electronics

MDPI

Article

Performance Analysis of Root Anycast Nodes Based on Active Measurement

Chao Li, Yanan Cheng, Hao Men, Zhaoxin Zhang * and Ning Li *

Faculty of Computing, Harbin Institute of Technology, Harbin 150001, China; 20b903094@stu.hit.edu.cn (C.L.); mrcheng0910@gmail.com (Y.C.); 21s030201@stu.hit.edu.cn (H.M.)
* Correspondence: zhangzhaoxin@hit.edu.cn (Z.Z.); li.ning@hit.edu.cn (N.L.)

Abstract: The root server is at the top of the domain name hierarchical structure. To improve root service performance, each root deploys anycast nodes worldwide. What is the actual service performance of these nodes after deployment? We analyze the service performance of the root anycast nodes deployed in China based on the active measurement data detected by the VPs of different ISPs in different geographical locations. From the analysis, we find that the resolution performance of the roots with anycast nodes deployed in China is higher than that of roots without deployment. However, users of different operators have significant differences in accessing the root servers, such as parsing time, hitting anycast nodes, and most anycast nodes only providing services for one operator, limiting the service scope and reducing the service performance. The analysis results can help the root management and introduction institutions master the actual service status of the root servers, which can be used to optimize the performance of the existing root anycast nodes and provide a basis for deploying new root anycast nodes in the next step. Finally, we find that 67 top-level domain names are hijacked on the resolution path based on the measured data.

Keywords: root server; DNS; anycast node; active measurement

Citation: Li, C.; Cheng, Y.; Men, H.; Zhang, Z.; Li, N. Performance Analysis of Root Anycast Nodes Based on Active Measurement. *Electronics* **2022**, *11*, 1194. https://doi.org/10.3390/electronics11081194

Academic Editor: Nurul I. Sarkar

Received: 9 March 2022
Accepted: 7 April 2022
Published: 9 April 2022

1. Introduction

The domain name system (DNS) is essential infrastructure for the Internet. As the entrance and foundation of Internet communication, it is responsible for resolving and mapping the domain names of websites and mailboxes to the corresponding IP addresses. The DNS is a distributed hierarchical system in which the root name servers are at the top level and provide the authoritative referral to name servers for generic top-level domains (e.g., gTLD, e.g., .com, and .org) and country-code top-level domains (e.g., ccTLD, e.g., .cn, and .us) [1].

The DNS uses the UDP protocol to provide domain name resolution service. Due to the limitation of UDP message size, there are only 13 root name servers globally. To alleviate the pressure of communication of the root servers and improve the resolution service quality and security, a large number of root-replicating instances has been deployed by anycast technology all over the world. As defined in RFC-1546 and RFC-4786, an anycast service operates on a single IP address, but multiple anycast nodes replicate that service at different physical locations. Based on the underlying BGP routing, the users' requests are directed to the root instance node that is the "closet" to them based on "best routing" (i.e., AS path) [2].

As we all know, anycast service is applied to root name servers that benefit from anycast's load balancing and ability to mitigate DDoS attacks [3]. However, what is the actual status of the name resolving performance for the deployed root anycast nodes? For users of different operators in different geographical locations, is the actual resolving performance the same as expected?

For the above two questions, some scholars have analyzed the service performance of a specific root based on passive data obtained from the root's server side [4–6]. These

studies concentrated on the operation of a single root server only and ignored the impact of user perspective factors on the root mirroring service, such as the operator to which the user belongs or geographical location. Some scholars have also analyzed the factors affecting the performance of root servers, using active measurement data obtained from public platforms like RIPE Atlas nodes [7–9]. These studies also did not consider the user aspect. At the same time, the vantage points (VPs) of RIPE are mainly distributed in Europe, North America, and other regions, especially Asia, which is not well covered geographically, a limitation of this data. In this paper, from the user's point of view, we analyze the current status of users' actual access to root mirror nodes by collecting the top-level authoritative record data requested by users of different geographical locations and different operators to 13 roots. In the study of root anycast nodes' performance, determining the specific geographic location from which the root server providing the response data come is the first issue. Only if the geographical location is correctly identified can the latter part of the study be carried out. In previous studies, most methods used were "traceroute + chaos" path analysis [4,10]. However, when only using this method, many nodes cannot be effectively identified and located because of the blind spot problem of the traceroute path. Therefore, the "traceroute + chaos" only method cannot effectively solve the problem of identifying and locating the root anycast nodes. To solve this problem, we especially proposed a method for identifying and locating mirror images based on multi-source information, which can identify and locate more mirror nodes.

In this paper, we first obtain the top-level domain name resolution data from the root anycast nodes by using measurements of the VP points from different operators in different geographic locations and then identify and locate the nodes that provide the response data. Finally, we analyze the actual service performance of the root servers. Our exploration reveals the actual performance status of the root anycast nodes, which is beneficial for root management and institutional introduction to improve the quality of root name resolving services. Concretely, the main contributions can be summarized as follows:

- Based on the active measurement data, we conduct a detailed analysis of the resolution latency of the anycast nodes, the actual situation of the anycast nodes hit by users of different operators in different geographic locations, and the factors affecting the service quality of the root servers. Those analyses are helpful for root management and introduction institutions to understand the actual service status of the anycast nodes and provide a specific basis for adjusting and improving the service quality of anycast nodes.
- An algorithm for locating the geographical position of the root anycast node based on multi-source information is proposed. Through this method, the anycast nodes can be located effectively, which provides data support for studying of the service scope and quality of the root anycast nodes, as well as the selection of sites for additional nodes.
- The active measurement data in this paper can also be used to discover top-level domain anomalies, such as hijacking events, to a certain extent, improving the security of the root resolution data.

The rest of this paper is organized as follows. Section 2 introduces the related work on recognition and performance studies of root servers. Section 3 introduces the method of data measurement and analysis. Section 4 presents the experimental results and discussion. Finally, Section 5 introduces the conclusions.

2. Related Work

IP anycast technology [11–15] is widely used in many applications. The two most important applications are root DNS servers [3,8,11] and content delivery networks (CDNs) [16–19]. Identification and geolocation of root anycast nodes. Fan Xun et al. identified and characterized the anycast nodes of root F by combining the chaos TXT record and traceroute [10]. However, due to the existence of blind spots in the last few hops of traceroute, many anycast nodes may not be identified, and the identification efficiency is often affected. Cicalese et al. studied the enumeration and city-level geolocation of the anycast prefixes by

using latency measurements based on the detection of speed-of-light violations [20]. Sarat et al. suggested that each anycast site has an announcement radius and found that 37–80% of the queries are directed to the closest anycast instance [21]. Zhang F et al. determined whether the root mirror node accessed was located in China or abroad by leveraging the DNS censorship mechanisms of China [22]. Although these methods of measuring images based on distance can locate some images, they cannot determine the identification of specific images.

Performance of root anycast nodes. Several studies have explored the performance of root anycast from the perspective of RTTs. Schmidt analyzed the factors affecting the latency of root anycast instances by comparing the optimal performance and resolving latency in the wild, and they concluded that having "a few sites" is enough to achieve nearly as good performance as having many sites [8]. J. Liang measured the latencies between about 20,000 open recursive resolvers and root DNS servers and found that the recursive servers with parsing times over 50 ms accounted for 40% [23]. Bellis carried out a comprehensive latency assessment in F Root's anycast using RIPE Atlas [9]. In addition, some studies have evaluated the effects of anycast through examining DNS traffic and BGP data. Zhihao Li found that the root anycast nodes do not serve the purpose of load balancing well by analyzing the D root long-term passive request data, which is mainly caused by the insufficient BGP routing cooperation, and gave a solution to improve the service quality by adding the geographic location hint information of anycast nodes in BGP announcements [4,5]. By analyzing global BGP routing information passive data, Rui Bian et al. pointed out that remote peering is a factor in the low performance of anycast mirror nodes [6]. Xiaodong Lee found that the unbalanced deployment of 13 root servers results in serious root zone file distribution latency and Border Gateway Protocol (BGP) convergence cost [24]. However, for those related works, little has been done to evaluate the practical effect of root anycast nodes on the side of clients in different locations served by different operators.

Our work differs from prior studies in several respects. First, we integrate a variety of different types of data to identify root anycast instances, such as string matching of the anycast instance name, domain resolving, traceroute pathing, and so on, instead of relying just on the source data. We can recognize more anycast servers by multiple source data and improve the recognition rate. Second, we obtain measurement data through vantage points deployed in different geographical locations in China. In many previous studies, active measurement was mainly through recursion servers or public measurement platforms such as RIPE Atlas [25] or PlanetLab [26]. Active datasets obtained through recursive servers are often inaccurate due to the cache. The vantage points of RIPE Atlas or PlanetLab are mainly distributed in Europe and North America and less so in other regions, especially in Asia. This bias makes the two platforms unsuitable for measuring data in those regions. Our measurement data can avoid the above two problems. Finally, our work focuses on the imbalance of the service performance of the root server to clients in different geographical locations and different operators, rather than on the delay or BGP routing of the root anycast. The summary of our work and related work is shown in Table 1.

Table 1. Overview of research on the service performance of root mirrors.

Content and Detection Method	User Perspective Analysis	Real-Time Active Detection Data	Identifying Anycast Nodes	Locating Anycast Nodes	No Testbed Data Bias	Related Work
"Traceroute + Chaos" path analysis	×	√	√	√	×	[4,10]
Study performance of root based on RTTs	×	√	×	√	√	[8,9,23]
Based on BGP data	√	×	√	×	×	[5,6,22]
Based on geographical path	×	√	×	√	√	[20,21]
Based on active distributed detection	√	√	√	√	√	Our work

3. Methodology

In this section, we describe the data measurement and analysis method. The main idea is to obtain the response data of 13 root servers to requests from clients of different operators in different geographical locations by active measurement, identify and locate the specific root anycast nodes, and then analyze the service performance of those nodes.

3.1. Vantage Points

In this paper, we obtain research data through active measurement and use our own measurement points instead of using public measurement platforms such as RIPE Atlas or PlanetLab, which have measurement points mostly in Europe and North America and less so in Asia, especially China. We deployed 66 vantage points in 22 provinces of China, as shown in Figure 1, each with three different ISPs, including Mobile, Unicom, and Telecom, to analyze how users of varying ISPs in different geographic locations actually access the root anycast nodes.

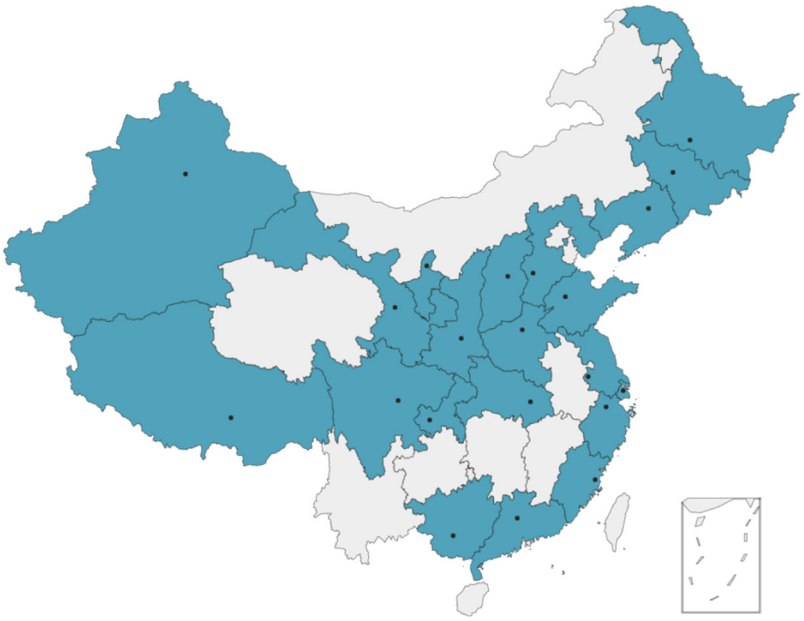

Figure 1. The geographical distribution of vantage points.

3.2. Measurement Method

The root is located at the top level of the DNS architecture and is responsible for resolving the top-level domain name through NS records as referrals to the authority servers. According to statistics, there are currently 1498 top-level domains [27]. In the measurement process, we used the vantage points described in Section 3.1 to request all top-level domains' NS records from 13 root servers, as shown in Figure 2. Although there are only 13 roots worldwide, each root actually corresponds to multiple mirror nodes because of the use of anycast technology, so there are two questions that usually need to be addressed for analyzing the performance of root anycast nodes:

1. Which root anycast node does each response datum come from?
2. How do we determine the geographic location of each root anycast node?

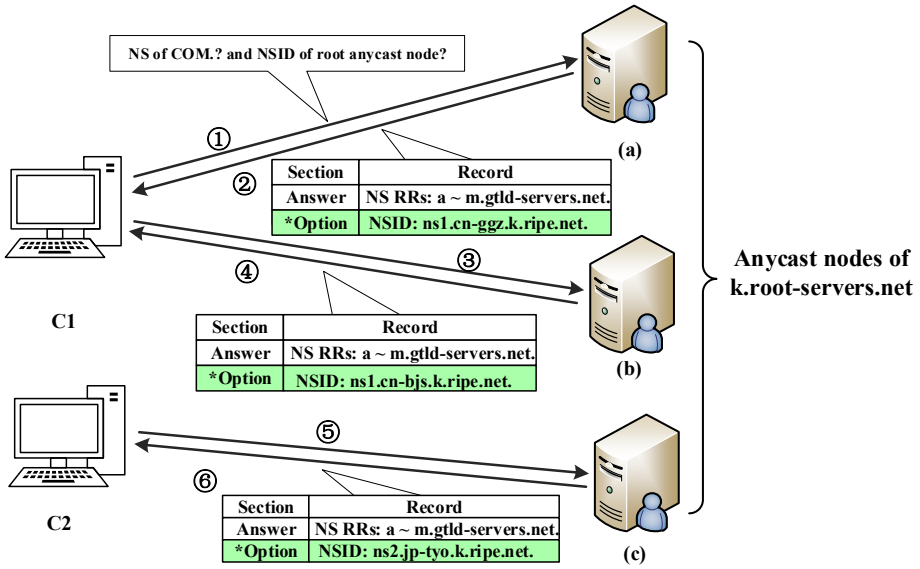

Figure 2. Example of obtaining the NS record of the top-level domain "com" and NSID of anycast nodes. The part marked with an asterisk is the NSID option, which identifies the mirror node from which this response comes.

For the first problem, the common method of obtaining anycast node names is to use DNS CHAOS TXT queries, but this method only obtains different anycast mirror names, which cannot be combined with the actual DNS request to determine which anycast node the actual DNS response comes from. Our solution is to enable the NSID option in the Extended DNS mechanism (EDNS) in the message of requesting the top-level domain NS record to the root servers. By setting this option, we can receive the NSID in the response message, which can uniquely identify the root anycast node, to distinguish which anycast node each response message comes from. To illustrate our proposed method in more detail, the following example is given, as shown in Figure 2.

The detailed process of the example shown in Figure 2 is as follows:

① Probe C1 queries the NS record of "com" to the K root and sets the NSID option in the query message. The request message is forwarded to the anycast node of the K-root via the BGP route, labeled as (a).

② The node (a) returns 13 authoritative servers' NS records of "com" in the answer section (e.g., a.gtld-servers.net, b.gtld-servers.net, . . . , m.gtld-servers.net), and returns the NSID (ns1.cn-ggz.k.ripe.net) in the Option section.

③ Probe C1 queries the K-root server for the NS record of "com" and the NSID option again, as in ①, but the request message is forwarded to the anycast node (b) of the K-root. Compared with ①, different mirror nodes are selected. This phenomenon may occur because the node server (a) has stopped working or because of the effect of a change in the BGP announcement routing policy, which will be described later in Section 4.2.

④ The node (b) returns the same NS RRs as ② and returns a different NSID: "ns1.cn-bjs.k.ripe.net".

⑤ Probe C2 queries the K-root server for the NS record of "com" and the NSID option, as in ①. The request message is forwarded by the BGP route to the anycast node, labeled as (c).

⑥ The node (c) returns the same NS RRs as ② and returns a different NSID: "ns2.jp-tyo.k.ripe.net".

Two requests from the same clients to the same root server or requests from different clients to the same root may be responded to by different root anycast nodes, which can be distinguished by the NSID in the Option section, as shown in Figure 2.

When requesting the NS records and NSID option, we simultaneously obtain the traceroute paths between the client and the root anycast node in parallel, as shown in

Figure 3, to assist in analyzing the second problem: how to determine the geographic location of each root mirror node. The solution to this problem is specifically described in Section 3.3.

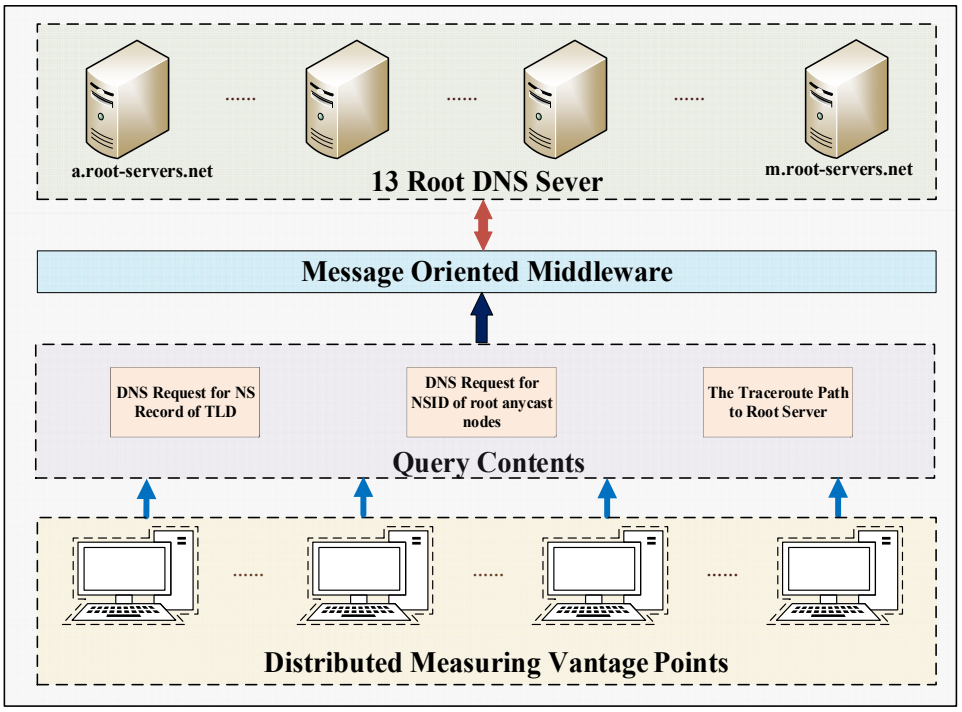

Figure 3. Schematic diagram of the active measurement structure.

In Figure 3, the distributed measuring vantage points were introduced in Section 3.1. When obtaining the query data, we use the dnspython module based on the Python language, which makes it easy to construct DNS query messages and parse response data, and for the message-oriented middleware, we use RabbitMQ, which can better realize the distribution of query messages and the recycling of response content.

3.3. The Geolocation of the Root Anycast Nodes

Through the NSID option in the EDNS, we can determine which root anycast node returns the response, and then we need to determine the specific geographical location of that node. At present, there are 1379 root anycast nodes deployed globally [28], including 22 nodes deployed domestically. The difference is that the local anycast nodes restrict the service to a specific AS through BGP routing restriction policy, while the global type has no such restriction, and any AS route can be reached. The number distribution of different kinds of root anycast nodes is shown in Figure 4.

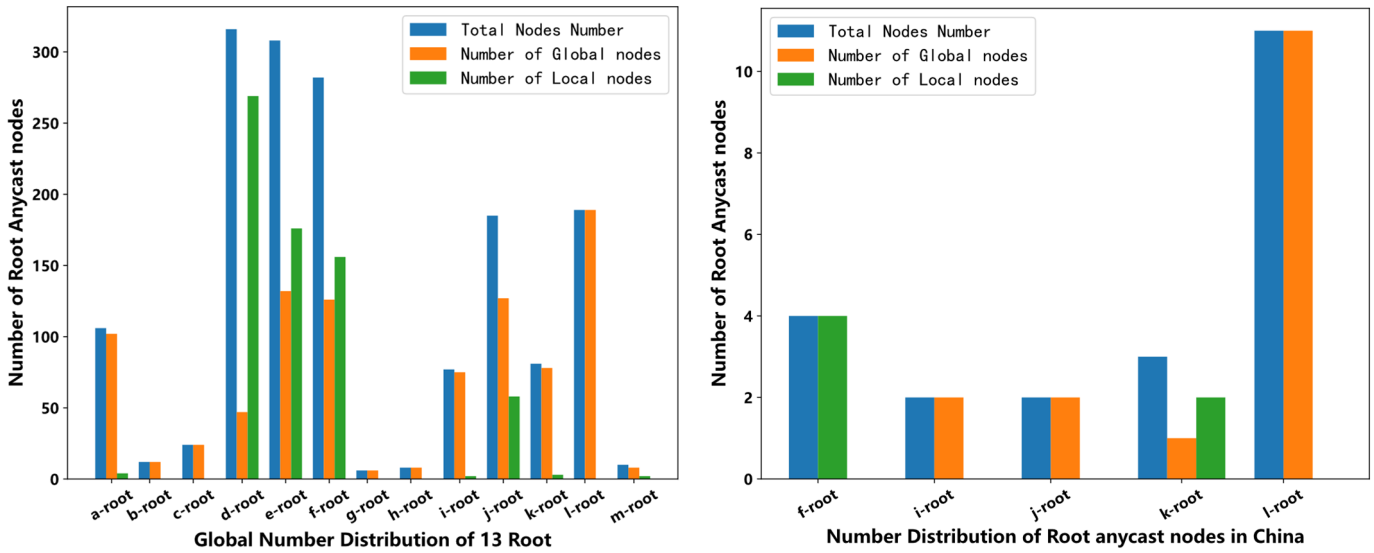

Figure 4. Global number distribution of root anycast nodes (left) and the distribution in China (right). (**a**) Global number distribution of 13 root. (**b**) Number distribution of root anycast nodes in China.

How to determine the geographical locations of these nodes is a challenging thing. Many prior studies are mostly implemented based on the method of traceroute paths. However, this method is not only limited by the number of detection points and geographical location. Due to the existence of blind spots in the last few hops of traceroute, many anycast nodes may not be located effectively. In this paper, to locate more anycast nodes, we not only use traceroute paths but also supplement a variety of other positioning methods, including string recognition, domain name resolution plus IP positioning, and latitude and longitude positioning of public datasets. Algorithm 1 describes the specific process of location.

Algorithm 1. The geographic location of root anycast nodes

Input: Three sets: set_{root_lg}, set_{id_ip}, set_{nsid};
Output: A set of geographic location of the anycast nodes: $set_{location}$;
While *nsid in* set_{nsid} **do**
  if *nsid in* set_{root_lg} **then**
    get the location L_{nsid} by its longitude and latitude, $L_{nsid} \rightarrow set_1$;
  else if *nsid* in set_{id_ip} **then**
    select the ip with the highest probability, and get the location L_{nsid} by that ip, $L_{nsid} \rightarrow set_2$;
  end
  Get the location of *nsid* by domain name resolution, $L_{nsid} \rightarrow set_3$;
  Get the location of *nsid* by string matching, $L_{nsid} \rightarrow set_4$;
end
$Set_{location} = set_1 \cup set_2 \cup set_3 \cup set_4$

First, we obtain the geographic latitude and longitude data corresponding to the root anycast nodes provided by VeriSign (the administrative organization of the "A" and "J" root servers) [28] through the web crawler and construct the mapping set set_{root_lg} between the root anycast node identifier NSID and the geographic coordinates. Using this set, we can obtain the geographical location of some anycast nodes through inverse geocoding technology. However, unfortunately, this set only provides a small number of anycast nodes' geographic coordinates, so most other nodes' geographic locations need to depend on the following methods.

Through the probe points we deploy, while requesting the NS records and NSID option, we also obtain the traceroute paths between the probe point and the root anycast node in parallel, as shown in Figure 3. In the traceroute path information, we first extract the IP address of the penultimate hop as the IP address of the anycast nodes identified by NSID. If the penultimate hop is empty, we take the IP address of the antepenultimate hop as the IP address of anycast nodes, and when both hop data are empty, they are regarded as invalid data, and we perform the next cycle measurement detection. After several consecutive detections, the probability set of NSID hitting different IPs is counted and recorded as set_{id_ip}. The IP with the highest probability in the set is taken as the IP of the anycast node. Finally, the geographical location of the anycast node is determined by an IP address location technique.

Although the setting form of the NSID of the anycast node has not been clearly specified, some operators of the root server are used to set a domain name as the NSID, and the IP address of that domain name is also the IP address of the anycast node. Therefore, the geographical locations of these nodes can be determined by domain name resolution and IP positioning.

Except for the above case, the naming rules of the root server identifier NSID are regular and generally contain the abbreviations of the country and city where the node is located. For example, in "cn-bjd-aa", "cn" is the abbreviation for China, and "bjd" is the three-character code abbreviation of Beijing, indicating that the node is located in Beijing, China. Therefore, the geographical locations of some root anycast nodes can be effectively identified by matching the NSID with the country-city abbreviation dictionary.

In general, by making extensive use of the above four identification methods, we first use each method to identify and locate the NSID of the root anycast nodes. Then, we take the union set, which can effectively locate more anycast nodes. Through the above analysis, our algorithm needs to cyclically locate all root mirrors' identifier NSID, so the time complexity of the algorithm is O(n). The variables in the algorithm use lists and sets, since the space complexity of the lists is O(n), and the memory complexity of the sets is O(1), so the memory complexity of the whole algorithm is O(n).

3.4. The Service Performance of the Root Anycast Nodes

From Sections 3.2 and 3.3, we can determine which root anycast node the user accesses and the geographical location of the anycast node. By combining the user-side information—the geographic location and operator, as well as the response message from the root anycast nodes—we analyze the actual service performance of the root anycast nodes deployed in China from the aspects of the geographic location, ISP, and the type of root anycast node in Section 4.

4. Result Analysis

In this section, we describe the experimental results and related analysis for the performance of root anycast nodes deployed in China, including the analysis of the difference between the mirrored and non-mirrored roots by counting the response times of 13 roots and the differences in the selection of root anycast nodes for different operators. Meanwhile, we also find that some top-level domain names are hijacked in the resolution path through the analysis of the measured data.

4.1. The Overall Response Times of the Root Servers

Based on 1 month of data measured by all the vantage points deployed in China, we counted the overall average response times of 13 root servers, as shown in Figure 5.

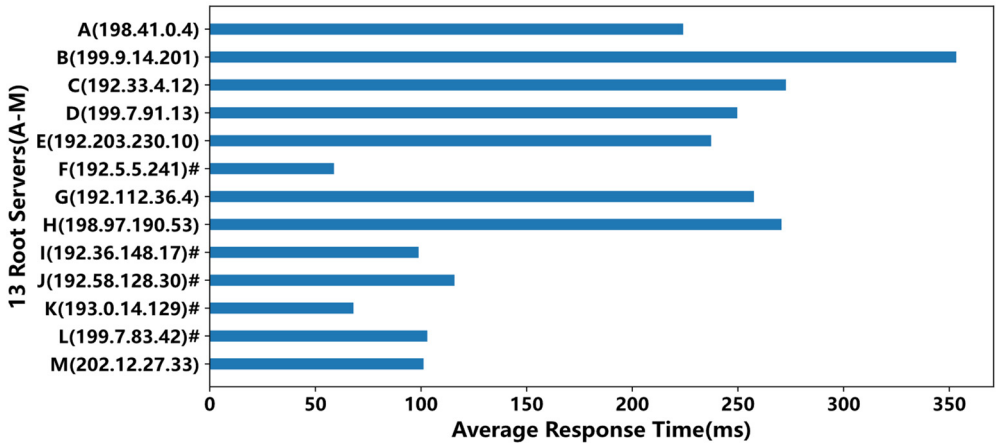

Figure 5. Overall average response times of root servers. The root marked with "#" indicates its mirror nodes have been deployed in China.

It can be seen from Figure 5 that the average response times of the roots with anycast mirror nodes deployed in China (i.e., the roots marked with "#") were significantly less than those of the roots without anycast nodes, except for the M root. The F root had the best performance with about 58 ms, and the worst was the B root with about 353 ms, which was about 6 times that of the F root. From Figure 6, we can see the standard deviation was consistent with the average response time in Figure 5. The standard deviations of the roots having anycast nodes in China were less than those of the roots without anycast nodes, except for the M root, indicating that the service performance was more stable for the roots with anycast nodes.

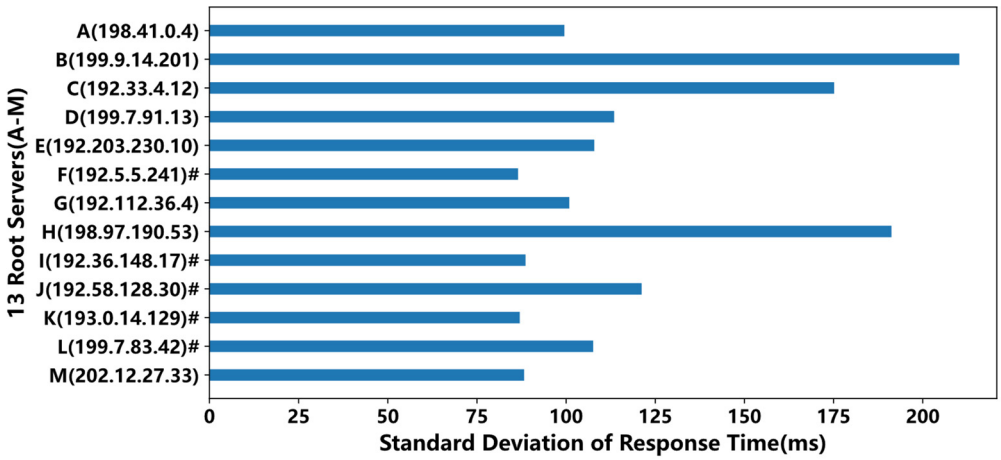

Figure 6. The standard deviations of overall average response times. The root marked with "#" indicates its mirror nodes have been deployed in China.

The response time of the M root was about 101 ms, which was close to the I root and even better than the L and J roots, but the M root had no anycast nodes in China. Through further analysis of the measured data, it was found that 99% of the responses of the M root were returned from a node in Tokyo, Japan, with an average response time of about 96 ms. In comparison, about 33% of the requests to the L root were routed to an anycast node in Australia, with an average response time of about 194 ms, so the overall average response time of the M root was better than that of the L root. Similarly, about 34.6% of the J root's responses were from a European node, with an average response time of about 216 ms, and 0.1% were from an American node at about 210 ms. Therefore, the overall average response time of the M root was also less than that of the J root.

4.2. Differences in Selecting Anycast Nodes between ISPs

The measurement results showed great differences in selecting the root anycast node for the vantage points of China Mobile, China Unicom, and China Telecom in different provinces, which led to a significant variation in time service performance.

In Figure 7, we count the average response times for three ISP vantage points in different provinces when accessing all roots (13 roots) and the average response time when only accessing the five roots (F, I, J, K, and L) with anycast nodes deployed in China. The second statistics correspond to the bar graph on the right and are indicated by black dots in the legend.

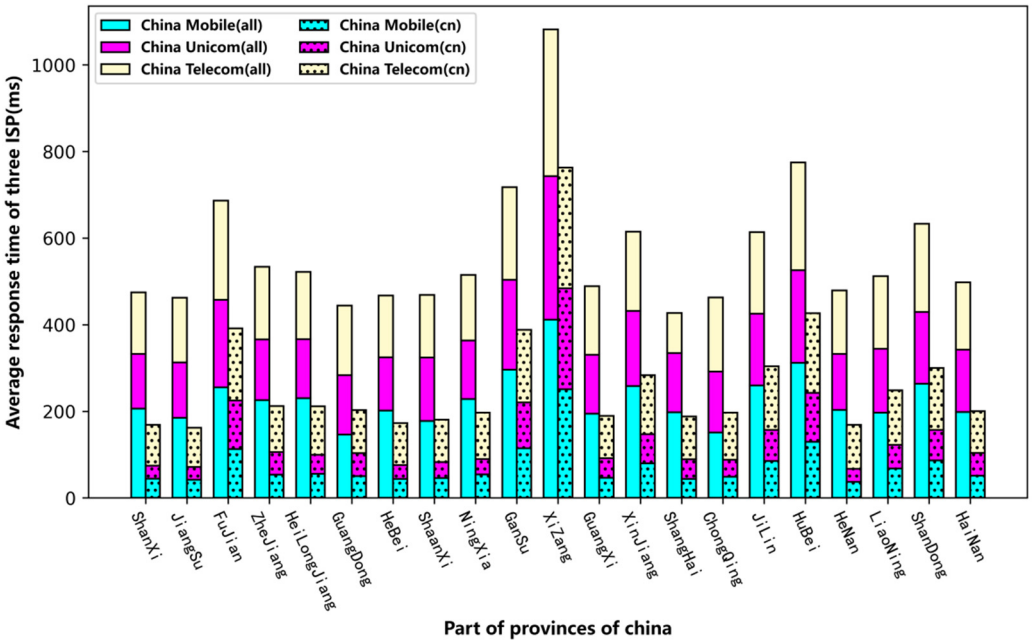

Figure 7. Average response times of accessing roots of three ISPs in different provinces.

We could characterize a province's response time from the root servers by calculating the average response times of the three ISPs. Therefore, it can be seen from Figure 7 that when accessing all the roots, Shanghai had the shortest time of about 426.7 ms, and the longest was from the Tibet Autonomous Region at about 1081.3 ms, which was 2.5 times that of Shanghai. The root response times of the provinces in the western region were significantly higher than those in the central and eastern provinces, which was caused by the long geographical distance between these provinces and the visited root anycast nodes and the poor network quality of these provinces. This has also been verified in the literature [21]. Compared with accessing all 13 roots, when only accessing the 5 roots having anycast nodes in China, the root average response times for the three ISPs in each province were significantly shorter. For example, the time of Shaanxi's China Mobile, China Unicom, and China Telecom operators accessing all 13 roots was 3.9 times, 4.0 times, and 1.5 times higher than those when accessing the 5 roots, respectively. This shows that the effect of introducing a root image in China was obvious, and the root resolution performance was improved as a whole. However, it also shows that the distribution of root anycast nodes is not reasonable enough, and new root anycast nodes need to be deployed in western China.

The root anycast nodes are divided into local and global types. Local instances attempt to limit their catchment area to their immediate peers only by announcing a supernet service with a no-export attribute. Global instances make no such restriction, allowing the BGP alone to determine their global scope and potentially provide service for the entire Internet. Among the anycast nodes of the five roots introduced in China, the F root is entirely the local type, I, J, and L are all the global type, and the K root has both local and global type nodes.

Among the anycast nodes of the five roots introduced in China, the nodes of the F root are all of the local type, the nodes of the I, J, and L roots are all of the global type, and the K root has both the local and global types. We chose the F root, L root, and K root as representatives to analyze the difference of their anycast node servers for the users of different ISPs in different provinces. For the convenience of subsequent descriptions, the global nodes belonging to China are labeled as "Global-in", those global nodes out of China are labeled as "Global-out", and the "Local" nodes refer to the local anycast nodes deployed in China.

4.2.1. F Root (Local)

The F root had only local anycast nodes in China and no global anycast nodes. We counted the response messages returned from the F root to 26,964 query requests from 3 vantage points in each province, and the number of responses from the local nodes and Global-out nodes are shown in Figure 8. In some provinces, the total responses were less than 26,964, which was because the network timeout responses were filtered out during the statistics.

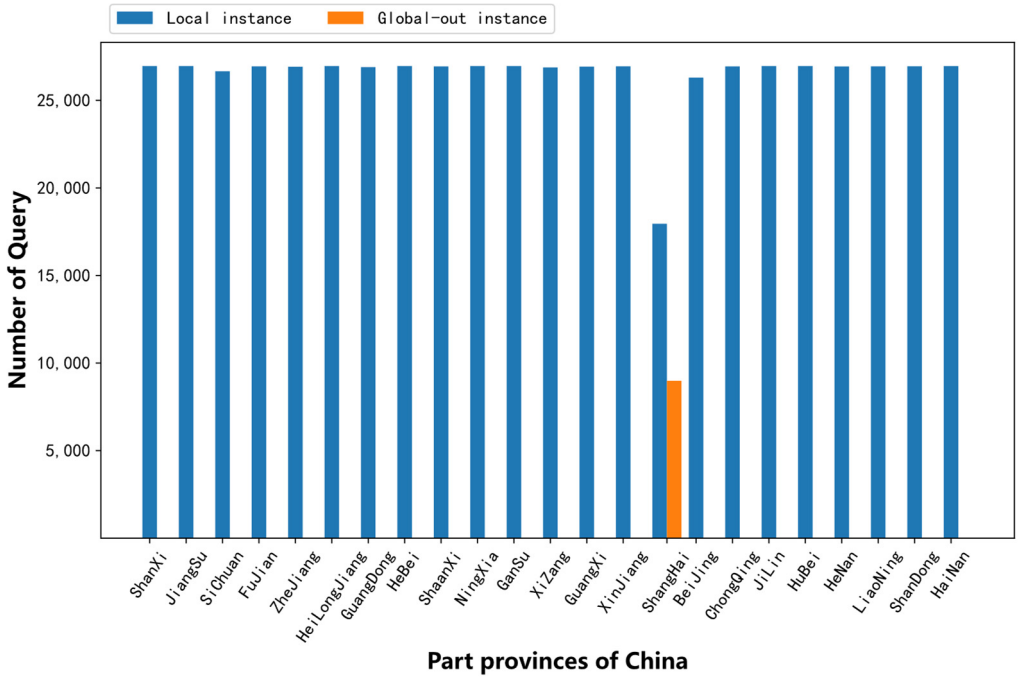

Figure 8. Query quantity to local and Global-out nodes of the F root.

It can be seen from Figure 8 that all provinces except Shanghai have access to the local anycast node in China, while Shanghai province has access to both local and global nodes. By further analyzing the query requests of the three vantage points deployed in Shanghai, we found that the VPs of China Mobile and China Unicom access the local anycast node in China, accounting for 2/3 of the requests, and the requests of Telecom access the Global-out nodes, accounting for 1/3 of the total requests. The average response time of the F root is shown in Figure 9, which is consistent with Figure 8. For Shanghai province, the response time to access the Global-out nodes was longer than that of the local node in China. From the above analysis, it can be seen that when accessing the root anycast node, some local operators rejected what was near and at hand and sought what was far away, which would increase the time for domain name resolution. That may be caused by unreasonable routing configuration from the local ISP.

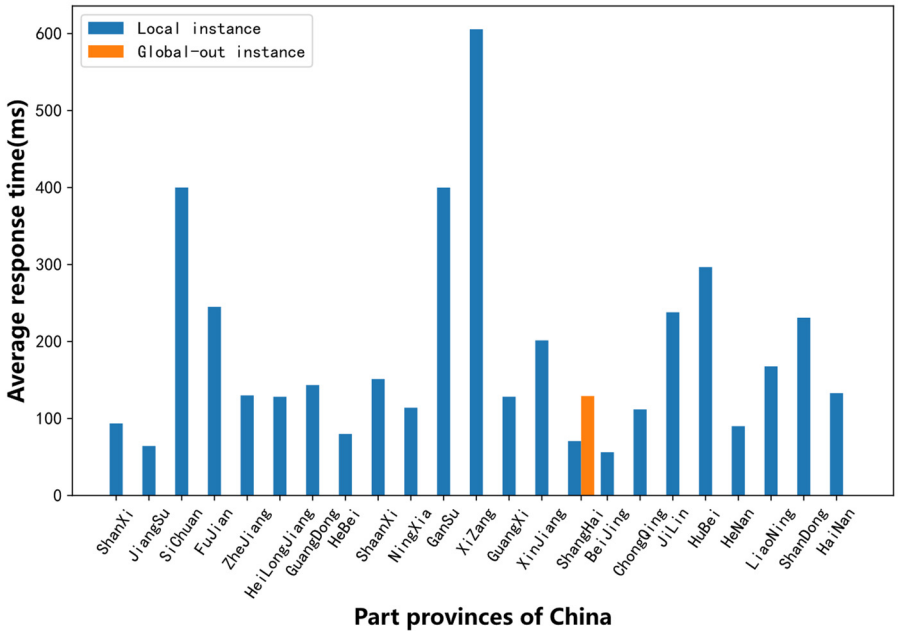

Figure 9. Average response times of two types of nodes of the F root in different provinces.

4.2.2. K Root (Local + Global)

The K root had both local and global anycast nodes in China. Figure 10 shows the number of query requests routed to the Local, Global-in, and Global-out anycast nodes of the K root in 26,964 requests from different provinces.

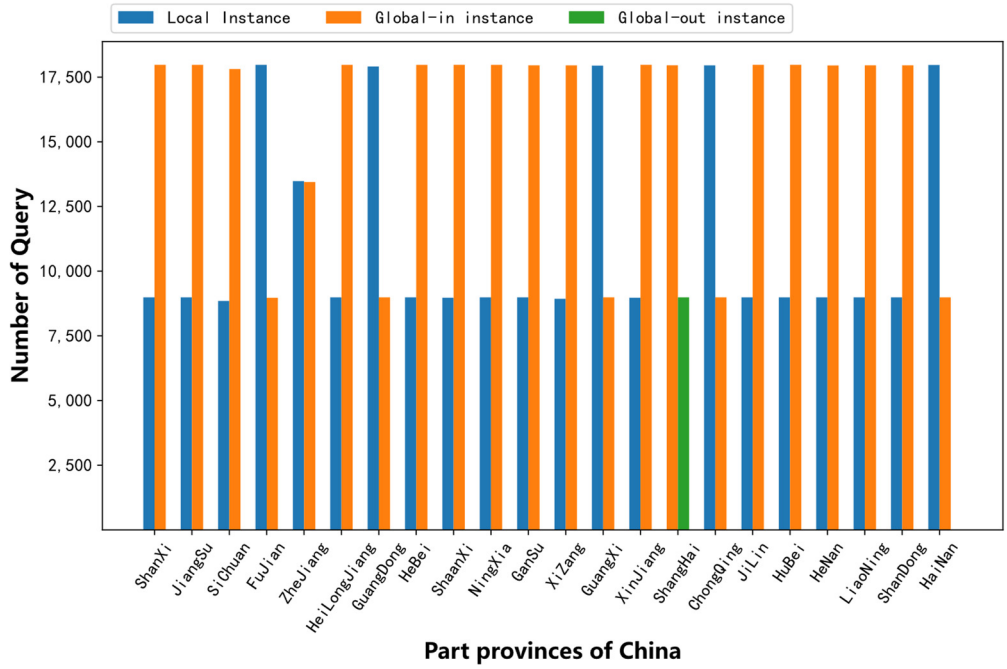

Figure 10. Query quantity to Local, Global-in, and Global-out nodes of K root.

As shown in Figure 10, only Shanghai province requested the Global-out node, and other provinces accessed the anycast nodes deployed in China. Among the 22 provinces, there were 15 provinces in which the number of query requests to the Global-in node was twice that of the Local node and 5 provinces in which the query number to the Local node was twice that of the Global-in node. Only Heilongjiang province had the same number for both types of nodes. The query requests from the 3 ISPs in 22 provinces were mostly

concentrated in one or two anycast nodes, as shown in Table 2. All the vantage points of China Unicom in 22 provinces requested the global node "ns1.cn-gya.k.rip.net" located in Guiyang, China. In 21 provinces, the vantage points of China Telecom all only accessed the local node "ns1.cn-ggz.k.ripe.net" located in Guangzhou, China. In addition, Shanghai Telecom requested three Global nodes in Tokyo, Japan. For the vantage points of China Mobile in 22 provinces, both Local and Global-in nodes were accessed.

Table 2. Statistics of K root anycast nodes requested by three operators.

Root Instance	China Mobile (22)	China Unicom (22)	China Telecom (22)	Country and City	Instance Type
ns1.cn-ggz.k.ripe.net	6	0	21	China, Guangzhou	Local
ns1.cn-gya.k.ripe.net	17	22	0	China, Guiyang	Global
ns1.jp-tyo.k.ripe.net	0	0	1	Japan, Tokyo	Global
ns2.jp-tyo.k.ripe.net	0	0	1	Japan, Tokyo	Global
ns3.jp-tyo.k.ripe.net	0	0	1	Japan, Tokyo	Global

Through the "traceroute + NSID" data analysis, we could determine that the "ns1.cn-ggz.k.rip.net" node was the Local node introduced by Telecom and the "ns1.cn-gya.k.rip.net" node was the Global image introduced by China Unicom. It can be inferred that the two operators may be trying to reduce the traffic settlement between cross operators so they do not access the root anycast node introduced by each other through methods such as routing bypass or restricting BGP routing announcements. This traffic barrier between operators limits the service range of the root anycast node and reduces the overall performance of the root resolution. It can be seen from Figure 11 that the response times of more than 70% of the provinces accessing the Local node were shorter than those of the Global-in node. It took less time for the vantage points of Shanghai province to query the Global-in nodes than the Global-out nodes. In general, it can be seen that the time performance was Local > Global-in > Global-out.

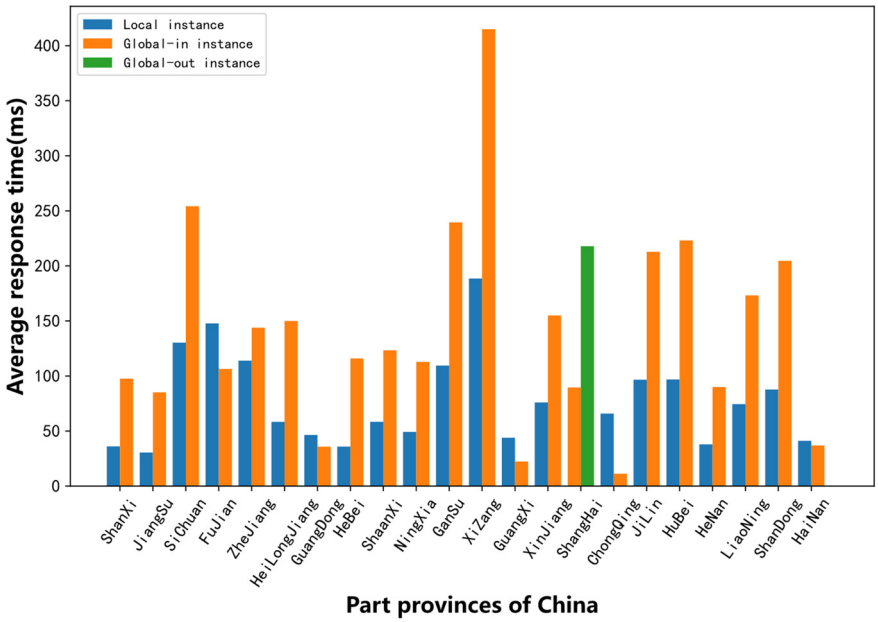

Figure 11. Average response times of three type nodes of K root in different provinces.

4.2.3. L Root (Global)

The L root only had a Global anycast node in China and no Local node. Table 3 shows which three operators queried the anycast nodes. It can be seen from the table that the vantage points of China Mobile mainly queried the two anycast instances in Zhenzhou, and the sum of the vantage points was greater than 22, which means that some vantage points accessed two or multiple different anycast nodes. China Unicom points in 21 provinces

requested the Beijing anycast node "cn-bjd-aa", and only 1 province requested the Wuhan node " cn-nhn-ab ". Through analysis, it can be found that the service scope of the L root Global nodes was also limited. Most nodes only served a specific ISP and could not be accessed by other ISPs, giving no full play to the advantages of the anycast technology, preferring the anycast node closest to the users. It can also be seen that all the detection points of China Telecom visited the overseas anycast nodes, 21 provinces requested the node "au-mas-aa" in Sydney, Australia, and 1 node visited the node "kr-icn-aa" in Incheon, South Korea.

Table 3. Statistics of L root anycast nodes requested by three operators.

Root Instance	China Mobile (22)	China Unicom (22)	China Telecom (22)	Country and City	Instance Type
au-mas-aa	0	0	21	Australia, Sydney	Global
cn-bjd-aa	0	21	0	China, Beijing	Global
cn-cgo-aa	21	0	0	China, Zhengzhou	Global
cn-cgo-ab	16	0	0	China, Zhengzhou	Global
kr-icn-aa	0	0	1	Korea, Incheon	Global
cn-nhn-ab	0	1	0	China, Wuhan	Global
cn-xnt-aa	1	0	0	China, Xining	Global
cn-xnt-ab	1	0	0	China, Xining	Global
cn-bjd-ab	0	0	0	China, Beijing	Global
cn-bjd-ac	0	0	0	China, Beijing	Global
cn-bjd-ad	0	0	0	China, Beijing	Global
cn-pvg-aa	0	0	0	China, Shanghai	Global
cn-nhn-aa	0	0	0	China, Wuhan	Global

Combined with Figure 12, it can be seen that the average response times of the Global-out nodes were significantly higher than those of the Global-in nodes. In other words, the introduction of L root anycast nodes did not improve the resolution performance for the users of China Telecom.

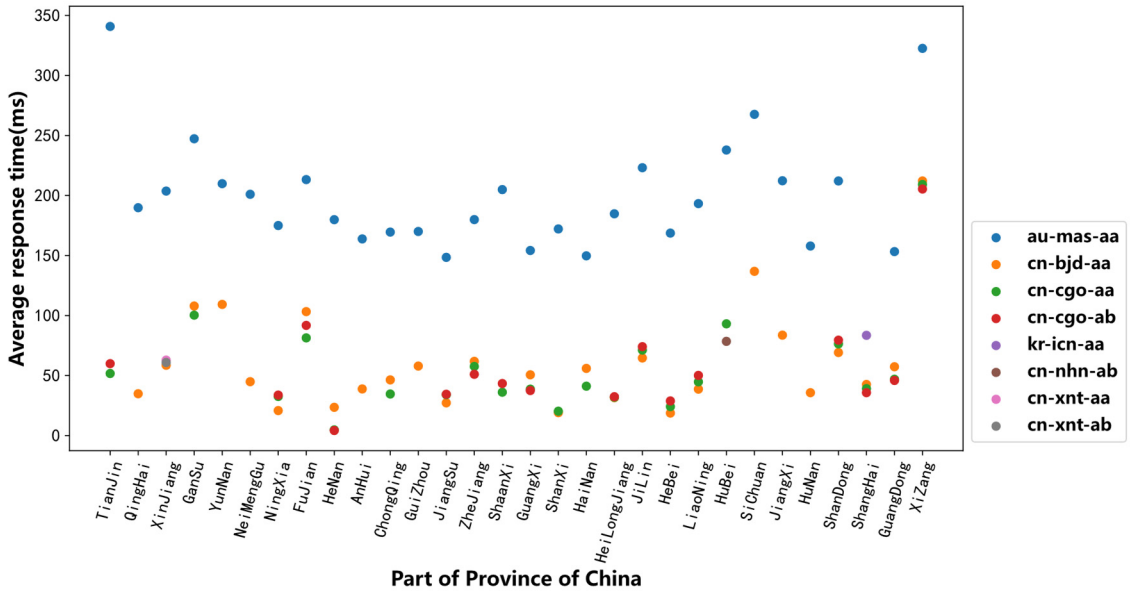

Figure 12. Average response times of three type nodes of L root in different provinces.

In addition, from the perspective of national data security, abandoning domestic instances and accessing overseas root instances will result in cross-border data traffic, which is inherently insecure. From Table 3, we can see that five anycast nodes were not accessed. Based on the analysis of Table 2, we speculate that the reason for this phenomenon may be that the autonomous system (ASes) to which the five mirror nodes belong does not

make BGP routing announcements to China Mobile, China Telecom, or China Unicom or that the three ISPs actively give up access to the five anycast nodes to reduce the traffic settlement between cross operators.

Through the analysis of the anycast nodes' performance of roots F, K, and L, we can see that the introduction of root mirror nodes in China improved the resolution performance of the root to a certain extent, and the effect of the local nodes was better than that of the global nodes. However, most anycast nodes only provide services for one operator, which not only limits the service's scope and but also reduces the service performance. In addition, when querying the L root, all telecom operators' vantage points in 22 provinces always request the overseas nodes. When accessing the F, K, and L roots, some ISPs' users will only access the overseas anycast nodes, such as Shanghai Telecom. That will increase the domain name resolution time and bring potential data security risks because of the cross-border data of the domain name resolution. To improve the root resolution performance and expand the service scope, we recommend that when deploying new root instances, the hybrid mode like the K root with both local and global anycast nodes should be adopted, which will bring the best performance.

4.3. Summary of Three Different Types of Root Analysis

Roots of different deployment mirror types have different actual access situations. The summary of the actual access of the F root, K root, and L root is shown in Table 4.

Table 4. Summary of access to the F root, K root, and L root.

Root	Deployment Type	Statistics on the Number of Vantage Points of Three ISPs						
		China Mobile		**China Unicom**		**China Telecom**		
F	Local	Local		Local		Local	Global-out	
		100%		100%		95%	5% (Shanghai Telecom)	
K	Local + Global	Local	Global-in	Local	Global-in	Local	Global-in	Global-out
		27%	77%	0	100%	95%	0	5% (Shanghai Telecom)
L	Global	Global-in		Global-in		Global-in		Global-out
		100%		100%		0		100%

It can be seen from Table 4 that for roots of different deployment types, the actual selection strategies of different operators were quite different. For example, for the K root, among the 22 vantage points of China Mobile, 27% selected the anycast nodes of the local type, and 77% selected the global type nodes. While all the vantage points of China Unicom selected the global type nodes, and the vantage points of China Telecom, except for the Shanghai node, all selected the local type, accounting for 95%. This is because the operator will preferentially choose to visit the root anycast node introduced by itself rather than other operators to reduce inter-operator settlement. The BGP routing information of the introduced root anycast node is often restricted because it is not mutually advertised among operators, limiting the service scope of the anycast nodes. In addition, we can see that when the vantage points of China Telecom accessed the F root, K root and L root, different proportions accessed the overseas root anycast nodes. For example, 100% of vantage points even accessed the overseas node when accessing the L root, which would not only increase the domain name resolution time but also bring potential data security risks because of the cross-border data of the domain name resolution.

4.4. Hijacking Abnormal Events

Based on the measured data in this article, while analyzing the service performance of the root anycast nodes, we also found that many top-level domain names were hijacked on the resolution path. We judged whether the top-level domain name was hijacked based on the following two points: (1) whether the response data contained the NSID option,

because when querying the NS records of the top-level domain name, we also requested the NSID option, so the normal response message should have returned both the NS records and the NSID of root anycast node, and (2) whether the TTL value of the response NS record was consistent with that in the root zone file record [27]. As shown in Table 5, for the top-level domain names "mh" and "xn–zfr164b", the returned response message had NS records, but the value of the NSID field was null. At the same time, the TTL values of their NS records were 30 s and 3600 s, respectively. However, we know from the root zone file that both TTL values should be 172,800 s.

Table 5. Examples of hijacking domain names.

TLD	NSID	NS Servers	NS_TTL	Position of Vantage Point	ISP
mh	null	ns.amarshallinc.com.; ns.nta.mh.	30; 30	Inner Mongolia	China Mobile
xn–zfr164b	null	ns12.sdc.org.cn.; ns13.sdc.org.cn.; ns14.sdc.org.cn.	3600; 3600; 3600	Guizhou	China Mobile

Therefore, we can speculate that the request messages of the two TLDs in Table 5 were hijacked on the resolution path. These request messages were sent to the hijacked DNS server and did not reach the root server, and the response message received by the measurement point came from the hijacked server rather than the root server.

We found that this hijacking occurred only in some specific vantage points in certain provinces. In addition to the two TLDs in Table 5, we also found 32 and 35 hijacked TLDs at the vantage point of China Mobile in Guizhou province and that of China Mobile in the Inner Mongolia Autonomous Region, respectively. At the same time, this hijacking phenomenon did not always exist, and we observed that after 24 h, the resolution data of 67 hijacked TLDs returned to normal. The specific reason for this phenomenon is still unclear, and it is speculated that it may be related to operators deliberately hijacking domain name traffic to reduce cross-provincial traffic settlement costs or conducting related network experiments.

5. Conclusions

In this paper, combined with our measurement data, we first proposed an algorithm for locating the geographical positions of the root anycast nodes based on multi-source information, which could effectively identify and locate the anycast nodes. Then, we analyzed the actual service performance of the root anycast nodes deployed in China from the aspects of geographic location, ISP, and the type of root anycast nodes, and we found that the advantage of anycast technology was not brought into full play in root mirrors' resolution efficiency and service scope. Due to the unbalanced deployment of root anycast nodes, the root response times of provinces in the western region were significantly higher than those in the central and eastern provinces. The overall response time performance of the root anycast nodes was as follows: Local > Global-in > Global-out.

At the same time, based on our active measurement data, we also found the phenomenon of hijacking top-level domain names on the root resolution path, which could bring a series of network security problems, such as fishing net websites and user privacy leaks.

To improve the resolution performance of the root servers, we give the following recommendations. First, it is necessary to deploy more root anycast nodes, especially in remote and economically underdeveloped areas. Then, we should pay attention to the factors affecting the resolution performance of the anycast node. For example, when deploying root mirror nodes, a hybrid model of local and global should be adopted, and the operating organization of anycast nodes should conduct peer-to-peer interconnection with different operators as much as possible to expand the anycast nodes' service scope.

In future work, we intend to further study and analyze the factors affecting the service scope of the root mirrors and explore other abnormal events that may occur on the root mirrors.

Author Contributions: Conceptualization, C.L., Z.Z. and N.L.; methodology, C.L. and Y.C.; software, C.L.; validation, H.M.; formal analysis, C.L. and Y.C.; investigation, H.M.; resources, Y.C.; data curation, H.M.; writing—original draft preparation, C.L.; writing—review and editing, C.L. and Y.C.; visualization, C.L.; supervision, N.L.; project administration, Z.Z.; funding acquisition, Z.Z. All authors have read and agreed to the published version of the manuscript.

Funding: This work is supported by the Natural Science Foundation of Shandong Province (Grant No. ZR2020KF009) and the Young Teacher Development Fund of Harbin Institute of Technology (Grant No. IDGA10002081).

Conflicts of Interest: The authors declare no conflict of interest.

References

1. Domain Names—Concepts and Facilities. Available online: https://datatracker.ietf.org/doc/html/rfc1034 (accessed on 1 December 2021).
2. Host Anycasting Service. Available online: https://datatracker.ietf.org/doc/html/rfc1546 (accessed on 1 December 2021).
3. Saridou, B.; Shiaeles, S.; Papadopoulos, B. DDoS Attack Mitigation through Root-DNS Server: A Case Study. In Proceedings of the 2019 IEEE World Congress on Services (SERVICES), Milan, Italy, 8–13 July 2019; pp. 60–65. [CrossRef]
4. Levin, D.; Zhi, L.; Spring, N.; Bhattacharjee, B. Longitudinal Analysis of Root Server Anycast Inefficiencies. Technical Report. 2017.
5. Li, Z.; Levin, D.; Spring, N. Internet anycast: Performance, problems, & potential. In Proceedings of the 2018 Conference of the ACM Special Interest Group on Data Communication, Budapest, Hungary, 20–25 August 2018; pp. 59–73.
6. Bian, R.; Hao, S.; Wang, H. Towards passive analysis of anycast in global routing: Unintended impact of remote peering. *ACM SIGCOMM Comput. Commun. Rev.* **2019**, *49*, 18–25. [CrossRef]
7. Cicalese, D.; Joumblatt, D.; Rossi, D.; Buob, M.; Augé, J.; Friedman, T. A Fistful of Pings: Accurate and Lightweight Anycast Enumeration and Geolocation. In Proceedings of the 2015 IEEE Conference on Computer Communications (INFOCOM), Hong Kong, China, 26 April–1 May 2015.
8. Oliveira Schmidt, R.D.; Heidemann, J.; Kuipers, J.H. Anycast latency: How many sites are enough? In *Passive and Active Network Measurement Conference (PAM)*; Springer: Berlin/Heidelberg, Germany, 2017; pp. 188–200.
9. Bellis, R. Researching F-root anycast placement using RIPE Atlas. *RIPE Labs* **2015**. Available online: https://labs.ripe.net/Members/raybel-lis/researching-froot-anycast-placement-using-ripe-atlas (accessed on 1 December 2021).
10. Dang, W.; Wang, H.; Wang, J. Evaluating performance and inefficient routing of an anycast CDN. In Proceedings of the 2019 IEEE/ACM 27th International Symposium on Quality of Service (IWQoS), Phoenix, AZ, USA, 24–25 June 2019; pp. 1–10.
11. Sommese, R.; Akiwate, G.; Jonker, M. Characterization of Anycast Adoption in the DNS Authoritative Infrastructure. In Proceedings of the Network Traffic Measurement and Analysis Conference (TMA'21), Virtual Online Conference, 14–15 September 2021.
12. Zhang, X.; Sen, T.; Zhang, Z. AnyOpt: Predicting and optimizing IP Anycast performance. In Proceedings of the 2021 ACM SIGCOMM 2021 Conference, Virtual Online Conference, 2021; pp. 447–462.
13. Li, Y.; Han, Z.; Gu, S.; Zhuang, G.; Li, F. Dyncast: Use Dynamic Anycast to Facilitate Service Semantics Embedded in IP address. In Proceedings of the 2021 IEEE 22nd International Conference on High Performance Switching and Routing (HPSR), Paris, France, 7–10 June 2021.
14. Metz, C. IP anycast point-to-(any) point communication. *IEEE Internet Comput.* **2002**, *6*, 94–98. [CrossRef]
15. Katabi, D.; Wroclawski, J. A framework for scalable global IP-anycast (GIA). *ACM SIGCOMM Comput. Commun. Rev.* **2000**, *30*, 3–15. [CrossRef]
16. Calder, M.; Flavel, A.; Katz-Bassett, E.; Mahajan, R.; Padhye, J. Analyzing the performance of an anycast CDN. In Proceedings of the ACM Internet Measurement Conference (IMC), Tokyo, Japan, 28–30 October 2015; pp. 531–537.
17. Giordano, D.; Cicalese, D.; Finamore, A.; Mellia, M.; Joumblatt, D.Z.; Rossi, D. A first characterization of anycast traffic from passive traces. In *International Workshop on Traffic Monitoring and Analysis (TMA)*; Springer: Berlin/Heidelberg, Germany, 2016.
18. Fan, X.; Katz-Bassett, E.; Heidemann, J. Assessing affinity between users and CDN sites. In *International Workshop on Traffic Monitoring and Analysis (TMA)*; Springer: Berlin/Heidelberg, Germany, 2015; pp. 95–110.
19. Fan, X.; Heidemann, J.; Govindan, R. Evaluating anycast in the domain name system. In Proceedings of the 2013 Proceedings IEEE INFOCOM, Turin, Italy, 14–19 April 2013; pp. 1681–1689. [CrossRef]
20. Cicalese, D.; Augé, J.; Joumblatt, D.; Friedman, T.; Rossi, D. Characterizing IPv4 Anycast Adoption and Deployment. In Proceedings of the ACM Conference on emerging Networking EXperiments and Technologies (CoNEXT), Heidelberg, Germany, 1–4 December 2015.
21. Sarat, S.; Pappas, V.; Terzis, A. On the use of anycast in DNS. In Proceedings of the IEEE International Conference on Computer Communications and Networks (ICCCN), Arlington, VA, USA, 9–11 October 2006; pp. 71–78.

22. Zhang, F.; Lu, C.; Liu, B. Measuring the Practical Effect of DNS Root Server Instances: A China-Wide Case Study. In *International Conference on Passive and Active Network Measurement*; Springer: Berlin/Heidelberg, Germany, 2022; pp. 247–263.
23. Liang, J.; Jiang, J.; Duan, H.; Li, K.; Wu, J. Measuring query latency of top level DNS servers. In *Passive and Active Network Measurement Conference (PAM)*; Springer: Berlin/Heidelberg, Germany, 2013; pp. 145–154.
24. Lee, X.; Yan, Z.; Chaparadza, R. Scaling the number of DNS root servers with Internet. In Proceedings of the 2014 IEEE Globecom Workshops (GC Wkshps), Austin, TX, USA, 8–12 December 2014; pp. 248–253.
25. RIPE. What is RIPE Atlas? 2021. Available online: https://atlas.ripe.net/about/ (accessed on 1 December 2021).
26. PlanetLab. What is PlanetLab? 2021. Available online: https://planetlab.cs.princeton.edu/about.html (accessed on 1 December 2021).
27. Root Zone. 2021. Available online: https://www.iana.org/domains/root/db (accessed on 1 December 2021).
28. DNS Root Server Datasets. 2021. Available online: https://root-servers.org/archives/ (accessed on 1 December 2021).