures. These infrastructures, which we refer to as *hosting infrastructures* throughout this paper, have multiple choices on how and where to place their servers.

Leighton differentiates between three options for Web content delivery [24]: (i) centralized hosting, (ii) data-center-based content distribution network (CDN), (iii) cache-based CDNs. Approaches (ii) and (iii) allow to scale content delivery by distributing the content onto a dedicated hosting infrastructure. This hosting infrastructure can be composed of a few large data-centers, a large number of caches, or any combination. In many cases, DNS is used by the hosting infrastructure to select the server from which a user will obtain content [20, 37, 7, 30].

The deployment of hosting infrastructures is dynamic and flexible in multiple ways, e.g.: increasing the size of the existing hosting infrastructure, changing peerings with ISPs, placing parts of the infrastructure inside ISP networks. Therefore, being able to identify and classify hosting infrastructures in an automated manner is a step towards understanding this complex ecosystem, and an enabler for many applications. Content producers can benefit from understanding the footprint of hosting infrastructures to place content close to their customer base. For CDNs, a map of hosting infrastructures can assist them in improving their competitiveness in the content delivery market. For ISPs, it is important to know which hosting infrastructures deliver a specific content and at which locations to make relevant peering decisions. The research community needs a better understanding of the evolving ecosystem of hosting infrastructures, given its importance as a driver in the evolution of the Internet.

As demand drives hosting infrastructures to make a given content available at multiple locations, identifying a particular hosting infrastructure requires sampling its location diversity. Previous work has attempted to discover specific hosting infrastructures in an extensive manner, e. g., Akamai [36, 35, 17]. Such studies rely on the knowledge of a signature that identifies the target infrastructure, e. g., CNAMEs in DNS replies or AS numbers. Labovitz et al. [22] inferred that a small number of hosting infrastructures are responsible for a significant fraction of inter-domain traffic. Unfortunately, this study observes only the traffic crossing AS boundaries, not traffic delivered directly from inside the monitored ISPs. As a consequence, important CDNs such as Akamai as well as data-centers deployed inside ISP networks are under-represented.

In this paper, we introduce *Web Content Cartography*, i. e., the

identification and classification of hosting infrastructures. We propose a lightweight and fully automated approach to discover hosting infrastructures based on DNS measurements and BGP routing table snapshots. Compared to previous work, our method is able to identify and classify new as well as existing hosting infrastructures without the need of a priori knowledge of their operation or deployment. To achieve such a degree of generality, we rely on the information that hosting infrastructures expose to end-users when requesting hostnames through DNS. We construct mappings between requested hostnames and IP addresses returned, and cluster the hostnames into hosting infrastructures with the help of network information such as IP addresses, prefixes and AS numbers.

Our contributions can be summarized as follows:

- *Identification of hosting infrastructures:* We propose a lightweight and fully automated approach to discover hosting infrastructures, based on DNS measurements and BGP routing table snapshots.
- *Classification of hosting infrastructures:* We classify individual hosting infrastructures and their different deployment strategies based on their network and location footprint.
- *Content replication:* We quantify the degree of content replication in the Internet and its impact on local content availability in different regions of the world. We introduce the *content monopoly index* that reflects the content an organization hosts, either replicated or exclusively hosted.
- *Revisiting AS rankings:* We derive content-centric AS rankings that complement existing AS rankings and shed light on recent observations about shifts in inter-domain traffic and the AS topology.

The remainder of this paper is structured as follows. We present our methodology in Section 2 and discuss our measurements in Section 3. In Section 4, we provide our results, and discuss the implications of our work in Section 5. We present related work in Section 6 and summarize the paper in Section 7.

2. METHODOLOGY

In this section we describe our approach to identify and classify hosting infrastructures in the Internet. The key idea is to collect the IP addresses that DNS returns for various popular and unpopular hostnames from geographically diverse vantage points. We use this information for several purposes: (i) to find the geographic location where popular content is available from, (ii) to find the network locations, e. g., prefixes and ASes, where content is available, and (iii) to find out by which hosting infrastructure a hostname is served.

2.1 Design Goals

To achieve our goals of mapping content and identifying hosting infrastructures, we design measurements tailored to our specific needs: (i) we target the hosting infrastructures that host content and (ii) we sample the network footprint of each of these hosting infrastructures in order to be able to classify them and study aspects such as content replication. We now elaborate on our choices and explain why they ensure that our measurements allows us to achieve the above goals.

Hosting Infrastructure Coverage.

To satisfy the first requirement, i. e., obtaining a wide coverage of popular hosting infrastructures in terms of traffic volume, one approach is to sample all possible hostnames. However, due to the sheer size of the Web – an estimated 92 million active domains

only for the .COM top-level domain [6] – querying all host names in the Internet would be way too cumbersome from a measurement perspective. Fortunately, there is high variation in the popularity of Web content. Given that Internet traffic at various levels of aggregation is consistent with Zipf’s law [13, 40, 38, 10], the hosting infrastructures that serve popular hostnames are likely to be responsible for a major part of today’s Internet traffic. Despite a lack of definitive figures about how many hosting infrastructures are responsible for most of the Web traffic, we believe that it is reasonable to assume that a limited number of highly popular Web sites is sufficient to cover the hosting infrastructures responsible for the majority of the HTTP traffic in the Internet. For example, Akamai claims to deliver about 20 % of the total Web traffic in the Internet [30]. Labovitz et al. [22] attribute up to 10 % of all Internet traffic to Google, more than 15 % to the top 10 hosting infrastructures and more than 40 % to the top 100.

Network Footprint.

The second goal—sampling the network footprint of hosting infrastructures—asks for measurements from multiple vantage points. By running measurements from vantage points that reside in different networks and countries, we benefit from the way hosting infrastructures use DNS to select the server from which a user obtains the requested content [20, 35, 36, 7]. CDNs rely on the network location of the recursive DNS resolver to determine the IP address returned by DNS [30, 28, 37]. In many cases, the hosting infrastructure assumes that the DNS resolver is close to the client and optimizes based on this assumption. Therefore, to sample the locations from which a given hosting infrastructure serves content, our approach relies on volunteers to sample from different networks, ASes, and countries around the world.

2.2 Network Features

The way hosting infrastructures are deployed in the Internet is not homogeneous. In Section 2.3 we leverage the “network footprint” of hosting infrastructures to map them. Now, we discuss features that can be used to distinguish between different deployment strategies of hosting infrastructures.

To this end, we extract the IP addresses obtained within the DNS replies, from geographically dispersed vantage points. The set of IP addresses returned for a particular hostname reveals the degree to which the corresponding hosting infrastructure is network-wise and geographically distributed. Hence, the natural choice for our features are *prefix*, *AS* and *location* of an IP address. For example, small data-centers will be located within a single AS in a single geographic location, having a limited number of /24 subnetworks, and a large number of IP addresses. A massively distributed CDN will rely on multiple ASes. Evidently, these features are correlated, and potentially differ in their power to discriminate between different types of hosting infrastructures. We leave this for further investigation, and prefer to rely on all features for now.

Throughout the paper, we rely on both the granularity of BGP prefixes as well as /24 subnetworks. /24 subnetworks have the advantage of better representing the actual usage of the address space by highly distributed hosting infrastructures such as Akamai. BGP prefixes on the other hand indicate at which granularity routing is performed and more closely match the address space usage of centralized hosting infrastructures such as data-centers.

To determine the AS for a given IP address, we use BGP routing information from RIPE RIS [4] and RouteViews [33], and assume that the last AS hop in an AS path reflects the origin AS of the prefix.

To infer the geographical location of an IP address, we rely on




Figure 1: High level view of our approach.

the Maxmind geolocation database [29]. We are aware that geolocation databases suffer from limited accuracy. However, they have been shown to be reliable at the country-level [32].

2.3 Clustering Algorithm

Our goal is to detect where hosting infrastructures are located by AS and country, and classify them according to their network footprint. The key idea is to cluster all hostnames that are served by the same hosting infrastructure. To this end, we rely on the assumption that each hostname is served by a single hosting infrastructure.

We are aware that counter-examples exist: Meebo, an instant messenger aggregator, which is running its own meta-CDN, distributes the content demand across different CDNs by using a DNS server under their control. Another case is Netflix, which offers video-on-demand streams and relies on both Level 3 and Limelight. Our approach accommodates such counter-examples by putting the respective hostnames into separate clusters.

By examining the resolved CNAME records for various hostnames, it is sometimes already possible to identify the CDN that delivers the respective content. For example, a CNAME to `akamai.net` clearly points to the Akamai CDN. Yet, finding content providers would require an extensive a-priori database. In addition, some CDNs do not use CNAMEs, and CNAMEs are also used in different contexts than CDNs. In contrast, our clustering approach achieves the important goal of identifying hosting infrastructures in the wild, and could be used to help build such a database. Moreover, our agnostic approach is able to separate hosting infrastructures if they are actually maintained by the same administrative entity (e.g., the Akamai CDN), but treat various hostnames differently in terms of replication.

In the rest of the section we present our algorithm that identifies hosting infrastructures based on our data. We choose a two-step algorithm as depicted in Figure 1. During the first step, we ensure that the prominent hosting infrastructures are identified. It also gives an upper bound on the size of the clusters. In the second step, the algorithm merges clusters that share network features. The first step prevents the second one from clustering small hosting infrastructures with large ones. This may happen for example when infrastructures share address space with others.

Step 1: Separating Large Hosting Infrastructures.

The goal here is to separate large hosting infrastructures from the rest. We rely on three network-based features: (i) the number of IP addresses, (ii) the number of /24 networks and (iii) the number of ASes a hostname is resolved to. We use the k -means algorithm [26] to partition the hostnames in up to k clusters in the feature space. The choice of the value of k is discussed at the end of this subsection. Clusters whose features have high values relate to widely-deployed infrastructures. On the other hand, smaller infrastructures that use very few /24 subnetworks and IP addresses are not sufficiently different, and therefore, can be found in the same cluster. Increasing the value of k in the clustering algorithm does not lead to improvements, as the feature space simply does not allow to differentiate them.

Step 2: Distinguishing Small Hosting Infrastructures.

The pre-clustering of hostnames in Step 1 does not take into account the actual network locations from where content is served, but only features that reflect the *size* of the hosting infrastructures. The goal of the second step is to build sub-clusters within each k -means cluster by identifying the hostnames that are hosted on similar network locations in terms of IP address space. To this end, we take into account the set of BGP prefixes the hostname maps to. Based on the similarity between the sets of prefixes of two similarity-clusters, we decide if they belong to the same hosting infrastructure, and if so we merge these clusters. For this, we define the similarity between two sets s_1 and s_2 as follows:

$$\text{similarity}(s_1, s_2) = 2 \cdot \frac{|s_1 \cap s_2|}{|s_1| + |s_2|} \quad (1)$$

where $|\cdot|$ denotes the size of the set. The purpose of the factor 2 is to stretch the image of the similarity function to the interval $[0, 1]$.

The second step of the algorithm is performed for each k -means cluster separately. Initially, we put each hostname contained in the current k -means cluster into its own sub-cluster, called a similarity-cluster. Then, we perform a pairwise comparison of all similarity-clusters of the current k -means cluster and merge them according to their similarity. We iterate the process until convergence to a fixed point. At this stage, each similarity-cluster identifies all hostnames used by a single content delivery infrastructure.

Tuning.

Generally, choosing k too high will lead to split large hosting infrastructures into smaller clusters, while choosing it too low may result in significant overlap between hosting infrastructures. As part of our case study in Section 4 we examine how sensitive the outcome of our clustering is to the choice of k . We find that the whole interval $20 \leq k \leq 40$ provides reasonable and similar results according to our verification, and therefore we decide to choose $k = 30$. Extensive tests reveal that merging thresholds of 0.7 on the similarity between two similarity-clusters work well for the second phase of the algorithm. We leave it for future work to advance our clustering techniques and to optimize the choice of parameters, and rather focus on the methodology for analyzing hosting infrastructures in the following.

2.4 Metrics: Content Potential and Monopoly

Now, we propose metrics and rankings that allow us to compare the obtained hosting infrastructures. We will use these metrics and corresponding rankings later (Section 4) to get insight about the geographic properties of content replication as well as the role of different organizations in the hosting of Web content, e.g., which organizations exclusively host content.

Content Delivery Potential.

Our goal is to provide intuition on the amount of content that is available for download in a geographic region (e.g., country, continent) or an AS. To this end, we define the *content delivery potential* as the fraction of hostnames that can be served from either a geo-

graphic region or an AS. Values close to 1 suggest that a major part of popular content is available locally. The disadvantage of the content delivery potential is that replicated content is counted as many times as there are different locations where it is hosted, introducing a bias in favor of replicated content.

Normalized Content Delivery Potential.

Therefore, we introduce the *normalized content delivery potential*, which takes into account the total number of locations from where content is available. Intuitively, a location does not exclusively deliver content, if the content has been replicated to a large number of locations. To take this into account, we calculate the normalized content delivery potential of a hostname as follows. First, we determine the weight of a hostname, which is 1 divided by the number of all hostnames. Second, we check how many different ASes, subnetworks, or regions serve this count, henceforth referred to as *replication-count*. To assess the contribution of a particular hostname to the normalized content delivery potential of an AS, subnetwork, or region, we take the weight from the first step and divide it by replication-count. The benefit of the normalized content delivery potential is a more balanced ranking in terms of hosted content, as it spreads the weight of distributed content infrastructure across all ASes, regions, or subnetworks that serve their hosted content.

Content Monopoly Index.

To distinguish between locations (ASes, geographic regions) that have exclusive content and those that host replicated content, we introduce the *Content Monopoly Index* (CMI). We define it as the ratio between the normalized content potential and the non-normalized content potential. An AS with a large CMI hosts a large number of hostnames that are not available in another AS.

3. MEASUREMENTS

In this section we present our approach to collect traces, i. e., active DNS measurements, in order to evaluate our methodology. To achieve our goal of identifying hosting infrastructures we compile a list of diverse hostnames and analyze DNS traces when resolving these hostnames as collected by end-users in commercial ISPs. Our experimental results advocate that our methodology is able to identify a significant fraction of hosting infrastructures network footprints, even with a small number of well-distributed vantage points.

3.1 Hostname Selection

To obtain a good coverage of the largest hosting infrastructures, we decide to include in our hostname list the top ranked ones according to Alexa [1]. Alexa relies on statistical sampling and determines its ranking by counting how many pages were visited by Internet users who have downloaded their toolbar. Note, Alexa itself is already accounting for various sampling biases of its user list¹. In order to check for potential differences and to scrutinize replication of content also for less popular hosts, we further add hosts that are at the bottom of Alexa's ranking.

Moreover, many web-pages contain embedded content, e. g., images, videos, and advertisements that the browser of the user has to download from different servers. In our study, such embedded content has to be taken into account, as it might be served from servers other than those serving the front page of a popular hostname. To give an example, the front page of `facebook.com` is served from Facebook datacenters, but the logo and other embedded objects such as the profile photo is served from the Akamai

content distribution network. In addition, to increase the chance of detecting the relevant infrastructures, we extracted hosts that are likely to be hosted on hosting infrastructures from the ranks 2001 to 5000 of the Alexa list. We identify such hosts by checking if they have CNAME records in their DNS answers.

Overall, we keep the 2,000 most popular and 2,000 from the least popular hostnames according to the Alexa ranking. Moreover, we include more than 3,400 embedded hostnames and 840 hostnames because of CNAMEs. This list leads to four subsets which we will refer to as TOP2000, TAIL2000, EMBEDDED, and CNAMES, respectively, for the remainder of the paper. Note, that several hostnames are used to deliver both embedded objects as well as popular websites. This leads to an overlap of 823 hostnames between TOP2000 and EMBEDDED.

3.2 Measurement Approach

Our measurement approach relies on volunteers to run a program that is publicly available for download on our project web-page along with instructions on how to perform the experiment. We initially announced the project during the IMC 2010 conference. In addition, we made a public call in several DNS-related mailing lists and invited friends to participate. This resulted in a total of 484 traces. Participants in our study run the program locally on their end-hosts. The program queries the locally configured DNS resolver, a Google Public DNS resolver and an OpenDNS resolver for the list of over 7400 hostnames, and stores the full DNS replies in a *trace* file. The traces collected with our program do not interact with any of the browsing or download history and activity of the user.

In addition to the DNS data, we collect meta-information that helps in sanitizing the measurement and in debugging. We report the Internet-visible IP address of the client every 100 DNS queries by asking a web-server running a custom script, and store information such as the operating system and the timezone, as well as the IP addresses of the DNS resolvers in use.

To sanitize the measurements, we check that the locally configured resolver is not a 3rd-party resolver such as Google Public DNS or OpenDNS. This information cannot always be derived from the resolver IP address, as the recursive resolver may hide behind a DNS forwarding resolver. Therefore, we perform queries for 16 additional names to domains under our administrative control. Their authoritative name servers are configured to reply to queries with the IP address of the querying resolver. This gives us the IP addresses of the resolvers directly, without having to further correlate logs from our authoritative name servers with the traces. To avoid receiving cached copies of the entries, we construct the names on-the-fly with the help of microsecond resolution timestamps and the Internet-visible IP address of the client.

The program is designed to collect all of the above data once every 24 hours, and write it to a trace file, until stopped by the user. This implies that there may be multiple traces per *vantage point*. We identify vantage points through the information contained in the trace files as well as meta information provided by the end-user running the program when uploading the trace files.

3.3 Data Cleanup

We perform a thorough cleanup process on the raw traces. We check for the following measurement artifacts.

We do not consider traces if the vantage point roams across ASes during the experiment, as we cannot determine the exact impact of the change. If the DNS resolver of the host returns an excessive number of DNS errors, or is unreachable, we do not consider the trace. If the DNS resolver of the host is a well-known third-party

¹http://alexa.com/help/traffic_learn_more

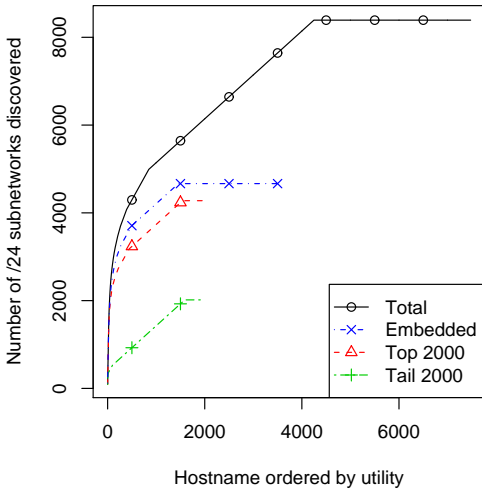


Figure 2: /24 subnetwork coverage by the hostname list.

resolver, e. g., OpenDNS, Google Public DNS, we do not consider the trace. We showed in previous work that using third-party resolvers introduces bias by not representing the location of the end-user [7].

When a vantage point provides us with repeated measurements over time, we only use the first trace that does not suffer from any other artifact to avoid over-representing a single vantage point. This is important to avoid bias when quantifying the content potential (cf. Section 2.4).

After removing all traces with the above artefacts, we have 133 clean traces that form the basis of this study Note, the cleanup process has limited impact on our hosting infrastructure coverage and sampling of the network footprint.

3.4 Data coverage

We next investigate the coverage that our hostnames and vantage points provide.

3.4.1 Network and Geographic Footprint of Vantage Points

We map the IP addresses of vantage points of the 133 clean traces to ASes and countries using the mapping methodology described in Section 2.2. This leads to a coverage of 78 ASes and 27 countries that span six continents. Our experiments include traces from major residential ISPs, e. g., AT&T Internet Services, Comcast, Verizon, Road Runner, Telefonica, Deutsche Telekom, British Telecom as well as smaller residential ISPs and some university and research networks.

3.4.2 Network Coverage by Hostname

Previous studies [17, 35, 36] were able to achieve an exhaustive coverage for a limited number of well known hosting infrastructures. In our study, we strive to achieve a wide coverage of the prevalent hosting infrastructures without targeting a-priori known hosting infrastructures. Thus, we investigate the scope of the network coverage of our study. For this, we analyze to which degree replies for different parts of our hostname list result in different network coverage. To identify the IP ranges utilized by hosting infrastructures we aggregate the returned IP addresses over /24 subnetworks. We argue that this is the right granularity as hosting infrastructures tend to deploy server clusters for resilience and load balancing. Aggregation on the prefix of the returned IP addresses

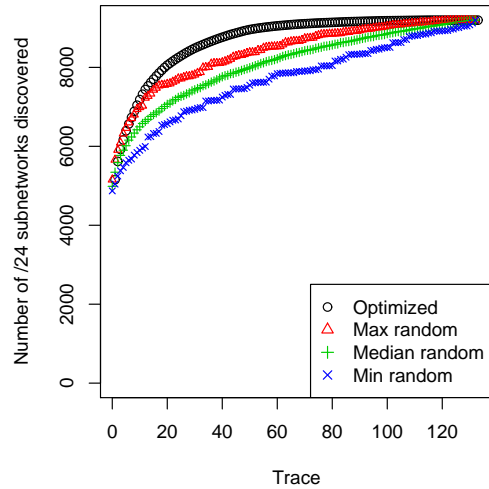


Figure 3: /24 subnetwork coverage by traces.

may lead to overestimation of the coverage, yet another indication why aggregation on /24 subnetworks is justified.

Figure 2 shows the total number of discovered /24 subnetworks when we stepwise add hostnames from our list (see Section 3.1) according to their utility. By utility we mean the number of new /24 subnetworks that are discovered by a hostname. The y-axis shows the total number of discovered /24 subnetworks as a function of the number of traces considered on the x-axis. In addition to the full list, we differentiate between three types of hostnames introduced in Section 3.1: TOP2000, TAIL2000, and EMBEDDED.

The curves in Figure 2 can be separated into three regions: a steep slope on the left, followed by a region with a slope of 1, and a flat region at the end. The steep slope region identifies hostnames with a high utility. These hostnames should be included to discover a significant fraction of the content infrastructure with limited probing effort. The region having a slope of 1 results from hostnames that positively contribute to the coverage but the utility is much lower than hostnames on the left. The third and flat region corresponds to hostnames that return redundant information about the hosting infrastructure, compared to the first two regions.

Let us now turn to the pairwise comparison of the three types of hostnames. While the hostname lists of TOP2000 and TAIL2000 are of equal size, the /24 subnetworks observed by TOP2000 and TAIL2000 exhibit a difference by a factor of more than two in the number of subnetworks they uncover. This unveils that popular content is served from more widely distributed hosting infrastructures than this of less popular content. Most of the difference in the cumulative utility between TOP2000 and TAIL2000 stems from a small number of popular hostnames. Furthermore, we observe that the hosting infrastructures that serve hostnames in EMBEDDED are well distributed.

To estimate the utility of additional hostnames we calculate the median utility of 100 random hostname permutations. We find that when adding the last 200 hostnames, the average utility is 0.65 /24 subnets per hostname, and 0.61 /24 subnets when adding the last 50 hostnames.

3.4.3 Network Coverage by Trace

Hosting infrastructures rely on geographic hints to serve content from servers close to the end user [24, 30]. We expect that traces in diverse regions of the world sample different parts of the host-




Figure 4: CDF of similarities for answers across different traces, for different sets of hostnames.

ing infrastructures. Therefore, we address now the utility of traces obtained from different vantage points.

Figure 3 displays the number of discovered /24 subnetworks of hosting infrastructures when we stepwise add traces from our dataset, see Section 3.2. The x-axis indicates the number of used traces while the y-axis shows the cumulative number of /24 subnetworks that can be identified based on these traces. The four curves of Figure 3 provide the results for different sequences in which we stepwise add traces: to obtain the curve with the black circles, we add in each step the trace that adds most additional /24 subnetworks to the set of already identified subnetworks (“Optimized”). In addition to the optimized cumulative utility (black circles) Figure 3 shows the maximum, median, and minimum cumulative utilities for 100 random permutations of the 133 traces.

In total, we find more than 8000 /24 subnetworks that are utilized by hosting infrastructures. We observe that every trace samples about half of these subnetworks (4800). About 2800 of these subnetworks are found in all traces. This relatively high fraction of common subnetworks among traces is the consequence of our choice of hostnames. It is not biased towards any given hosting infrastructure.

To estimate the potential benefits by including additional vantage points, we study the slope of the median curve (green crosses). Extrapolating the utility of the last 20 traces yields approximately ten /24 subnetworks per additional trace.

Furthermore, we notice that the traces that provide the highest utility (traces corresponding to the leftmost side in Figure 3), are actually located in different ASes and countries. For example, the first 30 traces belong to 30 different ASes in 24 different countries. The first 80 traces belong to 67 different ASes and 26 countries. This highlights the importance of utilizing vantage points that are geographically diverse and are hosted in different ASes.

To better understand both the need for diversity in vantage points as well as the underlying reasons behind the limited additional network coverage of hosting infrastructures by each trace, we perform a direct comparison of the traces. For this we re-use the similarity concept defined by Equation 1. For the same hostname, we define the /24 subnetwork similarity between two DNS replies as the similarity between their respective sets of /24 subnetworks. For two traces, we define their similarity as the average of /24 subnetwork similarities across all hostnames.

In Figure 4 we show the cumulative distribution of the similarity across all pairs of traces (TOTAL). We also show the similarity across traces when considering only one of the three subsets of the hostname list (EMBEDDED, TOP2000, TAIL2000). The high baseline value of similarity (always above 0.6) highlights the need for diversity to sample hosting infrastructures. It also confirms the slow increase in the utility of the traces shown in Figure 3.

As expected, the similarity for TAIL2000 is very high, indicating the limited location diversity for the corresponding hosting infrastructure. This is contrasted with the similarity for EMBEDDED, that is the lowest among the four curves. A low similarity for EMBEDDED is the consequence of the nature of the corresponding objects: typically they have a long lifetime and often are large. This makes them prime candidates for being hosted on distributed infrastructures, e.g., CDNs. TOP2000 lies in-between TAIL2000 and EMBEDDED. This indicates that the corresponding hostnames are hosted on a mix of centralized and distributed hosting infrastructures.

3.4.4 Summary

Our choice of a mix of different hostnames enables us to estimate the effect of our hostname list on sampling hosting infrastructures. Popular hostnames and embedded objects contribute most to discovering networks used by hosting infrastructures. Overall, studying data coverage we find that our set of popular and embedded hostnames is unlikely to miss large hosting infrastructures. The diversity of vantage points in terms of geographic and network location, however, is crucial to obtain good coverage.

4. RESULTS

In this section we examine our data set. First, based on the IP addresses we investigate where content can be obtained from. Next, we apply our clustering algorithm to characterize the resulting hosting infrastructure clusters. We gain insight on the deployment and hosting strategies of different infrastructures. We utilize our insight to derive content-centric AS-rankings and compare them with existing ones.

4.1 A Continent-level View of Web Content

Before delving into characteristics of hosting infrastructures, we want to understand which parts of the world serve Web content. In this section we choose the granularity of a continent, for two reasons: (i) the results directly reflect the round trip time penalty of exchanging content between continents, and (ii) our sampling is not dense enough to support country-level statistics. We quantify to which degree a user can find content in her own continent. This provides a view on the relative importance of different continents for Web content delivery as well the degree of replication of content.

4.1.1 Geographic Replication of Content

We first examine the relationship between the locations of content requester and content location as identified by DNS answers. We focus on TOP2000 in this section, and compare with other content types in the following section. Each line of Table 1 summarizes requests that originate from a given continent. Columns of Table 1 break down the requests among the continents from which the requested hostname is served. Each line adds up to 100%, while columns do not as they reflect the global importance of a continent. The shade of each entry of Table 1 is a visual aid, directly indicating its value (the darker the higher is the value).

At least 46% of the popular hostnames can be served from North America, 20% from Europe and 18% from Asia. The other three

Requested from	Served from					
	Africa	Asia	Europe	N. America	Oceania	S. America
Africa	0.3	18.6	32.0	46.7	0.3	0.8
Asia	0.3	26.0	20.7	49.8	0.3	0.8
Europe	0.3	18.6	32.2	46.6	0.2	0.8
N. America	0.3	18.6	20.7	58.2	0.2	0.8
Oceania	0.3	20.8	20.5	49.2	5.9	0.8
S. America	0.2	18.7	20.6	49.3	0.2	10.1

Table 1: Content matrix for TOP2000. Each line provides the percentage of all requests that originate from a given content. Columns indicate the continent from where content is served.

Requested from	Served from					
	Africa	Asia	Europe	N. America	Oceania	S. America
Africa	0.3	26.9	35.5	35.8	0.3	0.6
Asia	0.3	37.9	18.3	40.1	1.1	0.6
Europe	0.3	26.8	35.6	35.6	0.4	0.6
N. America	0.3	26.5	18.4	52.9	0.3	0.6
Oceania	0.3	29.2	18.5	38.7	11.3	0.6
S. America	0.3	26.4	18.2	39.3	0.3	14.2

Table 2: Content matrix for EMBEDDED. Each line provides the percentage of all requests that originate from a given continent. Columns indicate the continent from where content is served. The diagonal is more pronounced than for TOP2000 (Table 1).

continents, namely Africa, Oceania, and South America, do not appear to serve a lot of popular hostnames.

Another observation from Table 1 is a strong diagonal in the matrix, indicating that at least part of the hostnames are fetched from the same continent. Subtracting the minimum of a column from the corresponding element in the diagonal reveals that up to 11.6% of the hostname requests are served from their own continent. This locality of hostnames availability provides evidence that a considerable fraction of content is replicated in different regions of the world. Note, by choosing the granularity of countries, the existing diversity within continents is hidden. In addition, we observe an almost identical behavior for hostnames requested from Africa and Europe. Two factors are likely to cause this behavior: (i) a limited number of traces from Africa and (ii) the fact that Internet connectivity in Africa is mostly provided via Europe coupled to the lack of local content replication infrastructure. Oceania and Asia localize to a lesser degree than either Europe or North America.

4.1.2 Content-dependent Replication

Content varies in both popularity and type. This is the reason why we distinguish not only popular and less popular hostnames, but also different types of embedded objects (see Section 2). In this section, we refine the previous analysis of the relationships between the locations of content requester and content origin by comparing with the other two subsets of hostnames: TAIL2000 and EMBEDDED.

Surprisingly, the content delivery matrix for TAIL2000 (not shown) is almost identical to the one for TOP2000 (Table 1). The largest difference is a stronger concentration towards North America, with up to 1.4% points higher entries for TAIL2000. This indicates that the degree to which replicated hosting infrastructures are used in TOP2000 and TAIL2000 is very similar. How does that relate to the fact that top content has a far better utility in sampling large amounts of the address space than TAIL2000, cf. Section 3.4.2? The hosting infrastructures of highly popular content

Rank	#hostnames	#ASes	#prefixes	owner	content mix
1	476	79	294	Akamai	
2	161	70	216	Akamai	
3	108	1	45	Google	
4	70	35	137	Akamai	
5	70	1	45	Google	
6	57	6	15	Limelight	
7	57	1	1	ThePlanet	
8	53	1	1	ThePlanet	
9	49	34	123	Akamai	
10	34	1	2	Skyrock OSN	
11	29	6	17	Cotendo	
12	28	4	5	Wordpress	
13	27	6	21	Footprint	
14	26	1	1	Ravand	
15	23	1	1	Xanga	
16	22	1	4	Edgecast	
17	22	1	1	ThePlanet	
18	21	1	1	ivwbox.de	
19	21	1	5	AOL	
20	20	1	1	Leaseweb	

Table 3: Top 20 hosting infrastructure clusters by hostname count. The order of bars in the content mix column is: ■ only on TOP2000, ■ both on TOP2000 and EMBEDDED, ■ only on EMBEDDED, and ■ TAIL2000.

are distributed to a larger degree within each continent, when compared to hosting infrastructures predominantly used for less popular content.

When comparing the matrix of EMBEDDED (Table 2) with the others (TAIL2000 not shown, TOP2000 in Table 1), we observe that the diagonal is more pronounced for EMBEDDED. This indicates that embedded objects are, on a continent-level, more locally available than content from the other sets. We notice that Asia appears stronger for EMBEDDED compared to TOP2000 and TAIL2000, while North America appears weaker.

4.1.3 Summary

In this section, we analyzed the relative weights of Web content in different continents. We showed the prevalence of North America, Europe, and Asia in Web content presence, and how each region relies on each other. We observed a considerable local availability of content in most continents, implying that a considerable fraction of content is replicated across multiple continents.

4.2 A Portrait of Hosting Infrastructures

We turn our attention to the independent hosting infrastructures that are serving the hostnames from our list. In this section we identify the prominent hosting infrastructures, detect where they are actually located by AS and country, and classify them according to their network location footprint. Moreover, we study the geographic properties of hosting infrastructures and provide a ranking of countries according to their capability of serving popular Web content.

4.2.1 Clustering Validation

As first step, we validate the output of our algorithm of Section 2.3. Table 3 presents the top 20 clusters by hostname count. We find out by manually cross-checking that, indeed, all top 20 clusters correspond to networks that host large amounts of content.

Moreover, we leverage additional information about the Akamai and Limelight hosting infrastructures. In the case of Akamai we know the names present in the A records at the end of the CNAME chain inside DNS replies, which follow typical patterns. In the




Figure 5: Number of hostnames served by different hosting infrastructure clusters.

case of Limelight, we can use the same approach, and also verify that IP addresses belong to the AS number of Limelight. Scrutinizing this additional source of information, we find a few hostnames that are not included in the top clusters of Table 3. Typically, such hostnames are only available at a very small subset of the whole infrastructure, and are therefore separated in the first step of the clustering algorithm (Section 2.3) due to their unusual features. We conjecture these hostnames are intentionally treated differently by the hosting infrastructures.

4.2.2 Clustering Results

Overall, the output of our algorithm leads to the identification of more than 3000 potentially distinct hosting infrastructures. Figure 5 shows, for each hosting infrastructure cluster, the number of hostnames from our list that are served by the hosting infrastructure of that cluster, on a log-log scale. Hosting infrastructure clusters are ranked in decreasing order of hostname count. We observe that a few hosting infrastructure clusters are serving a large number of hostnames. Most of the hosting infrastructure clusters serve a single hostname. Hosting infrastructure clusters that serve a single hostname have their own BGP prefix. We infer that these are likely to be located in a single facility and in most of the cases serve non-replicated content. The top 10 largest hosting infrastructures are serving more than 15 % of the hostnames from our list. The top 20, still less than 1 % of all clusters, are serving about 20 % of the hostnames.

The resulting clustering allows us to make qualitative observations, namely that well-known hosting infrastructures are represented. Table 3 lists the top 20 clusters in terms of the number of hostnames from our list. Among them, we find well distributed CDNs such as Akamai, “hyper-giants” such as Google, and data-centers, such as ThePlanet. Note the different deployment strategies as coined by Leighton [24].

As can be seen in Table 3, we find multiple hosting infrastructure clusters run by the same infrastructure authority. Possible explanations include the observation that different types of services are hosted by hosting infrastructures, the geographic disparity of the infrastructure deployment, and acquisition or mergers. For Akamai and Google, the network footprints of their clusters are different

enough to be separated by the k -means step (first step) of the clustering algorithm (Section 2.3), while the ThePlanet clusters are only separated by the similarity step (second step).

For Akamai, the top two clusters have about twice as many locations as the other two clusters. Further investigation reveals that the first two Akamai clusters correspond to servers in the `akamai.net` second level domain (SLD), and the latter two are hosted on servers in the `akamaiedge.net` SLD.

Similar observations are made for the two Google clusters. The cluster with rank 3 is used for the standard Google services including Web search and the YouTube front page. The hostnames in this cluster are served from more than 200 IP addresses. The cluster with rank 5 hosts `doubleclick.net`, `googleapis.com`, `blogspot.com`, and several other Google related SLDs, which are served by 40 to 130 IP addresses each.

For ThePlanet, hostnames rarely map to more than one IP address. The clusters only emerge in step 2 of the clustering algorithm (Section 2.3), because the content is hosted on different BGP prefixes.

In addition, Table 3 shows the content-mix that is hosted on different hosting infrastructure clusters as bar-plot. We add CNAMEs (which come out of the Alexa Top 5000) to TOP2000 and report it as top content. In addition, we separate hostnames which are on both TOP2000 and EMBEDDED into a category of its own, top and embedded, to facilitate the interpretation of the content mix bar-plot. The order of the categories in the bar-plot (cf. caption of Table 3) is crafted to allow visually adding content of the top and embedded category to either the TOP2000 or the EMBEDDED category.

As expected, embedded content is prevalent on the top hosting infrastructures. Several hosting infrastructure clusters even almost exclusively serve embedded content, e.g., Limelight, Skyrock, Xanga, Edgecast, and the ad-service `ivwbox.de`. On others, e.g., the Akamai clusters, the first Google cluster, and AOL, top content accounts for a large fraction of the served hostnames. Some hosting infrastructure clusters serve mainly tail content, e.g., Wordpress, ThePlanet, and the second Google cluster. This may appear surprising at first, but it is a result of content consolidation. For example, blogs hosted by Google and Wordpress can be found amongst this content. This highlights how helpful less popular content can be in identifying and classifying the hosting infrastructures. The above mentioned results show that separating the hosting infrastructure of certain organizations into multiple clusters is both justified and necessary, because the infrastructures are not used homogeneously [30].

4.2.3 Geographic Deployment of Hosting Infrastructures

The clustering of the hosting infrastructures presented in the previous section is agnostic with respect to geographic locations. To provide insight about the geographic deployment of the different clusters as well as to better understand the degree of co-location in the content infrastructure, we map the clusters to the geographic locations of their prefixes.

Distinguishing between content infrastructures that rely on a few ASes or prefixes is tricky, especially because we do not have a priori knowledge about their signature in terms of ASes and prefixes. Indeed, some of these clusters might very well be present in a single location but for administrative reasons split their infrastructure into multiple ASes or use multiple prefixes due to multi-homing. One known example is Rapidshare [8], that relies on multiple ASes and prefixes yet whose facility is a single data-center.

Therefore, we estimate the number of countries in which a host-

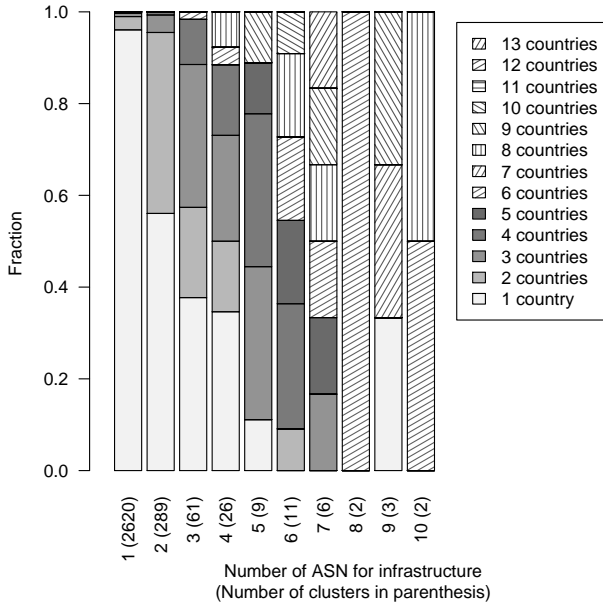


Figure 6: Country-level diversity of content infrastructure clusters.

ing infrastructure cluster is present based on geolocation information. For each cluster we check onto how many countries it is deployed. Figure 6 shows the resulting relationship in the form of a stacked bar-plot. On the x-axis, we show the number of ASes in which clusters are discovered. We annotate the number of different clusters found for a given number of ASes in parenthesis. On the y-axis, we show the fraction of clusters whose prefixes are located in a given number of countries (see legend).

As shown in Figure 6, most of the hosting infrastructure clusters that use a single AS are present in a single country. As a cluster’s footprint is on more ASes, the likelihood that it is present in multiple countries increases. At the same time, a significant fraction of hosting infrastructure clusters using multiple ASes are located in a single country. Because of the limited number of content infrastructure clusters located in 5 or more ASes (33 clusters), the fraction for these clusters in Figure 6 simply reflects a few instances of particular hosting infrastructures. Most of these clusters are present in several countries, thus, are probably CDNs.

4.2.4 Summary

In this section, we used our clustering algorithm to identify hosting infrastructures based on network features. We validate our classification based on additional information for two large CDNs. We present evidence of how Akamai and Google slice up their hosting infrastructures for different hosting purposes, and show which content is hosted on which hosting infrastructures. We distinguish between smaller hosting infrastructures based on both their features and their geographic deployment. We find a relationship between the number of ASes on which a hosting infrastructure is present and the multiplicity of its locations, giving a hint about their deployment strategy.

4.3 Mapping Hosting Infrastructures

To find out the geographic location of the hot-spots that serve most hostnames from our list, we compute both content potentials on a per country basis. Table 4 shows the results for both poten-

Rank	Country	Potential	Normalized potential
1	USA (CA)	0.254	0.108
2	China	0.128	0.107
3	USA (TX)	0.190	0.061
4	Germany	0.183	0.058
5	Japan	0.163	0.051
6	France	0.146	0.034
7	Great Britain	0.157	0.030
8	Netherlands	0.144	0.029
9	USA (WA)	0.135	0.027
10	USA (unknown)	0.164	0.027
11	Russia	0.038	0.027
12	USA (NY)	0.130	0.026
13	Italy	0.122	0.018
14	USA (NJ)	0.125	0.016
15	Canada	0.028	0.015
16	USA (IL)	0.116	0.014
17	Australia	0.118	0.013
18	Spain	0.116	0.013
19	USA (UT)	0.111	0.012
20	USA (CO)	0.113	0.012

Table 4: Geographic distribution of content infrastructure, ranked by the normalized potential.

tials. Note, for the USA only, we provide the state level. The lines of Table 4 are ranked by decreasing normalized content delivery potential and show the top 20 hosting infrastructures.


Despite the division into states, the USA leads the ranking with its hosting infrastructure in California. Indeed, in total 9 US states are among the top 20. On the second place we find China. Directly comparing California with China reveals that China’s delivery potential is a lot lower than California’s, yet the values of their normalized potential are quite close. Comparing China’s potential with its normalized potential indicates that a large fraction of the content served from China is only available in China. In total, China and California together count for over 23% of hostnames of our list in the normalized potential. Besides USA and China, 7 European countries are among the top 20, as well as Japan, Australia and Canada. In total we see content being delivered from 122 countries/US states, or 77 countries. The top 20 countries/US states presented here are responsible for 70% of all hostnames in our study.

4.4 Revisiting AS Rankings

Geographic hot-spots are insightful in that they reflect where large chunks of the hosting infrastructures are. However, they provide little insight to understand how content is delivered to Internet users. Therefore, we investigate where content resides at the AS-level.

To map hosting infrastructure clusters to ASes, we rely on the same approach as in Section 4.2.3. For each cluster, we take the prefixes from which it serves content, and map each prefix to an AS number using BGP data. This gives us a set of AS numbers for each cluster. Recall that a hosting infrastructure cluster is defined by a set of hostnames it serves. We reuse the notion of content delivery potential, as introduced in Section 2.4, but where “locations” are now ASes. The content delivery potential of an AS is the fraction of hostnames it can potentially serve from all the clusters that are hosted on this AS.

Figure 7 provides the top 20 ASes in terms of their content delivery potential. Unexpectedly, we find mostly ISPs in this top 20. Note that the CMI (cf. Section 2.4) is very low for all the top ranked ASes. The two genuine content hosters in the list are Akamai and Bandcon. There are two main factors explaining the unexpected top




Rank	AS name	CMI
1	NTT America	0.070
2	Tinet	0.029
3	Global Crossing	0.034
4	KDDI	0.025
5	Akamai Europe	0.019
6	TeliaNet Global	0.027
7	Deutsche Telekom	0.033
8	Korea Telecom	0.030
9	Qwest	0.036
10	Bandcon	0.045
11	Cable and Wireless	0.021
12	SingTel Optus	0.019
13	Akamai	0.018
14	France Telecom - Orange	0.017
15	Internode	0.017
16	Comcast	0.017
17	StarHub	0.018
18	nLayer	0.020
19	Beyond The Network	0.018
20	TATA	0.023

Figure 7: Top 20 ASes in content delivery potential.

20: (i) all these ASes host Akamai caches that boost their content delivery potential and (ii) all these ASes host some content that no other AS can provide. Given the widespread deployment of Akamai caches in carriers, the second factor is actually more important and explains why some ASes appear among the top and why others do not. A content-centric AS-ranking should be able to cope with the bias introduced by the deployment of highly distributed hosting infrastructures within ASes.

An AS-ranking based on the normalized content delivery potential does exactly this. It spreads the weight of distributed content infrastructure across all ASes that serve their hosted content. Figure 8 provides the top 20 ASes in terms of normalized content delivery potential. Our first observation is that the only overlap with the non-normalized ranking is NTT. The ASes that appear on the top of the normalized ranking do so because of the exclusiveness of the content they host as reflected by their CMI values. As expected, Google is among the top ranked ASes due to its importance in popular content. We also see data-center content infrastructures: ThePlanet, SoftLayer, Rackspace, 1&1 Internet, OVH, Amazon, Leaseweb, and Hetzner Online. A limited number of ISPs in China



Rank	AS name	CMI
1	Chinanet	0.699
2	Google	0.996
3	ThePlanet.com	0.985
4	SoftLayer	0.967
5	China169 Backbone	0.576
6	Level 3	0.109
7	China Telecom	0.470
8	Rackspace	0.954
9	1&1 Internet	0.969
10	OVH	0.969
11	NTT America	0.070
12	EdgeCast	0.688
13	GoDaddy.com	0.969
14	Savvis	0.785
15	China169 Beijing	0.706
16	Amazon.com	0.895
17	LEASEWEB	0.942
18	Cogent	0.687
19	Hetzner Online	0.962
20	AOL	0.932

Figure 8: Top 20 ASes in normalized content delivery potential.

seem to have a monopoly on a considerable fraction of popular content.

When comparing the normalized potential for all hostnames with TOP2000 and EMBEDDED, two more ASes enter the picture, Abitcool China and China Networks Inter-Exchange. In addition, slight re-rankings can be observed.

4.4.1 Content vs. Traditional AS Rankings

Many topology-driven rankings have been proposed [5, 22, 3, 2]. In this section we compare content-based rankings, described in the previous section, with topology-driven ones. Among the most well-known AS rankings are two from CAIDA [5]: one based on AS-degree and another on the size of the customer cone of an AS. Fixed Orbit [2] calculate an the Knodes Index, a centrality based metric. Labovitz et al. [22] provide an AS ranking based on the amount of traffic by monitoring inter-domain traffic from 110 ASes.

Table 5 compares seven different AS rankings: the CAIDA AS-degree (CAIDA-degree) and customer cone (CAIDA-cone) rankings [5], a ranking similar to CAIDA's by Renesys (Renesys) [3], the Knodes Index (Knodes) as reported by Fixed Orbit [2], the traf-

Rank	CAIDA-degree	CAIDA-cone	Reneysys	Knodes	Arbor	Potential	Normalized potential
1	Level 3	Level 3	Level 3	Level 3	Level 3	NTT	Chinanet
2	Cogent/PSI	AT&T	Global Crossing	Cogent	Global Crossing	Tinet	Google
3	AT&T	MCI	Sprint	Global Crossing	Google	Global Crossing	ThePlanet
4	MCI	Cogent/PSI	NTT	Sprint	*	Deutsche Telekom	SoftLayer
5	Hurricane	Global Crossing	Savvis	Tinet	*	KDDI	China169 backbone
6	Qwest	Sprint	TeliaSonera	NTT	Comcast	Telia	Level 3
7	Sprint	Qwest	Tinet	AT&T	*	Akamai	Rackspace
8	Global Crossing	Hurricane Electric	Verizon	Swisscom	*	Bandcon	China Telecom
9	tw telecom	tw telecom	AT&T	Hurricane	*	Cable and Wireless	1&1 Internet
10	INIT7	TeliaNet	China Telecom	Telia	*	Qwest	OVH

Table 5: Topology-driven AS rankings against traffic-driven and content-based AS rankings.

fic exchanges-based ranking by Labovitz et al. [23]² (Arbor), and finally our content-based rankings (potential and normalized potential). The purely topological rankings like the ones from CAIDA and Renesys tend to rank large transit carriers high. Besides the case of Google and Comcast, the top of Arbor’s ranking leads to similar results to topological rankings. Our content infrastructure-driven rankings on the other hand give more weight to those ASes that deliver a large amount of content. We notice that our normalized potential leads to similar top ranked ASes as topological and traffic-based rankings, while of course favoring ASes that host content.

We argue that no AS ranking captures all relevant aspects of the importance of an AS. All aspects of the Internet are important, i. e., topology, traffic, and content, and need to be taken into consideration to understand the Internet ecosystem.

4.4.2 Summary

We proposed two different ways to rank ASes based on their content potential. We showed that these rankings reveal different aspects of content hosted by ASes: replicated content and content exclusively hosted by a given AS. We proposed an index, called the content monopoly index, which measures the degree to which an AS hosts content not available elsewhere, compared to content that is replicated in other ASes. Finally, we related our content-centric rankings to those presented in the literature.

5. DISCUSSION

The deployment of hosting infrastructures is dynamic in multiple ways, e. g., by growing their existing infrastructure, by changing their peerings with ISPs, or by placing part of the infrastructure inside ISP networks. Our methodology is an automated tool that makes it possible to monitor the state of hosting infrastructures deployment in the Internet at a given point in time. Given the changing nature of the hosting infrastructures ecosystem, as observed by Labovitz et al. [22], it is important to have tools that allow the different stakeholders in the Internet to better understand the space in which they evolve.

With the commoditization of the content delivery and hosting landscape, content producers have more and more choice as to where they can place their content, for example to better target their user base. With the rise of user-generated content, e. g., through Facebook, Twitter, and YouTube, content producers also need to deliver a mix of different media for which multiple suitable content delivery platforms might be available. Meta-CDNs such as Meebo, an instant messenger aggregator, or Conviva³, a video delivery platform, distribute the content demand across different CDNs. To un-

²Some of the entries of the Arbor ranking were intentionally omitted by [22].

³<http://www.conviva.com/>

derstand the trade-offs involved in their business decisions, content producers need to be able to measure the actual server diversity visible to different users from different places in the world, as well as the performance delivered by different CDNs. Our work is a step in this direction.

Given the already wide deployment of hosting infrastructures around the world, existing and new CDNs need to understand the marginal utility of deploying more infrastructure. Indeed, competition among CDNs is intense, and exacerbated by the fact that ISPs have already deployed and are deploying more hosting infrastructures inside their network. The recently created Content Delivery Networks Interconnect (CDNi) [25] working group at the IETF, aimed at standardizing interconnection between CDNs, is a proof of the maturity of the CDN market that now considers collaboration.

Some ISPs nowadays deploy their own CDN, data-centers, caches, or even CDN boxes such as those from Akamai. The increased availability of content caching and hosting with ISP networks further complexifies the current and future hosting infrastructures landscape. With the significant server diversity from which a given content can be obtained [31], both outside and inside their network, the traffic engineering task of ISPs is becoming even more challenging. Separating intradomain traffic engineering from peering engineering is indeed becoming less and less relevant. ISPs need to think globally about the network and understand what content can be obtained from where before making traffic engineering and peering decisions. Web content cartography, combined with a better understanding of content delivery performance, can help ISPs to an adequate strategy to deliver content to their customers.

It is worth mentioning the critical role of virtualization for hosting infrastructures. Virtualization is one of the highly disruptive enablers that we believe will change the hosting infrastructures landscape in the near-future. Indeed, virtualization technology offers flexibility that can be exploited by content producers and providers, as well as ISPs, to dynamically move content. By exploiting virtualization, we expect to see a hosting infrastructures market that is richer, that better utilizes the available resources, and better serves end-users. Virtualization has both a positive and a negative impact on Web content cartography. On the negative side, the use of virtualization, together with collaboration between CDNs and meta-CDNs, will pose challenges to Web content cartography. On the positive side, it will make it more necessary to keep track of the dynamic mapping of content onto hosting infrastructures.

The research community also needs an understanding of the evolving ecosystem of hosting infrastructures, given its importance as a driver in the evolution of the Internet. In light of the recent observations by Labovitz et al. [22] that showed significant changes in the whole Internet ecosystem, it is important for the research community to be able to foresee the impact of hosting infrastructures to work on the most relevant technical challenges that the future Inter-

net may face. Our work is a first step in the direction of creating a comprehensive map of the whole Internet ecosystem that includes hosting infrastructures. A map of hosting infrastructures combined with a corresponding view of the paths (and their performance) to the end-users, would constitute a significant step forward in our understanding of Internet content delivery.

6. RELATED WORK

Our work on Web content cartography is motivated by recent studies that provide evidence in support of the significant rise of Web content traffic [15, 12, 22, 27]. Two major reasons are the growth of video traffic and the increasing penetration of broadband access. To cope with these changes, large-scale content distribution networks are being deployed [24, 30]. In addition, applications, such as file sharing, that used to rely on peer-to-peer delivery are nowadays increasingly served from data-centers [21], or One-click Hosters [8].

Labovitz et al. [22] observed consolidation of Web content traffic as well as a significant shift in peerings to better facilitate connectivity to content providers. They analyzed inter-domain traffic over a two year period, relying on data from 110 commercial ISPs and content providers. They detect global shifts in inter-domain traffic and the AS ecosystem. Our approach is complementary and focuses on hosting infrastructures, not inter-domain traffic. The advantage of our approach is our extensive coverage of popular hosting infrastructure, whose importance may be underestimated when analyzing only inter-domain traffic. For example, a significant fraction of the traffic delivered by a highly distributed infrastructure such as Akamai as well as data-centers located within ISPs, does not cross AS boundaries.

Shue et al. [34] observed, by looking at the .com and .net domains, that a vast majority of Web servers are co-located. Our results, on a more diverse set of domains, confirm that there is co-location of servers as well as hosting infrastructures.

Huang et al. [17], Su et al. [35], and Triukose et al. [36] leverage DNS requests to understand the distribution of Web content. However, their work is restricted to the study of specific CDNs, and does not try to detect different types of hosting infrastructures. Utilizing DNS replies of popular content in order to identify the location of hosting infrastructures shares similarities with work by Ager et al [7], by Krishnamurthy et al [20], by Feldmann et al. [14], and by Chang et al. [9]. Yet, their focus was on studying the performance of DNS resolvers [7, 20] or on deriving empirical models of inter-domain traffic matrices [14, 9], rather than leveraging DNS for Web content cartography.

Other approaches have been proposed to identify hosting infrastructures. Gill et al. [16] performed traceroutes towards large hosting infrastructures from distributed vantage points. They relied on DNS resolution from a single vantage point inside a university, providing a limited view of the network footprint of hosting infrastructures. Recently, Wang et al. [39] proposed a measurement methodology to estimate the distance between end-users towards two large CDNs. They rely on embedded measurement scripts within multiple websites and collect the measurements performed by the end-users. Our approach differs in that we do not require the targeted hosting infrastructure to be known a priori.

Clustering has been used in the past in different contexts, e. g., to group end-users from the perspective of routing [19] or DNS resolution [11]. Our approach does not focus on clustering of end-users, but rather on the identification and clustering of hosting infrastructures in the Internet.

7. CONCLUSION

In this paper, we introduce Web content cartography. We propose a lightweight and fully automated approach to discover Web content hosting and delivery infrastructures based only on DNS measurements and BGP routing table snapshots. To the best of our knowledge this is the first attempt to identify hosting infrastructures based on DNS replies and routing information, rather than relying on pre-identified signatures such as CNAMEs. The advantage of this method is that is general enough to identify new hosting infrastructures as well as cluster them based on their operation as revealed by DNS.

Our results show that by utilizing traces from a small number of well distributed vantage points it is possible to make qualitative observations for the deployment of hosting infrastructures and content replication. A key insight of our study is that a significant fraction of the content is exclusively delivered by hosting infrastructures such as Google or geographical regions, e. g., China. Furthermore, by deriving content-centric AS rankings that complement existing AS rankings we shed light on recent observations about shifts on the AS topology.

Our work is an important step towards answering crucial questions for content producers, content providers, ISPs, and the research community. For ISPs, knowing the locations from which popular content can be obtained is a key factor in peering decisions and network dimensioning. For content producers, geographic and network footprint of a CDN is an important factor in choosing how to deliver content best to their customer base. For content providers, Web content cartography can help them improve their competitiveness in the content delivery market. Moreover, the research community needs to understand and track the evolving ecosystem of hosting infrastructures, given their importance as a driver in the evolution of the Internet. Web content cartography is complementary to other maps of the Internet, e. g., router-level and AS-level maps, and enables further investigation of the shaping forces of the Internet. In particular it will allow researchers to investigate the interplay of content infrastructures with the Internet topology.

Acknowledgments

We want to thank the numerous contributors of DNS trace files.

8. REFERENCES

- [1] Alexa top sites. <http://www.alexa.com/topsites>.
- [2] Fixed Orbit Knodes Index. <http://www.fixedorbit.com/metrics.htm/>.
- [3] Renesys Market Intelligence. http://www.renesys.com/products_services/market_intel/.
- [4] RIPE Routing Information Service. <http://www.ripe.net/ris/>.
- [5] The CAIDA AS Ranking. <http://as-rank.caida.org/>.
- [6] Whois Source – Domain Counts & Internet Statistics. <http://www.whois.sc/internet-statistics/>.
- [7] B. Ager, W. Mühlbauer, G. Smaragdakis, and S. Uhlig. Comparing DNS Resolvers in the Wild. In *Proc. ACM IMC*, 2010.
- [8] D. Antoniadis, E. Markatos, and C. Dovrolis. One-click Hosting Services: A File-Sharing Hideout. In *Proc. ACM IMC*, 2009.
- [9] H. Chang, S. Jamin, M. Mao, and W. Willinger. An Empirical Approach to Modeling Inter-AS Traffic Matrices. In *Proc. ACM IMC*, 2005.

- [10] K. C. Claffy and N. Brownlee. Understanding Internet Traffic Streams: Dragonflies and Tortoises. *IEEE Commun. Mag.*, 2002.
- [11] C. D. Cranor, E. Gansner, B. Krishnamurthy, and O. Spatscheck. Characterizing Large DNS Traces Using Graphs. In *Proc. ACM IMW*, 2001.
- [12] J. Erman, A. Gerber, M. Hajiaghayi, D. Pei, and O. Spatscheck. Network-aware Forward Caching. In *Proc. WWW*, 2009.
- [13] W. Fang and L. Peterson. Inter-AS Traffic Patterns and their Implications. In *Proc. IEEE Global Internet*, 1999.
- [14] A. Feldmann, N. Kammenhuber, O. Maennel, B. Maggs, R. De Prisco, and R. Sundaram. A Methodology for Estimating Interdomain Web Traffic Demand. In *Proc. ACM IMC*, 2004.
- [15] A. Gerber and R. Doverspike. Traffic Types and Growth in Backbone Networks. In *Proc. of OFC/NFOEC*, 2011.
- [16] 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.
- [17] C. Huang, A. Wang, J. Li, and K. Ross. Measuring and Evaluating Large-scale CDNs. In *Proc. ACM IMC*, 2008.
- [18] V. Jacobson, D. Smetters, J. Thornton, M. Plass, N. Briggs, and R. Braynard. Networking Named Content. In *Proc. CoNEXT*, 2009.
- [19] B. Krishnamurthy and J. Wang. On Network-aware Clustering of Web Clients. In *Proc. ACM SIGCOMM*, 2001.
- [20] B. Krishnamurthy, C. Wills, and Y. Zhang. On the Use and Performance of Content Distribution Networks. In *Proc. ACM IMW*, 2001.
- [21] R. Krishnan, H. Madhyastha, S. Srinivasan, S. Jain, A. Krishnamurthy, T. Anderson, and J. Gao. Moving Beyond End-to-end Path Information to Optimize CDN Performance. In *Proc. ACM IMC*, 2009.
- [22] C. Labovitz, S. Lekel-Johnson, D. McPherson, J. Oberheide, and F. Jahanian. Internet Inter-Domain Traffic. In *Proc. ACM SIGCOMM*, 2010.
- [23] C. Labovitz, D. McPherson, and S. Iekel-Johnson. Internet observatory report, 2009. <http://www.nanog.org/meetings/nanog47>.
- [24] T. Leighton. Improving Performance on the Internet. *Commun. ACM*, 52(2):44–51, 2009.
- [25] Kent Leung and Yiu Lee. Content Distribution Network Interconnection (CDNi) Requirements. IETF draft, work in progress, draft-ietf-cdni-requirements-00, Sep 2011.
- [26] S. Lloyd. Least Squares Quantization in PCM. *IEEE Trans. Information Theory*, 1982.
- [27] G. Maier, A. Feldmann, V. Paxson, and M. Allman. On Dominant Characteristics of Residential Broadband Internet Traffic. In *Proc. ACM IMC*, 2009.
- [28] Z. M. Mao, C. Cranor, F. Douglis, M. Rabinovich, O. Spatscheck, and J. Wang. A Precise and Efficient Evaluation of the Proximity Between Web Clients and Their Local DNS Servers. In *Proc. Usenix ATC*, 2002.
- [29] MaxMind. <http://www.maxmind.com/app/ip-location/>.
- [30] E. Nygren, R. K. Sitaraman, and J. Sun. The akamai network: a platform for high-performance internet applications. *SIGOPS Oper. Syst. Rev.*, 44:2–19, August 2010.
- [31] I. Poese, B. Frank, B. Ager, G. Smaragdakis, and A. Feldmann. Improving Content Delivery using Provider-aided Distance Information. In *Proc. ACM IMC*, 2010.
- [32] I. Poese, S. Uhlig, M. A. Kaafar, B. Donnet, and B. Gueye. IP geolocation databases: unreliable? *ACM CCR*, 41:53–56, 2011.
- [33] University of Oregon Route Views Project. <http://www.routeviews.org/>.
- [34] C. A. Shue, A. J. Kalafut, and M. Gupta. The Web is Smaller than it Seems. In *Proc. of IMC*, 2007.
- [35] A. Su, D. Choffnes, A. Kuzmanovic, and F. Bustamante. Drafting Behind Akamai: Inferring Network Conditions Based on CDN Redirections. *IEEE/ACM Trans. Netw.*, 17(6):1752–1765, 2009.
- [36] S. Triukose, Z. Wen, and M. Rabinovich. Measuring a Commercial Content Delivery Network. In *Proc. WWW*, 2011.
- [37] P. Vixie. What DNS is Not. *Commun. ACM*, 52(12):43–47, 2009.
- [38] J. Wallerich, H. Dreger, A. Feldmann, B. Krishnamurthy, and W. Willinger. A Methodology for Studying Persistency Aspects of Internet Flows. *ACM CCR*, 2005.
- [39] Y. A. Wang, C. Huang, J. Li, and K. W. Ross. Estimating the Performance of Hypothetical Cloud Service Deployments: A Measurement-based Approach. In *Proc. IEEE INFOCOM*, 2011.
- [40] Y. Zhang, L. Breslau, V. Paxson, and S. Shenker. On the Characteristics and Origins of Internet Flow Rates. In *Proc. ACM SIGCOMM*, 2002.