

Latency Evaluation of Anycast CDN with Limited Cache Sets

Futo Masai* and Noriaki Kamiyama†

College of Information Science and Engineering, Ritsumeikan University, Osaka, 567-8570 Japan

*Email: is0635es@ed.ritsumei.ac.jp, †kamiaki@fc.ritsumei.ac.jp

Abstract—In CDN (Content Delivery Network), cache servers (CSs) are geographically distributed to reduce the distance between users and servers, reducing delivery delay time. In the Anycast CDN, CS selection depends on the routing policy of BGP (Border Gateway Protocol), which leads to suboptimal CS selection. In addition, the probability of selecting an optimal CS tends to decrease as the number of CSs constituting a CS set of anycast is increased. So, we explored the use of genetic algorithm to construct multiple CS sets, each consisting of a small number of CSs. However, the performance of this method in the real Internet has not yet been evaluated. In this paper, we measure and evaluate the effectiveness of this method using the measured RTT (Round Trip Time). We propose an evaluation method and use the measured RTT to compare our method with existing methods with a single CS set consisting of many CSs and clarify the effectiveness of our method.

I. INTRODUCTION

Content Delivery Networks (CDNs) enhance content delivery by geographically distributing cache servers (CSs) and delivering content from CSs close to users, reducing latency and distributing load. Anycast CDNs assign the same IP address to multiple CSs and route user requests via BGP to the closest CS. However, BGP routing policies often lead to suboptimal CS selection. Calder et al. point out that while anycast CDNs simplify operations compared to DNS-based CDNs by relying on BGP for CS selection, they lack precise control over CS selection [3]. Consequently, 10% to 15% of users route to distant CSs, contributing significantly to anycast performance degradation. Even when the closest CS is not selected, anycast delivery often routes requests to relatively nearby servers, which are presumed to provide good performance [3]. Therefore, to mitigate the impact of suboptimal routing by BGP, large-scale anycast operators deploy hundreds of CSs [5]. Schmidt et al. demonstrate that DNS root servers with more distribution sites excel in minimizing latency increases when routing to optimal locations fails[4]. However, these servers also exhibit a decreased probability of routing to optimal locations compared to those with fewer sites.

Microsoft’s anycast CDN maintains multiple sets of CSs, called rings, each of which is assigned a unique anycast address, to respond to different legal regulations in different countries [2]. User requests are always directed to the ring that meets the application’s regulatory requirements and contains the most CSs. However, in this scheme, CS sets are selected without considering the bias in the popularity of the content, which may reduce the cache efficiency. To address these limitations, we proposed a genetic algorithm-based method to construct rings with fewer CSs, optimizing for region-specific

content demand [1]. However, they only evaluated demand coverage, not latency reduction.

Therefore, this paper evaluates the latency reduction of multi-ring anycast CDNs, built with genetic algorithms, using measured RTTs. We measure RTTs for optimal and suboptimal CS selections, weighted by their occurrence probability. We then calculate delivery latency per content and average latency. Finally, we compare this method with a conventional single-ring anycast CDN to verify the genetic algorithm’s effectiveness.

II. ANYCAST CDN WITH LIMITED CACHE SETS

Anycast delivery faces the challenge of CS selection accuracy decreasing with an increasing number of CSs. In [1], we proposed a method to construct rings using a small number of CSs, considering the regionality of content popularity, using a genetic algorithm (GA). GA, an optimization algorithm inspired by biological evolution, is employed to efficiently construct multiple rings by combining geographically distributed ASs. Each AS contains CSs, and the collection of rings composed of multiple CSs as the genes. When a CD exists within the content demand country, it covers that country’s demand. Fitness is defined as the average coverage rate by the AS placement of genes among the top 5 demanding countries for each content. By constructing rings, i.e., CS sets, for content with demand ratios exceeding 90% in the top 5 demanding countries, efficient content delivery can be realized for most users.

III. METHODOLOGY FOR MEASURING AND ANALYZING DELIVERY DELAY

A. overview

This paper evaluates the performance of an anycast CDN with multiple rings constructed using GA through RTT measurements. However, we cannot directly measure RTT because the anycast CDN with GA-based limited CS sets is not deployed on the internet. Therefore, we propose a method for performance evaluation using RTT for an anycast CDN that is not deployed in a real-world setting.

To account for suboptimal CS selection in anycast delivery, we determine the optimal delivery AS and the AS selected in suboptimal cases for each content demand country and measure RTT for each. By weighting the measured RTTs with their respective probabilities of occurrence and summing them, we calculate the RTT that accounts for suboptimal server selection in anycast delivery.

B. RTT Measurement Methodology

The RTT was measured using the RIPE Atlas, an internet measurement platform [6]. The measurement points, referred to as *probes*, are distributed worldwide and considered by users requesting content. Ping measurements were conducted towards the AS to measure the RTT to the AS delivering the content. The IP addresses of the AS were selected from its prefix range, where three randomly chosen responsive addresses were used for the ping measurements. The average RTT of these three measurements was taken as the RTT to the AS.

The latency of each content delivery was calculated using the RTT values. Let $r_{m,i}$ denote the i -th country with demand for content m , $O_{m,i}$ represents the RTT from $r_{m,i}$ to the optimal AS, and $S_{m,i}$ is defined as the RTT to a suboptimal AS. Additionally, let $T_{m,r_{m,i}}$ be the demand proportion of $r_{m,i}$ for content m , and p be the probability of selecting a suboptimal AS. The latency L_m for content m is then defined as follows,

$$L_m = \frac{\sum_{i=1}^5 \left((1-p)O_{m,i} + pS_{m,i} \right) T_{m,r_{m,i}}}{\sum_{j=1}^5 T_{m,r_{m,j}}}. \quad (1)$$

C. Determination of Delivery AS

Determining the optimal delivery and suboptimal AS for each content demand country is necessary. In this study, to reduce the number of measurements, it is assumed that suboptimal selection occurs only when there are multiple ASs of the content demand country within the ring to which the content is assigned. Otherwise, $O_{m,i} = S_{m,i}$. The delivery AS is determined according to Algorithm 1.

Algorithm 1 AS Selection Algorithm

- 1: Search for an AS in the ring that belongs to the same country as the demand country.
- 2: **if** an AS in the same country exists **then**
- 3: **if** only one AS exists **then**
- 4: Selects that AS as the optimal and suboptimal AS
- 5: **else**
- 6: Measures the RTT for multiple ASes.
- 7: Selects the AS with the minimum RTT as the optimal AS and the AS with the maximum RTT as the suboptimal AS
- 8: **end if**
- 9: **else**
- 10: Calculates the distance to all ASes in the ring
- 11: Selects the AS with the minimum distance
- 12: **end if**

IV. PERFORMANCE EVALUATION

A. Evaluation Conditions

This paper compares the anycast CDN comprising multiple rings constructed from a small number of CSs using GA with an existing method consisting of a single ring composed of many CSs. For content, we obtained data from 261 popular webpages with demand ratios of 90% or more in the top 5 demanding countries from the June to August 2024 Similar-Web popular website rankings [8]. For AS data, we collected

the top 300 ASs from the CAIDA dataset and additionally 50 ASs to ensure that all ASs from the top 5 demanding countries were included [7]. These were used to construct rings using GA. The existing method is configured by combining the ring sets constructed by GA into a single ring. Table I lists the number of rings (N) and the number of constituent ASs (k) for the anycast CDN constructed in both the GA-based method and the corresponding existing method.

TABLE I
SETTING VALUE SETS OF N AND k

	s1		s2		s3	
	N	k	N	k	N	k
GA Method	20	5	20	10	20	15
Existing Method	1	84	1	154	1	198

B. Performance Evaluation Considering Suboptimal Selection

Figure 1 shows the average content latency against the probability p of suboptimal AS selection varies. The Anycast CDN constructed using GA demonstrates stability, while the existing method deteriorates with increasing p . In s1, the proposed method outperforms the existing method at $p = 0.35$, s2 at $p = 0.31$, and s3 at $p = 0.25$. Given that 55% of users are routed to geographically optimal CSs [3], the anycast CDN comprising multiple rings constructed from a small number of CSs using GA is effective.




Fig. 1. Average access latency of webpages against p , ratio of nonoptimally selecting AS

ACKNOWLEDGEMENT

This work was supported by a grant from the Takahashi Industrial and Economic Research Foundation.

REFERENCES

- [1] Chihito Kato and Noriaki Kamiyama, Designing Server Sets for Anycast CDN Using Genetic Algorithm, IEEE ICC 2024
- [2] Homas Koch, Ethan Katz-Bassett, John Heidemann, Matt Calder, Calvin Ardi, and Ke Li., Anycast in Context: A Tale of Two Systems, ACM SIGCOMM 2021
- [3] Matt Calder, Ashley Flavel, Ethan Katz-Bassett, Ratul Mahajan, and Jitendra Padhye, Analyzing the Performance of an Anycast CDN, ACM IMC 2015
- [4] R. d. O. Schmidt, J. Heidemann, and J. H. Kuipers., Anycast latency: How many sites are enough?, PAM 2017
- [5] Jan Marius Evang, Alojz Gomola and Tarik Cicic, Anycast Metrics and Performance Tuning, SoftCOM 2024
- [6] RIPE NCC web page, <https://www.ripe.net>
- [7] CAIDA web page, <http://www.caida.org/data>
- [8] Similar web web page, <https://www.similarweb.com>