IEEE Access
Multidisciplinary ⋮ Rapid Review ⋮ Open Access Journal

Received January 9, 2017; accepted February 6, 2017, date of publication March 24, 2017, date of current version May 17, 2017.

*Digital Object Identifier 10.1109/ACCESS.2017.2682190*

# CDNs Meet CN

## An Empirical Study of CDN Deployments in China

**JING'AN XUE[1], DAVID CHOFFNES[2], AND JILONG WANG[1], (Member, IEEE)**
[1]Tsinghua National Laboratory for Information Science and Technology, Institute for Network Science and Cyberspace,
Tsinghua University, Beijing 100084, China
[2]College of Computer and Information Science, Boston, MA 02115, USA

Corresponding author: Jing'an Xue (xja12@mails.tsinghua.edu.cn)

This work was supported by the National Key Research and Development Program of China under Grant 2016YFB0801301.

**ABSTRACT** Today's content delivery networks (CDNs) use large sets of globally distributed servers, advanced routing techniques, and dynamic server selection to provide users with low delay, reliable access to Web, and streaming content. These strategies allow CDNs to extend their reach and performance into large numbers of networks around the globe. However, nearly 20% of Internet users reside in China, where local regulations and network policies make it challenging for global CDNs to serve foreign content from servers in China and limit the effectiveness of their traditional strategies for low-latency content delivery. Recently, a number of China-specific CDNs have partnered with global CDNs to provide service in the country. However, little is known about how CDNs in China are implemented, nor their impact on performance. In this paper, we are the first to investigate the impact of China's unique policy and networking environment on CDN implementations and deployments. We find that deployments inside China exhibit different client-mapping behavior than global ones, namely, through static server selection for clients given a customer site, and through region-specific clients partitioning and mapping rather than using live network information. We also show that ignoring these properties can significantly impact server selection quality. Our results can be useful for optimization of CDNs in China.

**INDEX TERMS** Content distribution networks, measurement, performance evaluation, server selection.

## I. INTRODUCTION

Content delivery networks (CDNs) are used pervasively to serve Internet traffic, and have become critical Internet infrastructure [33], [39], [45], [51]. Conventional wisdom dictates that global CDNs use a large set of widely distributed servers, often deep inside ISP networks or in critical exchange points, to provide optimal Web and streaming performance, and near-perfect availability.

In reality, these CDNs perform a balancing act to provide *sufficient* performance and availability subject to the cost and policy constraints imposed by the networks in which they are deployed. Costs are determined by factors such as energy consumption and bandwidth charges, while policies include available BGP paths and prevailing laws of the jurisdictions where servers are located.

Perhaps nowhere are these trade-offs more prominent than in China. With nearly 680 million users, representing 20% of all Internet users [1], China is the fastest growing and soon-to-be the largest online market in the world. However, China has proven to be challenging for global CDNs tap into due to economic, technical and regulatory factors [27], [52], as described in Section II. To fill this void, several Chinese CDN providers have entered the market [6], [19] and recently a number of global CDNs have partnered with them to provide delivery from inside of China.

In this paper, we investigate the impact of China's unique environment on the server selection and performance of CDNs. Specifically, we identify different patterns of server selection dynamics, and leverage it to show the inefficiency of a multi-CDN selector and the customization of Chinese CDNs. Using a comparative analysis from China and the US, we also characterize the impact of not having local replicas on server selection quality. We focus on the CDNs hosting the most popular 1 million Web sites (according to Alexa) and use 16 million measurements from vantage points in China and the US.

This paper makes the following key contributions:

- We conduct a measurement study focusing explicitly on CDN behavior for clients in multiple Chinese provinces, and compare them with behavior seen from the US. We identify that Chinese CDN providers exhibit static server selection dynamics for clients, providing highly

customized service for Web sites relative to global providers that serve customer content from a large number of replica server locations.
- We find that one multi-CDN selector, Cedexis, provides suboptimal performance by not accounting for customer-specific replica selection used by Chinese CDNs with static server selection dynamics.
- We identify the serving infrastructure and find that due to Internet Content Provider license regulations in China, several global CDN providers have partnered with Chinese CDNs to serve content from replica servers inside China. We find that global CDNs who partner with Chinese providers can achieve significantly lower latencies (as much as an order of magnitude) than those who do not. However, several CDN providers in China have much lower availability than global ones.
- Focusing on global CDN providers, given a vantage point, we find that anycast-based redirection offers nearly same performance for all customers when all anycast BGP prefixes are announced from all of the CDN's data centers. However, when anycast CDNs selectively announce their prefixes, clients may receive significantly higher-latency paths for different Web site customers.
- Further, we find that DNS-based redirection policies that work well in the US lead to suboptimal selection in China due to lack of close replicas. Specifically, if a CDN chooses one of the best $n$ replica locations in China, it might send clients to the US even when a replica is available in Asia. In the US, the best $n$ locations tend to all exhibit low latency.

In summary, we find that China's unique policy and network environment offers many challenges to incumbent CDN providers, leaves much room for improvement in terms of client-perceived latency, and different optimization methods may be employed by CDNs in China due to the relatively static mapping.

The rest of the paper is structured as follows. Section II reviews the background and related work. Methodology and dataset are presented in detail in Section III. We investigate the unique characteristics of Chinese CDNs in Section IV and compare overall performance of all targeted commercial CDNs in Section V. We then examine the behavior of global CDNs in Section VI and conclude in Section VII.

## II. BACKGROUND AND MOTIVATION

In this section, we briefly review the primary CDN redirection techniques in deployment today, how China's unique network and policy environment poses challenges for global CDNs, and how previous research relates to this paper.

### A. CLIENT MAPPING

An important goal of a CDN is to map clients (*e.g.*, users' Web browsers) to one or more replica servers that are both available to serve the requested content, and that can do so with relatively low-latency (*e.g.*, due to low server load and short network paths). To achieve this goal, CDNs generally use *DNS-based* and/or *anycast-based* mappings.

*DNS redirection* works as follows. When a client requests name resolution for a CDN-hosted domain, its local DNS server (LDNS) will contact the authoritative DNS server for that domain. Through techniques that include using CNAME entries, the LDNS eventually contacts a CDN's DNS server, which returns a set of IP addresses corresponding to available replica servers [26], [51]. An advantage of this approach is that a CDN can direct a client *dynamicallly* at request time, based on *live information* about network conditions and server load. However, a potential downside is that these redirections are based on a client's LDNS, which may not be colocated with clients [48], [50] and thus may lead to suboptimal performance. The EDNS0 client subnet [28] extension can potentially address this issue, but it has not yet been widely adopted except for Google servers [38].

When using *anycast* [32], a CDN's DNS server also returns one or more IP addresses to the client's LDNS server. In this case, however, the IP addresses returned will correspond to replica servers along the most-preferred BGP path between a client's ISP and the CDN's data centers announcing the corresponding prefix. The advantage of this approach is that it does not require a sophisticated DNS-redirection and measurement infrastructure, and the adopted paths are optimal from the ISP's perspective. However, the approach can be challenging to configure properly at a global scale [36], [37] and generally does not respond to varying network conditions without additional functionality. With anycast, CDN customers may be served a from a single IP prefix globally like Cloudflare [12], or from different prefixes in different locations like Edgecast [41] and Fastly [21], but the address will **not** change over time.

### B. CHINESE NETWORKS, POLICIES, AND CDNs

China has a large online population, which accounts for half of its total population [1] and nearly 20% of the Internet population. However, the speed and reliability of Internet infrastructure in China leaves much room for improvement.

In contrast with network topologies in Europe and the US, which have "flattened" due to extensive peering at IXPs [34], [46], China's network topology remains highly hierarchical. Specifically, China has a small number of ISPs, each with a large AS, built atop a nationwide backbone [52]. Inter-ISP links are a well-known bottleneck in China [42], with average latencies between different ISPs in major Chinese provinces registering latencies larger than 90 ms.

Many links between ASes in China are staticly configured, with BGP generally used only between larger Chinese ISPs. Only recently was the first IXP in China built [9]. This unique environment makes approaches like anycast fairly ineffective for CDN deployments in China [15].

In addition to technical challenges, CDNs face regulatory hurdles for serving content from China. Namely, the Internet Content Provider (ICP) Beian [27] is a regulation put forth by the Ministry of Industry and Information Technology (MIIT) of China that requires Web sites wishing to deliver content within China to register with the applicable Chinese governmental agency for an ICP license. As a result, CDNs serving a global set of customers cannot generally use replica servers in China because they might contain content without a corresponding ICP license.

Perhaps the easiest way to address this restriction is for a global CDN to partner with a Chinese CDN, so that the latter can handle the ICP license application process. As a result, few global CDN providers have infrastructure in mainland China (as we will confirm explicitly later); rather, clients are redirected to partner servers in China or to servers outside of China.

As demand for Internet content surges in China, a growing number of Chinese content providers are contracting with CDNs to provide users with good performance. The Chinese CDN market includes three dominant, incumbent providers (ChinaCache, ChinaNetCenter and FastWeb), as well as a few relatively new entrants in the end of 2015, AliCDN (Ali Kunlun), Baidu (Yunjiasu), and Tencent [3], [5], [31]. These new players are disrupting the CDN market in China by leveraging their existing large-scale infrastructure deployments to provide transparent SLAs at relatively low cost compared to the incumbents.

### C. RELATED WORK

CDNs are pervasive, responsible for large fractions of global Internet traffic and for allowing content providers to scale quickly and gracefully to millions of users. Given their prominent role, a number of studies have investigated their design and implementation.

Early work [44] evaluated Akamai and Digital Island in terms of redirection effectiveness, showing that they were able to avoid unavailable servers although not always with shortest latency. Several other efforts [43], [51], [53] focused on the infrastructure deployments, implementation of redirection, performance, and availability of large CDNs. More recent work focuses on CDN performance and expansions, including those of Akamai [40], [49], Yahoo [35], Google [38], [45], and Bing [39].

Our work builds upon and extends this prior work. We leverage previous studies of CDN implementations to inform our measurements of CDNs in China. Unlike prior work, we focus on differences between server selection dynamics and quality of CDN in China relative to those outside of China.

In closely related work [54], the authors reported the practice and experience of commercial CDN deployment in China around 2011. Our work is different in that they focus on general geographical coverage and market share while we identify server selection behavior that is based on customer sites, and its impact on performance. Additionally, they target the Chinese top 100 websites from China while we target the Alexa top 1 million (include Chinese and non-Chinese) websites from China and the US, which broadens the scope of our findings.

## III. METHODOLOGY AND DATASET

In this section, we present our methodology of CDN identification, validation, and measurements. Our goals are to measure how CDNs provide custom service to individual customers hosting popular Web sites, and compare the differences between US and Chinese CDN server selection dynamics.

Specifically, we focus on:

- Identifying which popular Web sites are delivered by CDNs and which techniques are used for mapping clients to servers.
- Determining how CDNs provide differential service to different customers.
- Measuring the relative performance of the different approaches in terms of latency and availability.

It is important to note several non-goals for this work. First, our goal is not to cover all CDNs, networks, or countries. In fact, we found that getting broad coverage in China was difficult because existing platforms such as RIPE Atlas or PlanetLab (There are very few available nodes after the membership of CERNET was withdrawed in 2009.) do not provide hosts that meet our needs. Instead, we focus on three hosts located near prominent Chinese IXPs and thus allow us to receive mappings to a replicas in a variety of Chinese regions [52]. Rather, we focus on a small number of vantage points and CDNs serving the most popular 1 million Web sites according to Alexa. The 6 chosen Chineses CDNs account for about 90% Chinese markets [6] and the corresponding customer sites we identified are all Chinese sites. We found this to be sufficient to reveal interesting differences between US and Chinese server selection dynamics. We expect server selection dynamics under the same management domain, *i.e.,* CDN, to show some consistency across vantage points and customer sites.

Second, we do not provide a comprehensive study of each CDN considered, nor do we attempt to provide a complete characterization of CDN performance. We focus primarily on potential replica server selection quality and the corresponding latency performance, which is critical for many aspects of CDN performance, and is most impacted by China's network and regulatory framework.

### A. MEASUREMENTS

The goal of our measurements is to provide a comparative analysis of representative CDNs from a geospatial, temporal, and customer point of view. We started by determining how popular content is served by CDNs. Specifically, we resolved all Web sites from the Alexa World Top 1 million sites (as of April 2, 2016) from three vantage points China (Beijing, Hefei and Xi'an) and 1 vantage point in the US (Boston). Picking three widely distributed and topologically core exchange points in China allows us to identify region-specific CDN behavior, and the location in Boston is used to compare CDN behavior outside of China. Because Web sites frequently join and leave commercial CDNs, we first

**TABLE 1.** Summary of target CDNs. Some CDNs assign unique CNAME (custom) for each customer while some put many customers on the same CNAME (shared). Some CDNs allow customers to choose limited regions for delivery service as noted in "Regional". * indicates that we could not confirm using public sources. + Regional service depends on the selected CDN.

| CDN provider | CDN Country | # found | # measured (used) | last CNAME | Redirection | Regional |
|---|---|---|---|---|---|---|
| ChinaNetCenter | CN | 876 | 876 (778) | shared/custom | DNS | Y |
| ChinaCache | CN | 304 | 304 (288) | shared/custom | DNS | Y |
| FastWeb | CN | 246 | 246 (232) | shared/custom | DNS | Y |
| Baidu (Yunjiasu) | CN | 501 | 501 (455) | custom | DNS | N |
| AliCDN | CN | 456 | 456 (430) | custom | DNS | Y |
| Tencent | CN | 201 | 201 (185) | shared/custom | DNS | Y |
| CDNetworks | KR | 355 | 355 (339) | custom | DNS | Y* |
| Cloudflare | US | 17022 | 50 (49) | custom | Anycast | N |
| Akamai | US | 12319 | 125 (127) | custom | DNS | Y* |
| AkamaiCN | US | 388 | 388 (381) | custom | DNS | Y* |
| Fastly | US | 1813 | 170 (168) | shared/custom | DNS/Anycast | Y |
| Cloudfront | US | 1430 | 50 (48) | custom | DNS | Y |
| Edgecast (Verizon) | US | 486 | 486 (405) | shared/custom | Regional Anycast | Y* |
| Level3 | US | 355 | 355 (347) | custom | DNS | Y* |
| Cedexis (CDN selector) | US | 43 | 43 (43) | N/A | DNS | + |
| Total | | 36795 | 4606 (4275) | | | |

had a bootstrap domain name resolution of all 1M DNS host names, identified and sampled target sites in order to balance budget probes and fine-grained interval dynamics. Then we measured target sites and conducted a data sanitization before the final analysis.

Note that we also filtered out possible censored Web sites using end-to-end test. We first tried to fetch the landing page of each Web site, if we get error response status or connection-timeout, we discard that site. In total, we filtered out 1035 sites (among them only 133 sites using CDNs, which has little impact) and the bootstrap resolution identified 36,795 websites using CDNs (Table 1).

Each CDN serves multiple customer sites, and exhibits similar redirection behavior for subsets of those customers. Thus, to reduce the number of measurements (which allows us to improve measurement time-granularity), we sampled sites served by Akamai, Cloudflare, Cloudfront and Fastly in a way that captured different redirection behavior, as we explain below. This resulted in 4,606 sites that we monitor for the remainder of the study.

We measure each of these 4,606 Web sites approximately once every two hours from three vantage points in China (Beijing, Hefei and Xi'an) and one vantage point in the US (Boston) from April 22nd to May 2nd in 2016. In total, we gathered approximately 2.8 million DNS resolutions (including intermediate name resolutions) and 13.2 million RTT measurement records. We describe these measurements in detail later in this section.

### *a: SAMPLING*

We sampled CDN-served Web sites as follows. Cloudflare uses global anycast based on BGP prefixes, so we selected 50 customer sites that cover all of the 29 routable prefixes we observed.

For CDNs that use CNAME redirection, we first grouped Web sites according to the last CNAME resolved (Table 1). We expect that Web sites that use the same authoritative domain or the same last CNAME will exhibit similar characteristics. These CNAMEs often include information that identifies the CDN customer, *e.g.,* Akamai resolves `www.apple.com` to `e6858.dscc.akamaiedge.net` where `6858` is the customer identifier and the remainder of the content identifies the type of service [43].

To sample such CDN customers, we picked Web sites that covered a diverse set of CDN service offerings subject to the constraints of probing only a few hundred sites per CDN. For example, we observed that Fastly serves 1,813 sites but uses only 170 unique CNAMEs, and thus measured only the most popular site corresponding to each CNAME.

In the case of Akamai, there were 5,847 unique CNAMEs (out of 12,319 total) and 68 unique authority names. Among them, 57 are used by more than one customer, so we sample the most popular two customers using each such authority name. Table 1 summarizes the identifiers we found for each CDN and how many we sampled.

### *b: MEASUREMENT DETAILS*

We have the same measurement process across all vantage points. In each round spaced by two hours, we first resolve all target Web sites using `dig` and keep the entire DNS resolution chain and server set for later analysis. This allows us not only to understand CDN redirection dynamics, but also tells us if a Web site changes CDN providers during our study.

Immediately after resolving each DNS name for each site, we record 15 round-trip time (RTT) latency of TCP-SYN pings on port 80 to the resolved IP(s) and all the IPs previously resolved for that site within the past recent week from that vantage point. We call this set of IPs the **candidate set**. This allows us to understand the latency and availability [43] not only to the currently selected servers, but also to any previously selected servers that could have potentially served the site's content. Before starting our 10-day measurement study, we resolved Web site names for three days as part of a "warming up" phase to ensure similar-sized candidate sets for the duration of our experiments.

At the conclusion of the measurement study, we checked whether any sites changed CDN providers. In fact, we found that 331 out of 4,606 measured sites changed CDN providers or stopped using CDNs. We filtered these sites from our results. The final set of measured sites is summarized in Table 1.

### B. IDENTIFYING CDN-HOSTED SITES

Commercial CDNs provide several ways for customers to enable their service, including CNAME redirection, domain hosting, and/or URL rewriting [47]. Based on this observation, we identify CDN-hosted sites based on naming conventions and semantic regular expression patterns similarly as previous work [43] and/or posted publicly by CDNs *e.g.,* [23].

A drawback to this approach is that we may incorrectly assume that a site is using a CDN to deliver content, when in fact it uses a CDN only for DNS resolution (*e.g.,* to defend against DDoS attacks). Further, a Web hosting company (*i.e.,* a company that serves multiple domains from a single location) may exhibit similar characteristics to a CDN in that multiple domains are hosted by a third party. To distinguish these cases from sites using CDNs to deliver content, we use ground truth data and conservative filters as follows:

1) We manually inspect the suspected CDN domain's Whois information to ensure it belongs to a known CDN company (*e.g., cdngc.net* and *panthercdn.com* belong to CDNetworks).
2) We use publicly disclosed IP ranges [4], [16], [17], [22] and CDN-provided tools [11] to identify CDN replica server IPs.
3) We use the number of IPs seen from multiple vantage points to filter out cases of Web hosts (*i.e.,* those that host multiple independent domains at a single location). Some Chinese CDNs serve customers from one set of servers for all our vantage points, so we further investigate each site using larger set of public DNS servers [30]. If we still see a small number of IPs from a large number of DNS vantage points, we discard the site as using a Web hosting provider and not a CDN. The exceptions to this rule are Edgecast, Cloudflare and (partially) Fastly, who use anycast. Another exception is known Chinese CDNs that serve customers from a single region.

Given the paucity of ground truth that maps Web sites to CDN providers, we cannot evaluate the accuracy of these mappings. However, our subsequent analysis reveals that the customers assigned to each CDN by our methodology exhibit clusters with similar dynamics and performance, suggesting that our mappings are likely correct.

### C. REGION-SPECIFIC CDN SERVICE

Many CDNs today provide service that varies according to the region and customer. For example, Fastly provides options to limit content delivery in only US/EU [24] and CloudFront has three-tier pricing model with different content delivery regions [18] to allow customers to balance between cost and performance. Further, most Chinese CDN providers by default only provide content delivery service in China, providing global service only on an opt-in basis.

In our subsequent analysis that compares server selection quality of CDNs with and without PoPs in China, it is important to distinguish cases where a customer signs up for CDN service that is China-specific. Unfortunately, identifying that this is the case is not trivial because neither customers nor CDNs are required to make this information public. We use a simple *heuristic* to identify and account for this when evaluating CDN server selection: if a Web site is consistently resolved to edge servers in the same country/region from both China and the US, we determine that it is using a region-specific CDN service. We first rule out cases where content delivery service is only provided in a single country other than China.

For instance, even though CDNetworks has servers in China and the US, our vantage points consistently are directed to CDNetworks servers in a single, third country for certain customers (*e.g.,* Japan and Korea). Manual checks show that most of them are served in the same language as the country it is served from. Another case is that some Chinese CDNs serve customer content from replicas in a country outside of China because the customer did not have an ICP license at the time we measured.

We also found 362 websites that use different CDN providers in China and the US. In many cases, a site uses a Chinese CDN inside China and a global CDN outside China. In addition, there are 43 sites that use a CDN selector provider, Cedexis.

For the remainder of the paper, we focus only on sites that use global and China-specific CDNs due to our focus on China. For the global CDNs that contract with Chinese CDN providers, we use different names to identify the China-specific deployment: AkamaiCN and CDNetworksCN. We did not observe other global CDNs with presence in mainland China.

## IV. CHINESE CDN BEHAVIOR

We first investigate how Chinese CDNs, which are adapted to China's unique environment, exhibit significantly different behavior from global ones.

### A. SERVER SELECTION DYNAMICS

When using DNS-based redirection, CDNs can dynamically send clients to replicas in a way that optimizes criteria such as client-perceived latency and bandwidth costs. Previous work [51] in fact showed that one large CDN redirected clients primarily based on network conditions. Given the limited network diversity and low bandwidth availability in China, an important question is whether such performance-based redirections are present and effective in Chinese CDNs.

We first examine the replica server selection dynamics of each CDN provider to show how Chinese CDN

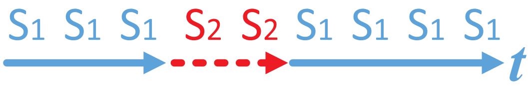

**FIGURE 1.** A simple example of normalized rate of server set changes. If a client observes this times series of server set for a customer site, the rate for the site is 0.25.

providers adapt to China's network environment. As previously described in Section III, we issue a DNS request for each targeted website about once every two hours, and store the entire resolution chain for analysis. Thus for each website, we have time series of resolved server IPs in each round for each web site, which we call *server sets*.

This section focuses on redirection behavior that impacts network paths to replicas, not those within a cluster of servers located in the same PoP, so we first map all server IPs to routable prefixes [20] and aggregate server sets by prefix. If two consecutive resolved server sets differ in routable prefix, we count it as a change in server sets; otherwise, we label them as staying the same. We expect server selection dynamics under the same management domain, *i.e.*, CDN, to show some consistency across vantage points and customer sites.

For each Web site, we define the *normalized rate of server set changes* as: the number of times the resolved server set changes divided by the number of total measurements minus one. A simple example is in Fig. 1. If the resolved server set does not ever change, the rate is 0; if it changes with every measurement round, the rate is 1. We then plot the CDF of server set change rates for each CDN provider, with each point referring to the rate for each CDN customer (Fig. 2).

The four subfigures in Fig. 2 show that server set dynamics for about 2600 Web sites hosted by Chinese CDNs tend to always see a fixed server set while sites hosted by global CDNs tend to have much more dynamic server sets. In stark contrast to global CDN providers, incumbent Chinese CDN providers (ChinaCache, ChinaNetCenter, FastWeb) are nearly *static* (*i.e.*, fixed), with the exception of 7% of ChinaCache customers seeing relatively significant dynamics. We manually checked the DNS records for these customers, and identified that ChinaCache appears to have a new dynamic service that uses CNAME of the form `newdynamic023.cnc.ccgslb.net`. Thus, we believe the reason for these dynamics is a customer-selected service agreement that uses dynamic redirections based on factors such as geolocation, ISP, and network conditions.

Among Chinese CDNs, the new entrants to the CDN market (Baidu, AliCDN and Tencent) are slightly more dynamic than the incumbents. We can only speculate why this might be. For example, Baidu cooperates with Cloudflare and its customers are sometimes redirected to Cloudflare's IPs. These may be used for DDoS mitigation, for example.

We did not plot Edgecast and Cloudflare in the figure because they use anycast and did not change anycast IPs, *i.e.*, normalized rate of server set changes is 0, for each site during our study. However, we note that for 388 (out of 405) customers of Edgecast, all three Chinese vantage points are sent to the same anycast IP block, while the US vantage point is sent to another. Thus, Edgecast uses a form of regional anycast.

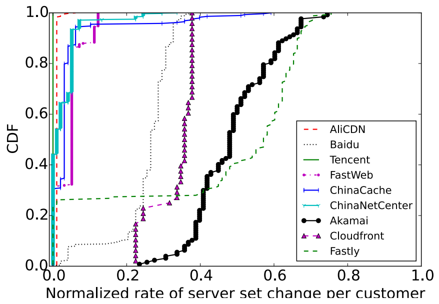

(a) Beijing

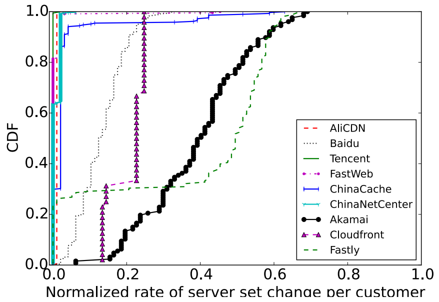

(b) Hefei

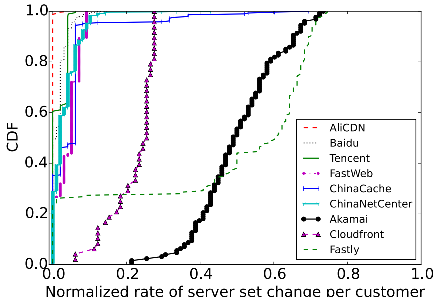

(c) Xi'an

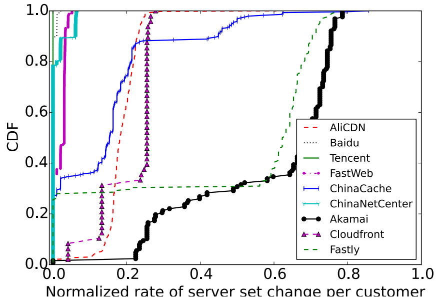

(d) Boston

**FIGURE 2.** Normalized rate of server set changes over customers of each CDN. The figures show that ∼ 2600 Web sites hosted by Chinese CDNs are nearly static (*i.e.*, the rates are close to 0) and global CDNs are much more dynamic. Interestingly ChinaCache and AliCDN behave differently (using a higher rate of server set change) when accessed from outside of China.

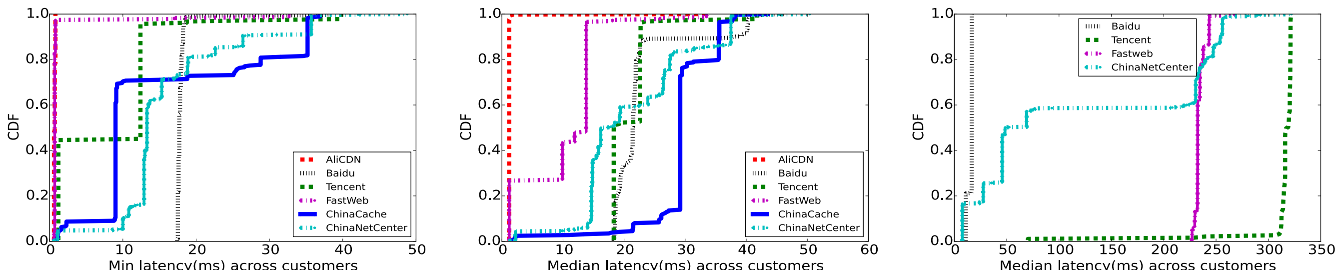

(a) Minimum latency in Beijing (b) Median latency in Beijing (c) Median latency in Boston

**FIGURE 3.** Customer latency performance of static Chinese CDNs, showing that different customers are assigned to servers with different latencies.

Fastly also provides anycast service [21] if customers are hosting their sites using an apex domain.[1] Thus for these Fastly customers, their rate of server set changes is 0 as in Fig. 2. Like Edgecast, Fastly uses regional anycast, *i.e.,* all three vantage points in China see the same IP for a given customer while different IP is seen from the US.

## B. CUSTOM SERVICE

The recent entry of AliCDN, Baidu, and Tencent has disrupted the CDN market in China, and led to interesting public disclosures and claims [10] about competing Chinese CDN offerings. One such revelation is that the incumbent Chinese CDNs (ChinaCache and ChinaNetCenter) provide highly customized and configurable content delivery service [8], [25], where customers are allowed to choose where their content is served from according to the location of the bulk of the users and their price they are willing to pay. As a result, we expect that different customers contact different CDN replicas, and this should lead to observable differences in the latency performance.

To evaluate this hypothesis, we plot a CDF of the minimum and median latencies for each CDN customer, as seen by our vantage point in Beijing as in Fig. 3 (cases in Hefei and Xi'an are similar). For comparison, we also plot the median latencies seen from outside of China.

As expected, Ali and Baidu provide the same service for all of its customers as disclosed in public documents. The median latency for about 10% Baidu's customers is 40ms. Raw data show that this median latency comes from that some customers changed their IPs to Baidu's partner – Cloudflare's infrastructure (more in Section V-A), which is very likely due to DDoS mitigation.

In contrast, ChinaCache and ChinaNetCenter exhibit distinct behavior, with clusters of customers receiving different latency performance. Importantly, these latency differences are very unlikely to be due to redirection dynamics, because we found in the previous section that the mappings of clients to replicas was relatively static. As a result, the likely reason for this behavior is that clients are sent to replica sets that depend on the CDN customer, and that different CDN customers choose different levels of service from the provider.

We also investigated static Chinese CDNs performance for our US vantage point. Fig. 3 (c) shows that Baidu and Cloudflare (Fig. 8) have nearly identical performance due to that Baidu's customers are served by Cloudflare's infrastructure through partnership. ChinaNetCenter and Tencent provide a small number of replicas outside of China, but the figure suggests that their customers rarely pay for the opportunity to use them. For example, we found only two Tencent customers using service provided by servers outside of China (using the CNAME x2.tcdn.qq.com).[2] We did not observe any FastWeb servers outside of China.

### a: Replicas Intersection of Static Chinese CDNs

Given *static* server selection for clients of Chinese CDNs, the clusters of customers receiving similar performance indicate there exist different classes of service, especially when some are resolved to the same shared CNAME. We now investigate whether the replica server sets assigned to a customer are specialized according to client location, or are more likely specialized according to a customer-selected service region [25] (*e.g.,* three vantage points all see the same replica).

To identify this behavior, we focus on whether different Chinese vantage points see different replica server sets for each customer. If they all see one server set, then it is likely that the customer has opted for content to be served from a single CDN location. If they see different server sets, then it is likely due to using multiple CDN locations to serve content. The larger the number of unique server sets, the more opportunity there is for the CDN to provide clients with low latency.

We plot this using the probability distribution function (PDF) of the number of replica server sets seen for customers of each CDN, across all three Chinese vantage points (Fig. 4). For each customer, we find the distinct server sets seen more

---

[1] Also called a "root," "bare," or "naked" domain, *e.g., example.com* is an apex domain but *www.example.com* is not

[2] Tencent owns qq.com

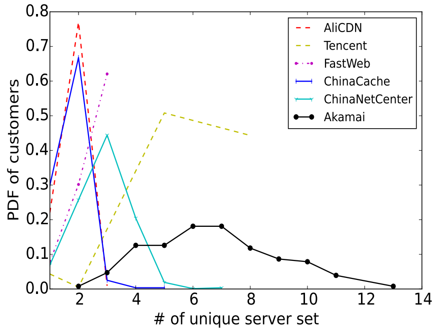

**FIGURE 4. PDF of number of unique server sets assigned to each customer across all three Chinese vantage points. Curves with more density to the left indicate customers using regional service limited to a small set of replicas, while those that are flatter indicate customers being served from a large set of replicas.**

than 5% of the time by each vantage point.[3] We then count the number of distinct server sets across all three Chinese vantage points for each customer. Curves closer to $x$=1 represent a large set of customers choosing to be served from a small set of regions, while those that extend to the right of the figure indicate those served by a diverse set of replica-server locations.

The figure shows that AliCDN and ChinaCache exhibit the lowest replica diversity, with few (but not zero) customers served from different replicas across all three regions. Customers of Tencent have a slightly higher level of server diversity. By comparison, Akamai has substantially more server diversity, likely due to the typical case where customers do not select service regions within a country.

### C. MULTI-CDN SELECTOR BEHAVIOR

Several Web sites are served by a multi-CDN selector, Cedexis. This provider contracts with multiple CDNs on behalf of a customer, and according to their documentation, they select which CDN to serve clients based on real time latency information [29].

For the 43 sites that using Cedexis (Section III-C) we found that for 16 of them, only a single Chinese CDN is ever seen by our vantage points. We focus our analysis on the seven sites that use both ChinaCache and ChinaNetCenter, to observe the impact of Chinese CDNs on multi-CDN selector performance.[4]

First, we investigated redirection dynamics in terms of which CDN the selector chooses at any point in time. We found that the seven sites using two Chinese CDNs all saw identical redirection dynamics (Fig. 5(a)). Thus, if Cedexis changed a CDN for one customer, it changed it for all customers using the same service.

Recall, however, that Chinese CDNs provide highly customized service to different customers, meaning that switching between CDNs for all customers may lead to different performance impact for these customers. When we investigated the latency performance for these customers, we found that while one customer saw improved latency, the other saw *worse* latency. In Fig. 5 (b), RTT means latency to the selected server and RTT difference means latency of selected server minus minimum latency among **candidate servers** in the same round. Importantly, to optimize latency performance when using Chinese CDNs like ChinaNetCenter and ChinaCache, one must take *the customer* into account when making CDN selections. By failing to do so, Cedexis cannot guarantee optimized performance for all its customers.

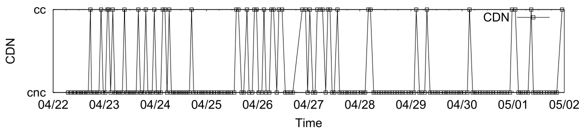

(a) Cedexis alternates between using ChinaNetCenter (cnc, bottom) and ChinaCache (cc, top), and makes identical decisions for all customers.

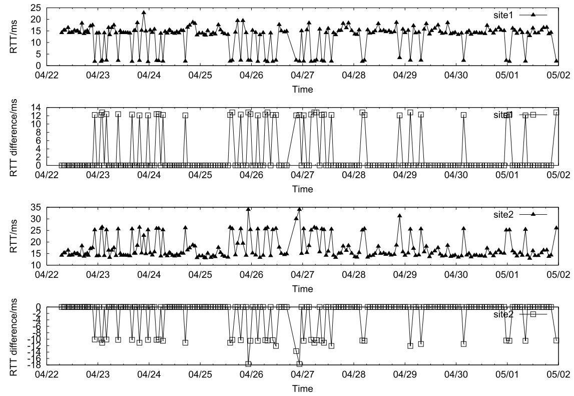

(b) As a result of the above behavior, the same client will see better (top) or worse (bottom) latency depending on the customer.

**FIGURE 5. Time series showing that Cedexis makes decisions at the CDN granularity, but leads to suboptimal performance due to Chinese CDNs serving different customers from different replicas.**

### D. SUMMARY

We examined redirection dynamics and found that a large fraction of Chinese CDNs provide relatively fixed mappings between clients and replica servers for each CDN customer. This is in stark contrast to the behavior exhibited by many global CDNs in the same period. We also investigated that Chinese CDNs provide highly customized service at customer granularity, allowing region-specific replica selection.

## V. CHARACTERISTICS SUMMARY

This section takes a preliminary look at how CDN providers deploy and implement CDNs in mainland China. We compare latency and reliability analysis of major global and Chinese CDNs. By showing an intuitive difference, we will dive deeper into its impact on server selection dynamics of global CDNs in Section VI.

---

[3] Because we are focusing on relatively static mappings of customers to replicas, we do exclude customers of ChinaCache with normalized rate of server set change greater than 0.05.

[4] The remaining 20 sites do not use any Chinese CDNs.

### A. OVERALL LATENCY PERFORMANCE

We first investigate the performance of CDNs in terms of the RTT latency between our vantage points and CDN-selected servers. This analysis reveals the potential impact of ICP licensing requirements on the minimum latency that CDNs provide their clients in China.

We begin by focusing on measurement results from Beijing, given that it is one of the *best-connected locations* in China, and is a *major exchange point with external networks*. As a result, it should experience low latency to any CDN replicas, if existing, in China. For each CDN, we aggregate the latency measurements for each /24 prefix, under the assumption that servers in the same /24 are co-located and should exhibit similar network-induced latencies.[5] We then plot the CDF of minimum latencies[6] for each /24 prefix measured from Beijing, for all Chinese CDNs and three typical global CDNs. We omit other global CDNs to make the graph more readable, and described them below.

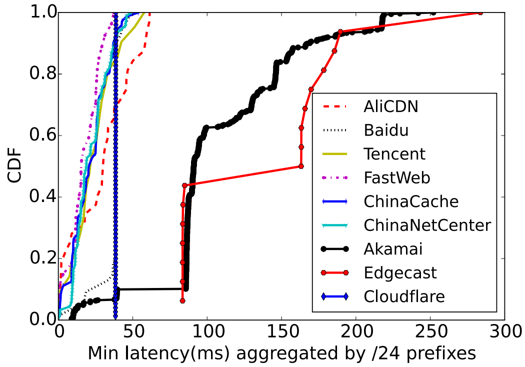

**FIGURE 6.** **CDF of minimum latency between clients and replica servers, grouped by CDN customer, for a vantage point in Beijing. The figure clearly shows a gap between 50–90 ms that separates CDNs with presence in China (*e.g.,* ChinaCache and ChinaNetCenter) and those that do not (*e.g.,* Edgecast). This gap is likely caused in part by the ICP license policy, which makes it difficult for CDNs to host content in China without partnerships with Chinese companies. Interestingly, Akamai uses servers in China for only a small subset of its customers (those with latency less than 40 ms).**

Fig. 6 shows the CDF of latencies, indicating a clear boundary in latencies between Chinese CDNs and global ones. The latency profile in Hefei and Xi'an is similar with Beijing. Chinese CDNs map our Beijing client to nearby replicas, demonstrating the benefit of having presence in mainland China. Cloudflare uses anycast to map clients to replicas, leading to identical latency across all its /24s. Reverse DNS lookups on traceroute hops suggest these servers are located in Hong Kong.

Interestingly, Akamai has a small fraction of low latency clusters, all of which are resolved using a name of the form *.tl88.net.[7] This is due to Akamai's ChinaCDN module option [2].

Cloudflare claims to partner with Baidu to provide service in China as of September, 2015 [15]. We did not see evidence of Cloudflare IPs located in China, but we expect that this is due to their approach to providing service in China. Specifically, they appear to use Baidu's DNS infrastructure to serve content, which prevents us from identifying which Baidu-hosted sites are Cloudflare customers.[8]

We now summarize CDNs omitted from the figures. These CDNs often use a replica in an Asia-Pacific PoP and occasionally use one that is located more than 100 ms away. Fastly uses servers in Hong Kong and Japan, but its lowest latency was 70ms. CDNetworks uses servers in Korea (40 ms), and Cloudfront's nearest observed servers are in India (120 ms). CDNetworks exhibits performance similar to Akamai, and also provides a ChinaCDN option [7]. Level3 claims to have a Beijing site, but we did not observe one; rather, most latencies from Level3 are ≈150ms from Beijing, with a minimum of 50 ms (via a server in Hong Kong). No other commercial CDN providers claim to have presence in China according to their public documents.

In summary, when serving content for Chinese clients, there is a clear advantage in terms of latency for CDN providers that can use replicas located in China. Thus, the ICP license regulation can have a significant impact on performance for global CDNs and motivates the need for cooperation with Chinese CDNs to improve response times for Chinese users.

### B. OVERALL RELIABILITY

This section analyzes CDN reliability in terms of ping response failure rates for each customers, defined as the fraction of port 80 TCP-SYN pings that do not receive a response. We compute the failure rate for each CDN customer as seen from Beijing, and plot a CDF of customers' failure rates for each CDN.

Fig. 7 shows a CDF of failure rates, using a log scale on the x-axis. Note that all CDNs exhibit generally high reliability, with most exhibiting between 3 and 4 nines of reliability. In fact, for most CDNs, the majority of customers see 100% availability, except that some CDNs have worse overall reliability than others. Baidu, CDNetworks, and Akamai have lower than 3 nines of availability for more than 10% of their customers.

We also examine the average failure rate across all customers observed for each CDN. Table 2 shows that CDNetworksCN, ChinaNetCenter, Akamai, and Baidu all have lower-than-average reliability.

---

[5] We do not explore the impact of server load on latency.

[6] Note that the CDF of median latencies was qualitatively identical.

[7] Akamai has 388 (out of 12,319) customers that exhibit this behavior. For example, www.apple.com is resolved to e6858.dscc.akamaiedge.net from the US, and to e6858.e19.s.tl88.net in China. We confirmed that all tl88.net servers are in fact hosted on Chinese IPs (from China Unicom and China Telecom).

[8] Cloudflare prevents a single customer from being served by both Cloudflare and Baidu at the same time [14].

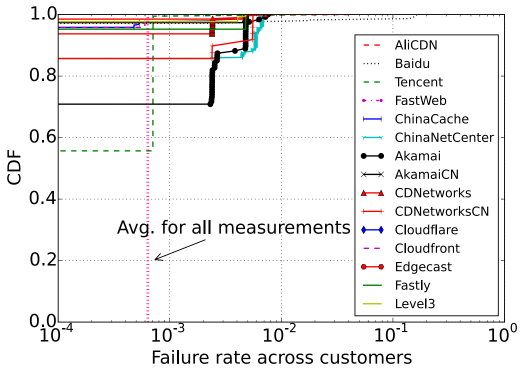

**FIGURE 7.** Several CDNs exhibit lower than 3 nines of reliability for more than 10% of their customers, though the majority servers for all CDNs have 100% availability.

**TABLE 2.** Average failure rate (and standard deviation) for each customer in Beijing. All Chinese CDNs except for FastWeb have higher failure rate. Akamai also exhibits a high failure rate, likely due to the long paths to replicas.

| CDN | Average | Standard dev. |
|---|---|---|
| FastWeb | 0 | 0 |
| Cloudflare | 0 | 0 |
| Cloudfront | 3.74x10$^{-5}$ | 1.46x10$^{-4}$ |
| Edgecast | 8.35x10$^{-5}$ | 6.30x10$^{-4}$ |
| Level3 | 1.12x10$^{-4}$ | 6.87x10$^{-4}$ |
| ChinaCache | 1.16x10$^{-4}$ | 7.05x10$^{-4}$ |
| AkamaiCN | 1.38x10$^{-4}$ | 8.02x10$^{-4}$ |
| CDNetworks | 2.05x10$^{-4}$ | 8.42x10$^{-4}$ |
| AliCDN | 2.21x10$^{-4}$ | 2.20x10$^{-3}$ |
| Fastly | 2.56x10$^{-4}$ | 1.08x10$^{-3}$ |
| Tencent | 4.17x10$^{-4}$ | 1.43x10$^{-3}$ |
| **Overall average** | **6.54x10$^{-4}$** | **1.25x10$^{-3}$** |
| CDNetworksCN | 7.12x10$^{-4}$ | 1.74x10$^{-3}$ |
| ChinaNetCenter | 8.48x10$^{-4}$ | 2.09x10$^{-3}$ |
| Akamai | 1.08x10$^{-3}$ | 1.84x10$^{-3}$ |
| Baidu | 2.61x10$^{-3}$ | 1.92x10$^{-2}$ |

### C. SERVER SET DIVERSITY

Previous work [51] found that CDN redirection dynamics were primarily driven by providing low latency. If this is the case, we would expect to see redirection dynamics to lead to replicas in multiple network locations. We now test this hypothesis by looking at the diversity of replica server network locations. We will evaluate how this diversity impacts the client-perceived latency for each CDN in Section VI.

We begin by mapping the replica server sets for each CDN customer to their BGP routable prefixes [20], and find the number of unique prefixes for each customer during the measurement period. Then we group the data by CDN, and compute the mean (and standard deviation) over the CDN's customers (Table 3).

We find that some global CDN providers exhibit a high rate of server set changes for Chinese vantage points as in Fig. 2 (Fastly and CDNetworks), but these redirections are to server IPs in a small number of network locations. For most CDNs, the number of unique prefixes used is relatively small (2-4 prefixes) in Beijing, with a few exceptional cases for other vantage points. Notably, Akamai and Tencent have the greatest server diversity in terms of unique BGP prefixes.

**TABLE 3.** Average (standard deviation) of unique routable prefixes per customer for each CDN. Although global CDNs dynamically select servers for each customer site, with few exceptions they serve only from a small number of unique BGP prefixes.

| CDN | Beijing | Hefei | Xi'an | Boston |
|---|---|---|---|---|
| Edgecast | 1.00 (0.05) | 1.00 (0.05) | 1.01 (0.14) | 1.00 (0.04) |
| Cloudflare | 1.53 (0.5) | 1.53 (0.5) | 1.53 (0.5) | 1.57 (0.58) |
| CDNetworks | 2.32 (0.89) | 2.31 (0.86) | 2.39 (0.91) | 2.08 (0.9) |
| Fastly | 2.38 (0.87) | 2.35 (0.82) | 2.35 (0.82) | 2.52 (0.84) |
| Fastweb | 2.38 (0.89) | 1.44 (0.58) | 2.81 (0.71) | 2.54 (0.67) |
| Level3 | 2.56 (1.63) | 2.77 (1.61) | 4.31 (2.16) | 2.56 (1.45) |
| AliCDN | 2.84 (0.84) | 2.83 (0.77) | 1.01 (0.11) | 10.5 (3.13) |
| CDNetworksCN | 3.2 (2.4) | 6.69 (2.39) | 2.97 (1.66) | 2.75 (2.45) |
| ChinaCache | 4.02 (3.78) | 10.21 (8.78) | 5.38 (3.78) | 12.7 (8.14) |
| Cloudfront | 4.08 (0.27) | 3.64 (0.93) | 4.45 (0.82) | 6.29 (1.42) |
| Baidu | 4.38 (0.91) | 4.15 (0.85) | 1.75 (0.85) | 2.08 (0.35) |
| Chinanetcenter | 4.81 (3.33) | 2.75 (2.33) | 4.36 (3.17) | 2.41 (1.99) |
| Tencent | 5.72 (1.36) | 7.11 (1.94) | 6.62 (1.47) | 7.13 (2.1) |
| AkamaiCN | 6.85 (1.69) | 7.45 (2.11) | 6.82 (2.26) | 13.5 (3.2) |
| Akamai | 7.10 (2.73) | 9.19 (3.44) | 10.71 (4.44) | 11.28 (5.2) |

We also found that Cloudfront exhibits relatively high IP-level diversity (approximately 60 IPs per vantage point) but a relative low diversity at the routable prefix level, likely indicating load balancing within a data center.

Two Chinese CDNs, ChinaCache and AliCDN, exhibit vastly different behavior when accessed from China versus from the US. They dynamically select a wide range of servers when accessed from outside China, but use a static, small set of replicas for Chinese clients.

Anycast CDNs exhibit clear and unsurprising behavior, since Edgecast assigns one (anycasted) IP address for each customer and Cloudflare (as well as Baidu's service outside China, which uses Cloudflare) usually configures two IPs per customer.

### D. SUMMARY

This section provides a preliminary look into CDN behavior when viewed from Chinese vantage points. We found that China's ICP licensing requirement can have a significant impact on client-perceived latency in China. Specifically, global CDNs that partner with Chinese ones can achieve 50% or more reduction in latency compared to using servers outside China. In terms of reliability, we found that most CDNs offered 100% availability for most of their customers during our study, but a number of CDNs provided much lower availability (*e.g.,* only 2 nines) for some of their customers.

## VI. GLOBAL CDNs IN CHINA

In this section we explore the behavior of global CDNs and their impact on performance in China. Specifically, we look at anycast and dynamic DNS-based CDNs. Latency profiles to these foreign CDNs at customer granularity in Beijing, Xi'an and Hefei are very similar, likely due to a lack of network diversity. Consequently, in the following we simply use Beijing in our measurements.

### A. ANYCAST CDN CLIENT-MAPPING

Recent work [41] highlights that BGP anycast is often used to support TCP applications at large scale, including CDNs. We found that Edgecast, Cloudflare and some of Fastly customers are served via anycast addresses.

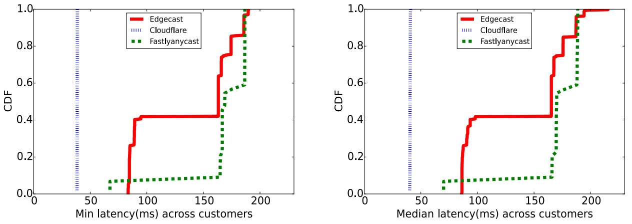

(a) Minimum latency in Beijing (b) Median latency in Beijing

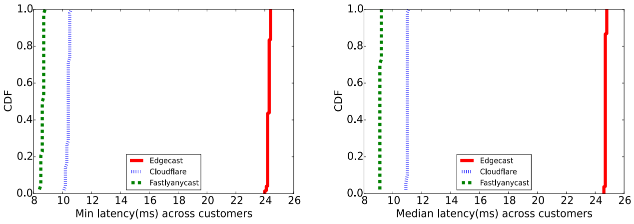

(c) Minimum latency in Boston (d) Median latency in Boston

**FIGURE 8.** **Minimum and median latencies for anycast CDNs observed from Beijing and Boston. Regional anycast causes inconsistent performance for different customers.**

To understand the impact of anycast on performance, we plot the minimum and median latency across customers of these three CDNs, from Beijing and the US, in Fig. 8. The minimum and median latencies are nearly identical, indicating stable performance for clients of these anycasted services. This is likely explained by findings in previous work [39], [41] showing that for the anycast CDNs they measured, most clients show high *front-end affinity*, *i.e.,* clients are sent to the same front-end over time.

We now compare regional and global anycast performance. Cloudflare announces all of its prefixes from all data centers [13] and our clients saw the same IPs for a given customer from all four vantage points. This implementation leads all of our vantage points to see consistent performance for all customers.

In contrast, Edgecast and Fastly use regional anycast, *i.e.,* configure customers with different IPs for different regions and announce different IP prefixes from different data centers. Interestingly, while this has nearly no impact on performance for clients in the US, we find substantial differences for clients in China. For example, some Edgecast customers receive 40% worse latency than others. To understand why this is the case, we used BGP feeds and traceroutes with reverse DNS lookups to isolate the likely location of data centers that our vantage points were redirected to. We found that the lowest latency PoP was located in Singapore ($\sim 80ms$ from China), with the remaining cases resulting from clients being routed to Las Vegas, Los Angeles ($\sim 170ms$), and San Jose ($\sim 200ms$).

### B. DYNAMIC REDIRECTION QUALITY

In this section, we evaluate server selection quality of DNS-based dynamic CDNs in terms of the RTT latency between clients and servers. In particular, we seek to understand whether each CDN is redirecting clients based on the best possible latency it can provide for a given customer. This is in general a difficult problem, because we do not know in advance the complete set of replicas from which a CDN is willing to serve a customer's content. As such, comparing the latency of a selected replica against *all of a CDNs replicas for all customers* is potentially unfair.

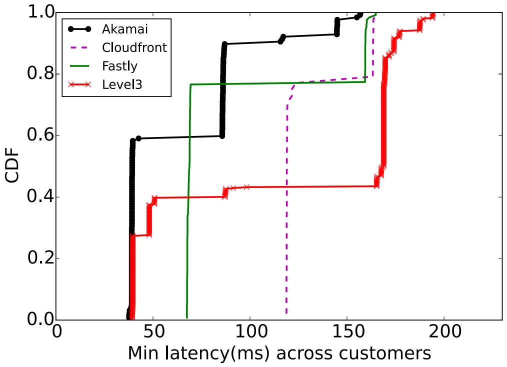

**FIGURE 9.** **Minimum latency across customers of dynamic CDNs, which we use to identify region-specific groups in Table 4.**

**TABLE 4.** **Categorization of customers by regional content delivery service. Each category of CDNs ordered from lowest minimum latency to highest.**

| Provider | A | B | C |
|---|---|---|---|
| Akamai | 0-50ms | 50-100ms | 100ms+ |
| Fastly | 0-100ms | 100ms+ | |
| Level3 | 0-100ms | 100ms+ | |
| Cloudfront | 0-150ms | 150ms+ | |

Instead, we conduct our latency comparisons on a *per-customer* basis. We plot the minimum latencies seen for each customer in Fig. 9. The graph shows that each CDN offers a small number equivalent classes in terms of latency performance, which we approximate using Table 4. We analyze each CDN's redirection quality according to these classes.

Specifically, for each round, we compute the difference between the latency of selected replica and the minimum latency of candidate set as defined in Section III-A. (We group replicas by their /24 prefixes and use the minimum latency for the group.) We then plot the average of these differences for each customer in Fig. 10, with each subfigure representing different CDNs, and each curve representing different classes for each CDN (Table 4).[9] A detailed algorithm description is in Algorithm 1.

The figures show that classes with higher minimum latencies (classes B and C) see smaller latency differences among replica servers, *i.e.,* when latency is already poor, it doesn't get much worse. On the other hand, for customers that *can observe low latency from a replica*, they are sent to a non-optimal, *i.e.,* higher latency, replica more than half

---

[9] We omit AkamaiCN because it oscillates between nodes in the US and China using a diurnal pattern. While peculiar, it is not the kind of dynamic redirection we focus on in this section.

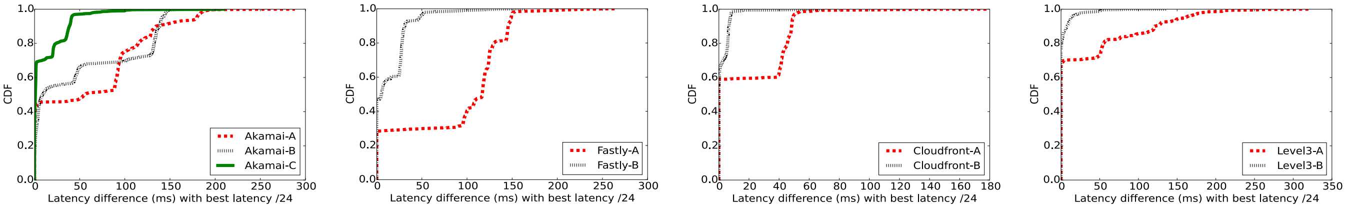

(a) Akamai (b) Fastly (c) Cloudfront (d) Level3

**FIGURE 10. Server selection quality for different regional delivery categories of each dynamic CDNs. Customers choose large coverage may experience worse server selection in terms of latency with best server clients could have been redirected to.**

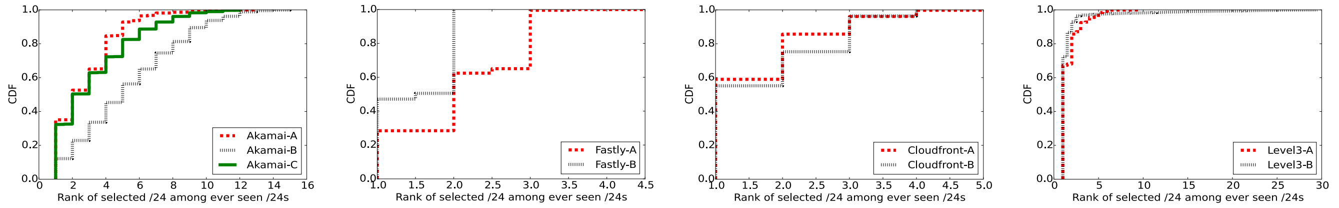

(a) Akamai (b) Fastly (c) Cloudfront (d) Level3

**FIGURE 11. Server selection quality in terms of rank, where a value of 1 represents the best replicas, 2 represents the second best replicas, and so on. For Chinese vantage points, using one of the top-ranked replicas instead of the best one can lead to significant performance loss (Fig. 10).**

**Algorithm 1** Evaluate the Quality of Server Selection

```
Input:
  $m$: a measurement round
  $S$: a set of candidate server IPs
  $L_i, i \in S$: latency of IP $i$
  $I_i, i \in S, I_i \in \{0, 1\}$: in current measurement round whether the server $i$ is selected (1) or not (0)
Output:
  $R$: ranks of selected server
  $D$: latency differences with the best in candidate set

  for each $m$ in a given class do
    $P = \{GetPrefix(i) \mid i \in S\}$
    for each $p,\ p \in P$ do
      $Latency_p = \min\{L_i \mid GetPrefix(i) == p,\ i \in S\}$
      $I_p = \bigcup I_i$, where $GetPrefix(i) == p,\ i \in S$
    end for
    $CurLatency = \text{median}\{Latency_p \mid I_p == 1,\ p \in P\}$
    $CurRank = GetRank(CurLatency, \{Latency_p \mid p \in P\})$
    $Diff = CurLatency - \min\{Latency_p \mid p \in P\}$
    $R$.add($CurRank$)
    $D$.add($Diff$)
  end for
  return $R,\ D$
```

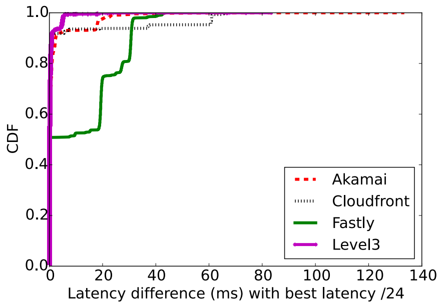

(a) Latency difference

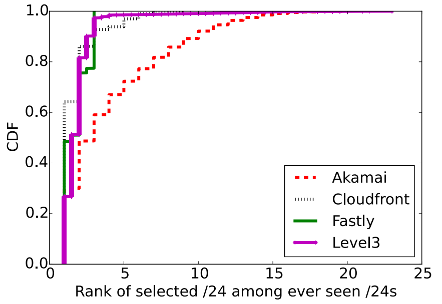

(b) Rank of selected CDN

**FIGURE 12. CDFs of latency differences and rank for selected servers seen from Boston, to compare with Figs. 10 and 11. Note that the CDF of server ranks seen from Boston is similar to those seen from China; however, in the case of Boston, the latency difference between top-ranked servers is not as large.**

the time for Akamai and Fastly. One possible explanation is that these CDNs are not selecting the *best* replicas for client, but instead select one of the top $n$ replicas, for some number of replica locations $n$. If this policy is applied globally, it is likely to have little impact on performance in the US (where there are many replica locations) but potentially devastating performance impact for Chinese clients due to the paucity of server locations near or inside China.

To evaluate this hypothesis, in each measurement round for a given customer, we also investigated the ''quality'' of redirections in terms of the rank of servers returned by each CDN. The lowest-latency server is given rank 1, second-lowest rank 2, and so on. Fig. 11 shows a CDF of selected server ranks for each CDN and CDN class. By comparison we show the same for the US in Fig. 12. We find, for example, that the Akamai does not always pick the best replica locations, but rather simply one of the best. This has little effect in the US (latencies for all customers are low) but leads to large latency differences in China. Thus, there is a potential to significantly improve performance for Chinese clients by selecting better ranked servers.

## VII. CONCLUSION

Though China represents a large fraction of Internet users and their ranks continue to grow, little research has focused on how CDNs—which are responsible for delivery some of the most frequently accessed content—perform in the country. We find that China's unique network, policies, and CDN business models lead to a wide range of CDN behavior and performance delivered to clients.

For example, licensing requirements make it difficult for non-Chinese CDNs to deploy content in China on their own, and those CDNs with access to servers in China (often through partnerships with Chinese CDNs) have a marked advantage over those who do not. We also find that, unlike global CDNs, Chinese CDNs offer a high level customization with respect to which replicas serve customer content. Given a customer site, these relatively static client-mappings along time and region-specific replica assignment can save content providers money, but also potentially leave significant performance on the table. When focusing on the performance of global CDNs, we identify several behaviors that lead to poor server selection quality in China, including regional anycast, arbitrary selection of ''relatively good'' replicas, and customer-agnostic CDN selection. We believe this paper provides useful information for optimization of CDNs in China.

## VIII. ACKNOWLEDGMENT

The authors would like to thank Shenchao Liu for help, and the anonymous reviewers for their comments.

## REFERENCES

[1] *Internet Population*, accessed on May 10, 2016. [Online]. Available: http://www.internetworldstats.com

[2] *Akamai China-CDN Module*, accessed on May 10, 2016. [Online]. Available: https://www.akamai.com/us/en/solutions/intelligent-platform/china-content-delivery-network.jsp

[3] *AliCDN SLA*, accessed on May 10, 2016. [Online]. Available: https://help.aliyun.com/document_detail/27109.html

[4] *Baidu Yunjiasu IP Range*, accessed on May 10, 2016. [Online]. Available: http://su.baidu.com/api-doc/

[5] *Baidu Yunjiasu SLA*, accessed on May 10, 2016. [Online]. Available: https://www.qcloud.com/doc/product/228/cdn%E4%BA%A7%E5%93%81%E6%9C%8D%E5%8A%A1%E4%BF%9D%E9%9A%9C%E5%8D%8F%E8%AE%AE(SLA)

[6] *CDN Whitepaper, China Academy of Telecommunication Research of MIIT*, accessed on May 10, 2016. [Online]. Available: http://www.catr.cn/kxyj/qwfb/bps/201412/P020141205319913647751.pdf

[7] *Cdnetworks China-CDN Module*, accessed on May 10, 2016. [Online]. Available: https://www.cdnetworks.com/products/china-acceleration/

[8] *ChinaCache Custom Solution*, accessed on May 10, 2016. [Online]. Available: http://en.chinacache.com/contact

[9] *ChinaCache Forms Strategic Cooperation to Build China's First IXP*, accessed on May 10, 2016. [Online]. Available: http://ir.chinacache.com/releasedetail.cfm?ReleaseID=955286

[10] *ChinaCache's Cooperation on Time Division Multiplexing Usage of CDNs*, accessed on May 10, 2016. [Online]. Available: http://companies.caixin.com/2015-06-10/100818027.html

[11] *ChinaNetCenter IP Verification Tool*, accessed on May 10, 2016. [Online]. Available: https://cloudedge.chinanetcenter.com/cn/index.html

[12] *Cloudflare*, accessed on May 10, 2016. [Online]. Available: https://www.cloudflare.com/

[13] *Cloudflare Anycast*, accessed on May 10, 2016. [Online]. Available: https://support.cloudflare.com/hc/en-us/articles/203491930

[14] *Cloudflare China-CDN Module*, accessed on May 10, 2016. [Online]. Available: https://www.cloudflare.com/china/

[15] *Cloudflare Extended Into China*, accessed on May 10, 2016. [Online]. Available: https://blog.cloudflare.com/how-weextended-cloudflares-performance-and-security-into-mainland-china/

[16] *Cloudflare IP Range*, accessed on May 10, 2016. [Online]. Available: https://www.cloudflare.com/ips/

[17] *Cloudfront IP Range*, accessed on May 10, 2016. [Online]. Available: http://docs.aws.amazon.com/general/latest/gr/awsip-ranges.html

[18] *Cloudfront Price Class for Regional Content Delivery*, accessed on May 10, 2016. [Online]. Available: http://aws.amazon.com/cloudfront/pricing/

[19] *Commercial CDN Market Share*, accessed on May 10, 2016. [Online]. Available: http://www.globaldots.com/china-andasia-content-delivery-network-cdn/

[20] *Cymru Team IP to BGP Prefix*, accessed on May 10, 2016. [Online]. Available: http://www.team-cymru.org/IP-ASNmapping.html

[21] *Fastly Anycast Option for Content Delivery With Apex Domains*, accessed on May 10, 2016. [Online]. Available: https://docs.fastly.com/guides/basic-configuration/using-fastly-withapex-domains

[22] *Fastly IP Range*, accessed on May 10, 2016. [Online]. Available: https://api.fastly.com/public-ip-list

[23] *Fastly Public CNAME Naming Scheme*, accessed on May 10, 2016. [Online]. Available: https://docs.fastly.com/guides/basic-setup/adding-cname-records

[24] *Fastly Service: Limiting POP Use to the US and EU*, accessed on May 10, 2016. [Online]. Available: https://docs.fastly.com/guides/performance-tuning/enabling-globalpops#limiting-pop-use-to-the-united-states-and-european-union

[25] *Fastweb Custom Solution*, accessed on May 10, 2016. [Online]. Available: http://www.fastweb.com.cn/a/english/products/en_static.html

[26] *How Akamai Works?*, accessed on May 10, 2016. [Online]. Available: http://research.microsoft.com/en-us/um/people/ratul/akamai.html

[27] *ICP Beian Policy*, accessed on May 10, 2016. [Online]. Available: https://support.cloudflare.com/hc/en-us/articles/209714777

[28] *IETF Draft of EDNS-Client-Subnet*, accessed on May 10, 2016. [Online]. Available: http://www.afasterinternet.com/ietfdraft.htm

[29] *Multi CDN Selector Based on Real Time Latency*, accessed on May 10, 2016. [Online]. Available: http://www.cedexis.com/blog/why-geo-routing-misses-the-point/

[30] *Open DNS*, accessed on May 10, 2016. [Online]. Available: https://dns-test.warrenguy.me/

[31] *Tencent SLA*, accessed on May 10, 2016. [Online]. Available: https://bce.baidu.com/doc/cdn/cdn_SLA.html

[32] J. Abley and K. Lindqvist, *Operation of Anycast Services*, document RFC 4786, 2006.

[33] V. K. Adhikari, Y. Guo, F. Hao, V. Hilt, and Z. L. Zhang, ''A tale of three CDNs: An active measurement study of Hulu and its CDNs,'' in *Proc. INFOCOM*, Mar. 2012, pp. 7–12.

[34] B. Ager, N. Chatzis, A. Feldmann, N. Sarrar, S. Uhlig, and W. Willinger, ''Anatomy of a large European IXP,'' in *Proc. SIGCOMM*, 2012, pp. 163–174.

[35] M. Al-Fares, K. Elmeleegy, B. Reed, and I. Gashinsky, ''Overclocking the Yahoo!: CDN for faster Web page loads,'' in *Proc. IMC*, 2011, pp. 569–584.

[36] Z. Al-Qudah, S. Lee, M. Rabinovich, O. Spatscheck, and J. Van der Merwe, ''Anycast-aware transport for content delivery networks,'' in *Proc. WWW*, 2009, pp. 301–310.

[37] H. A. Alzoubi, S. Lee, M. Rabinovich, O. Spatscheck, and J. Van der Merwe, ''Anycast CDNs revisited,'' in *Proc. WWW*, 2008, pp. 277–286.

[38] M. Calder, X. Fan, Z. Hu, E. Katz-Bassett, J. Heidemann, and R. Govindan, "Mapping the expansion of Google's serving infrastructure," in *Proc. IMC*, 2013, pp. 313–326.
[39] M. Calder, A. Flavel, E. Katz-Bassett, R. Mahajan, and J. Padhye, "Analyzing the performance of an anycast CDN," in *Proc. IMC*, 2015, pp. 531–537.
[40] F. Chen, R. K. Sitaraman, and M. Torres, "End-user mapping: Next generation request routing for content delivery," in *Proc. SIGCOMM*, 2015, pp. 1–15.
[41] D. Cicalese, J. Augé, D. Joumblatt, T. Friedman, and D. Rossi, "Characterizing IPv4 anycast adoption and deployment," in *Proc. CoNEXT*, 2015, p. 16.
[42] Y. Hao, Z. Tong-Yu, and L. Chuang, "Multimedia networking, from content delivery networks to future Internet," *Chin. J. Comput.*, vol. 35, no. 6, pp. 1120–1129, 2012.
[43] C. Huang, A. Wang, J. Li, and K. W. Ross, "Measuring and evaluating large-scale CDNs," in *Proc. IMC*, 2008, pp. 1–4.
[44] K. L. Johnson, J. F. Carr, M. S. Day, and M. F. Kaashoek, "The measured performance of content distribution networks," *Comput. Commun.*, vol. 24, no. 2, pp. 202–206, 2001.
[45] R. Krishnan *et al.*, "Moving beyond end-to-end path information to optimize CDN performance," in *Proc. IMC*, 2009, pp. 190–201.
[46] C. Labovitz, S. Iekel-Johnson, D. McPherson, J. Oberheide, and F. Jahanian, "Internet inter-domain traffic," in *Proc. SIGCOMM*, 2011, pp. 75–86.
[47] J. Liang, J. Jiang, H. Duan, K. Li, T. Wan, and J. Wu, "When HTTPS meets CDN: A case of authentication in delegated service," in *Proc. IEEE S&P*, May 2014, pp. 67–82.
[48] Z. M. Mao, C. D. 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, pp. 229–242.
[49] E. Nygren, R. K. Sitaraman, and J. Sun, "The Akamai network: A platform for high-performance Internet applications," *ACM SIGOPS Oper. Syst. Rev.*, vol. 44, no. 3, pp. 2–19, 2010.
[50] J. S. Otto, M. A. Sánchez, J. P. Rula, and F. E. Bustamante, "Content delivery and the natural evolution of DNS: Remote DNS trends, performance issues and alternative solutions," in *Proc. IMC*, 2012, pp. 523–536.
[51] A.-J. Su, D. R. Choffnes, A. Kuzmanovic, and F. E. Bustamante, "Drafting behind Akamai (travelocity-based detouring)," in *Proc. SIGCOMM*, 2006, pp. 435–446.
[52] Y. Tian, R. Dey, Y. Liu, and K. W. Ross, "Topology mapping and geolocating for China's Internet," *IEEE Trans. Parallel Distrib. Syst.*, vol. 24, no. 9, pp. 1908–1917, Sep. 2013.
[53] S. Triukose, Z. Wen, and M. Rabinovich, "Measuring a commercial content delivery network," in *Proc. WWW*, 2011, pp. 467–476.
[54] H. Yin, B. Qiao, Y. Luo, C. Tian, and Y. R. Yang, "Demystifying commercial content delivery networks in China," *Concurrency Comput., Pract. Exper.*, vol. 27, no. 13, pp. 3523–3538, 2015.

**JING'AN XUE** received the B.E. degree in computer science and technology from Xi'an Jiaotong University. She is currently pursuing the Ph.D. degree with Tsinghua University. Her research interest includes network measurement and management, especially on content delivery network measurement, performance evaluation, and anomaly detection.

**DAVID CHOFFNES** is currently an Assistant Professor with the College of Computer and Information Science, Northeastern University. His research is primarily in the areas of distributed systems and networking, focusing on mobile systems, privacy, and security. He is a co-author of three textbooks, and his research has been supported by the NSF, Google, the Data Transparency Lab, M-Lab, and a Computing Innovations Fellowship.

**JILONG WANG** is currently a Professor with the Institute for Network Science and Cyberspace, Tsinghua University. He is also the Director of Network Operation Center of several large-scale network infrastructures, including Tsinghua campus network TUNET, China next generation Internet backbone CNGI-CERNET2, and Trans-Eurasia Information Network TEIN. His research interests focus on network management and measurement.

• • •