階層型CDNに対するリアルタイム最適キャッシュ汚染攻撃

Optimum Cache Pollution Attacks on Real-Time Hierarchical CDN

Jiaqi Liu　Noriaki Kamiyama

Ritsumeikan Univ

2026.03.04

Content Delivery Network

- Content Delivery Network (CDN)
  - Origin servers: Provide the original version of the content
  - Cache servers: Cache the copy of contents, and they are responsible for delivering that content to nearby users.
  - DNS servers: Respond user's request with the name of a cache server from which the content can be served faster.

- The feature of CDN
  - Serves a large portion of the Internet content
  - Provides a faster and high-performance experience
  - Reduce bandwidth costs

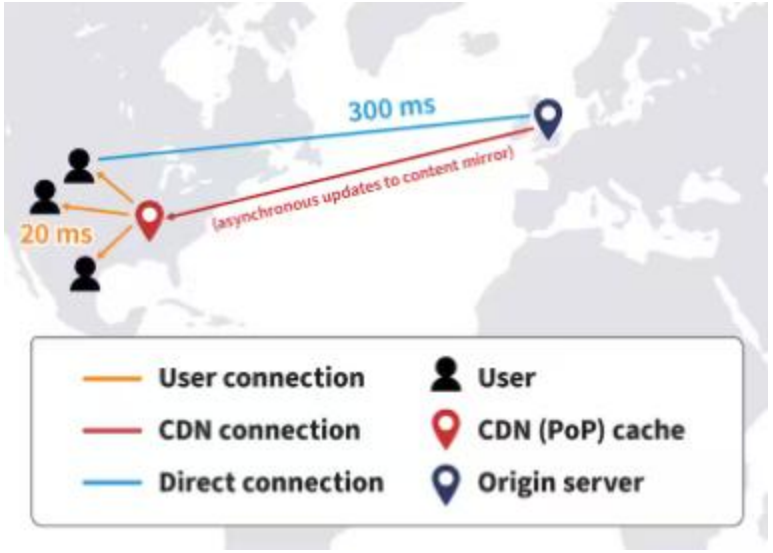

Potential Threat

- DDoS Attack
  - Sending a large number of requests makes it impossible for the server to process them
  - Edge servers can prevent them from attacking the origin server
  - Simple and efficient, but easy to be detected

- Cache pollution attack（CPA）
  - Sending a small number of requests for low-popularity content to occupy cache space
  - Requests sent by normal users will also be carried to the origin server
  - Complex and efficient, but difficult to be detected

Purpose of research

- Propose an algorithm for optimizing cache pollution attacks
- Identify the vulnerabilities of the CDN

Model structure

- Hierarchical cache CDN
  - Users connect to the local edge server
  - Requests that cache missed on the edge server are sent to the local Mid-server
  - Requests that cache missed on the Mid-server are sent to the origin server

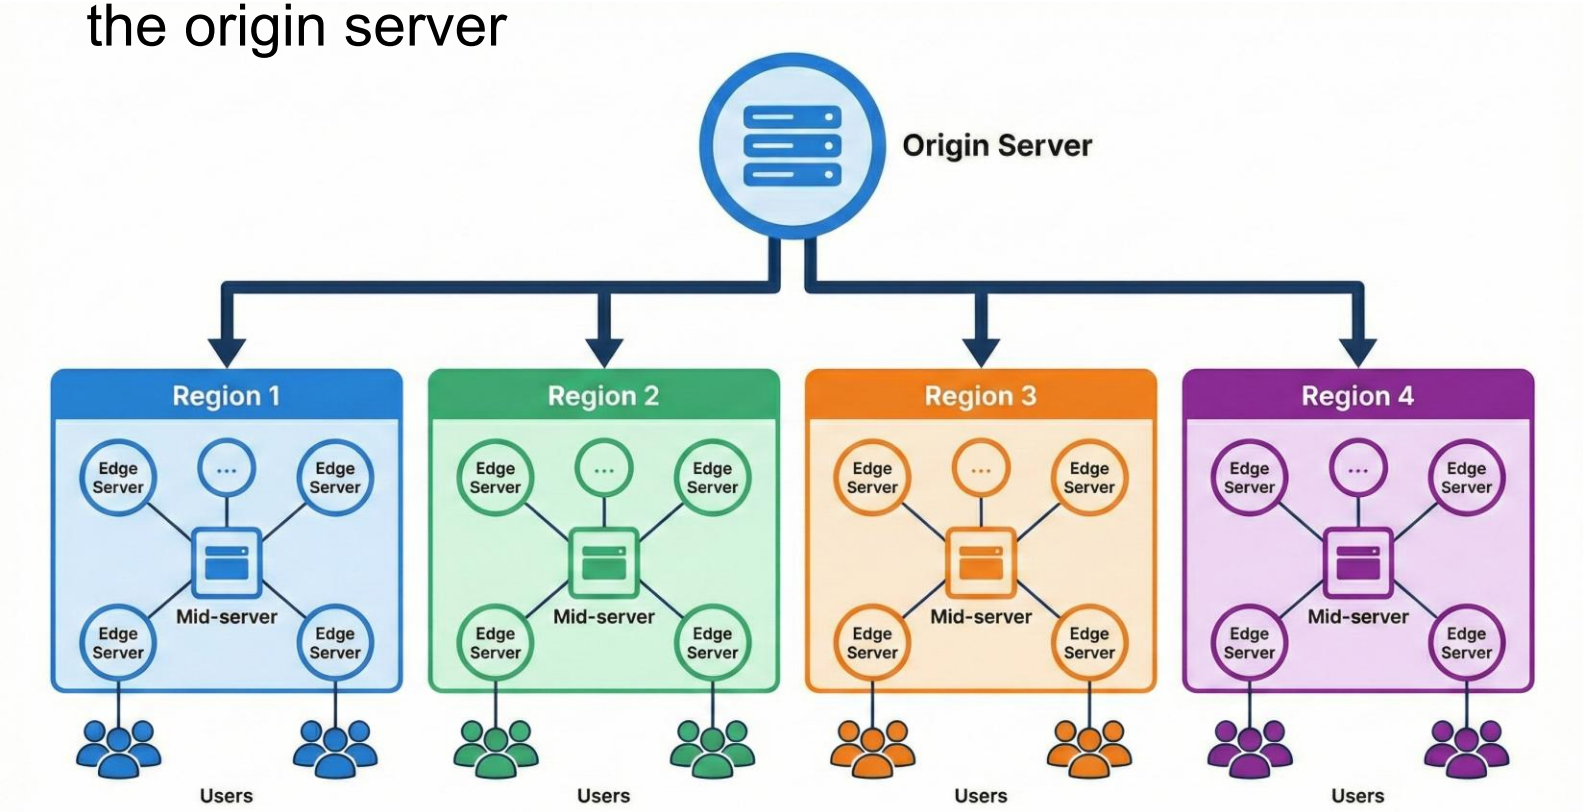

Adaptive Attack Strategy

- The "Perturb-and-Observe" Mechanism
  - Perturbation: Superimpose Gaussian noise on the current strategy $\tilde{\theta} = \theta + \epsilon, \quad \epsilon \sim \mathcal{N}(0, \sigma^2 \mathbf{I})$
  - Action Generation: $\mathbf{A}_t = C_{max} \cdot \mathrm{Softmax}(\tilde{\theta})$
  - Evaluation: $R_t = W_t \times (1 + D_t)$
  - Greedy Update: **if** $R_t > R_{best}$ **then**
    $\theta \leftarrow \theta + \eta \cdot \epsilon$
    $R_{best} \leftarrow R_t$

Adaptive Attack Strategy

- Black-Box Environment
  - Challenge: The attacker is located at the network edge and has no access to the cache table or topology structure within the CDN
  - Solution：Abandoning the traditional gradient-based reinforcement learning, we adopt zeroth-order optimization - a perturbation-based stochastic hill climbing algorithm
  - Objective：Under a fixed attack bandwidth budget C_{max}, find the traffic distribution matrix A_t that can maximize the cost of the origin server

Adaptive Attack Strategy

- What Did the Agent Learn
  - Agent's Discovery: The optimal solution is traffic dispersion
  - Mechanism
    - The algorithm automatically flattens the Zipf distribution
    - A sharp increase occurs in the active working set, which exceeded the cache capacity of the mid-tier servers
    - It triggers severe cache thrashing, disrupting the system's statistical regularity
  - The algorithm discovers that dismantling the pattern of requests is far more devastating than merely increasing the volume.

Evaluation and Analysis

- Packet Drop
  - We assume that the server will keep the waiting time within 5s
  - Requests in the queue will be dropped when their waiting time exceeds 5s
- Even when the attack traffic is very small (30/second), there is still an obvious attack effect when the attack occurs
- As the attack traffic increases, the packet loss rate increases linearly.

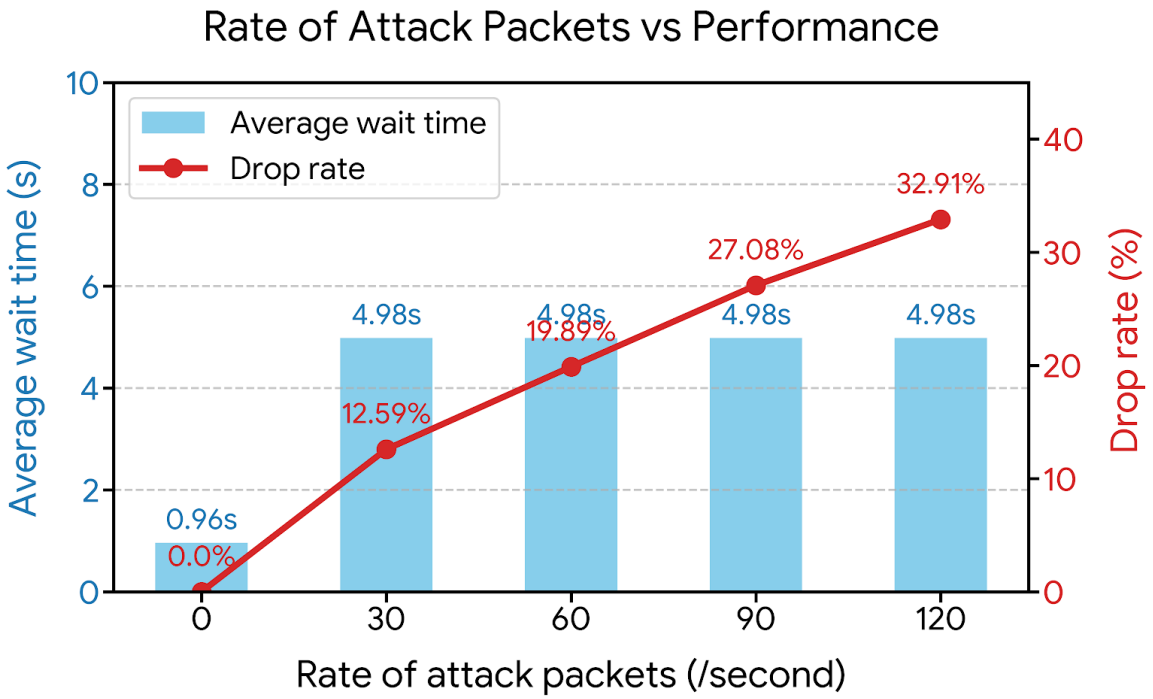

Evaluation and Analysis

- Cache Hit Ratio
  - Edge servers cache high popular content
  - Mid-server cache medium to high popularity content

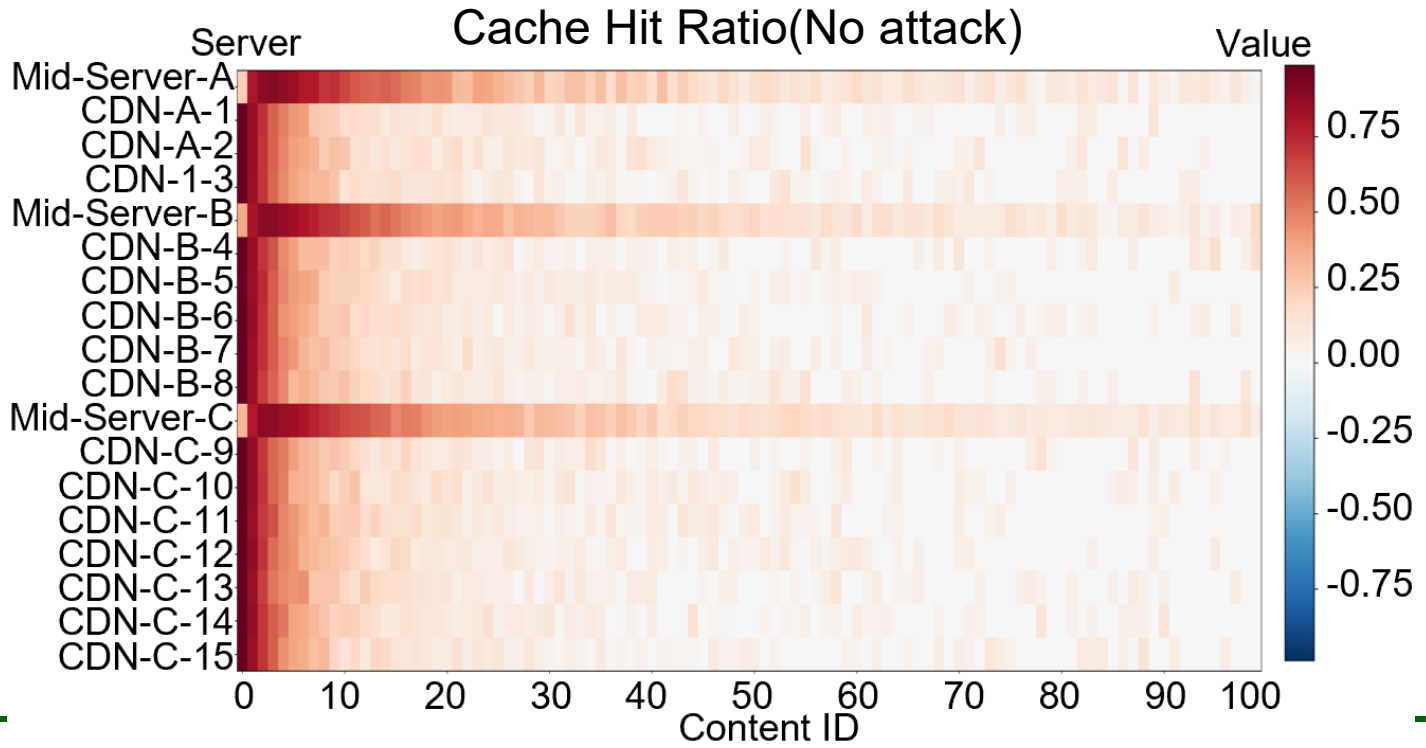

Evaluation and Analysis

- Cache Hit Ratio Change
  - The attack decrease the hit ratio of medium to high popularity content
  - Cache missed popular content that is frequently requested will be sent to the origin server

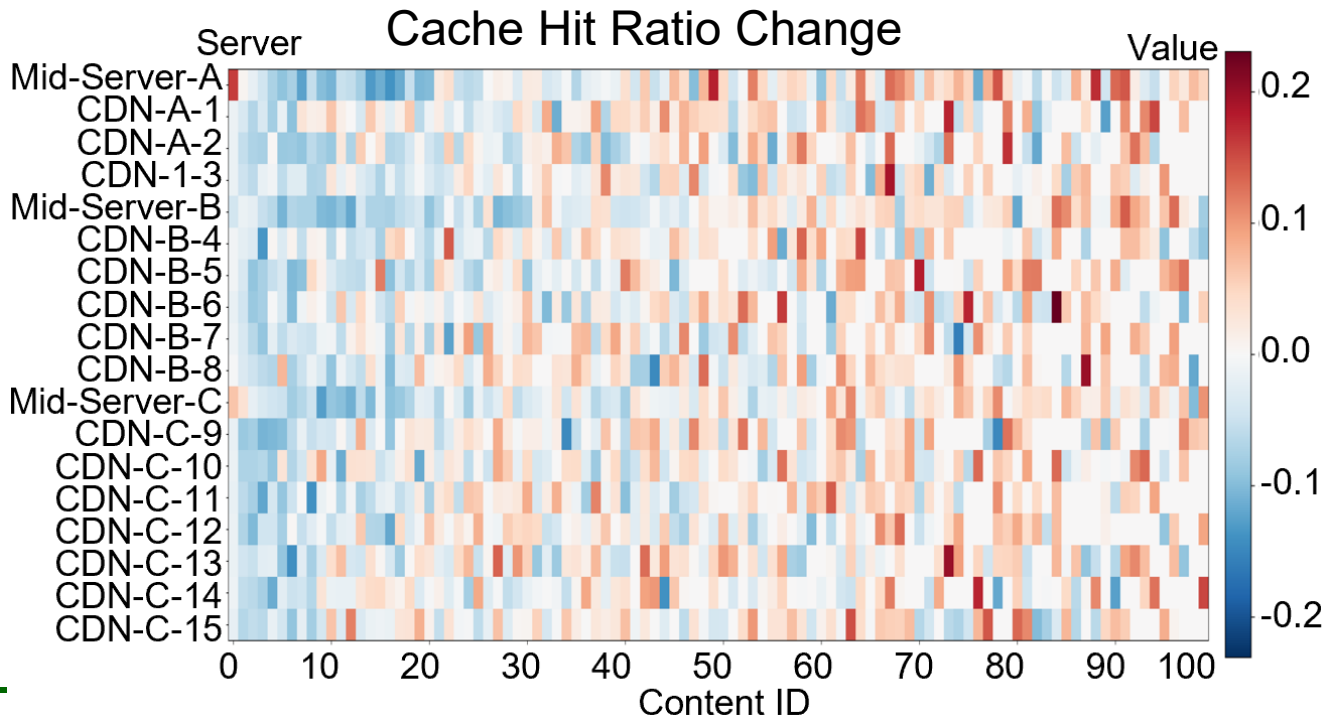

Evaluation and Analysis

- Request Input Rate
  - In a well-functioning CDN, most requests will remain at the edge servers
  - A small number of cache missed requests will also remain in the mid-server
  - Although the origin server serves a lot of users in multiple regions, only a small portion of the requests are sent there

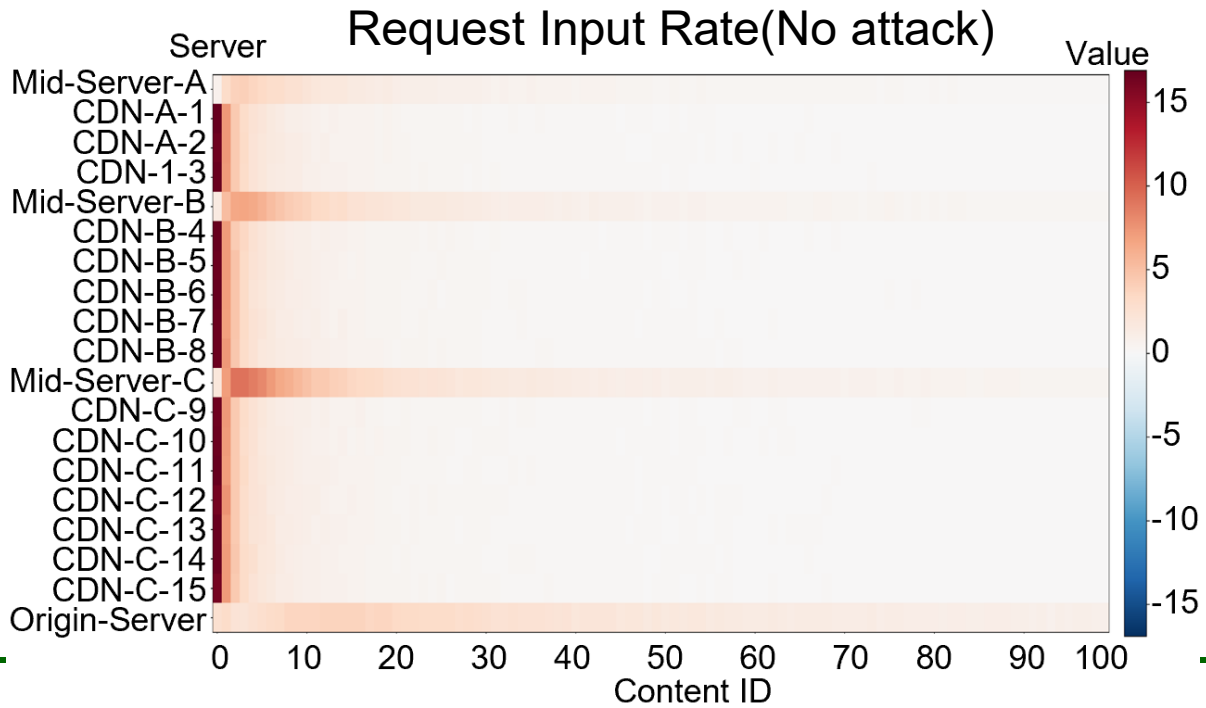

Evaluation and Analysis

- Request Input Rate Change
  - Due to the attacker's disruption of the cache space, the number of requests forwarded to the mid-server has increased significantly
  - This also leads to a lot of requests eventually flowing to the origin server

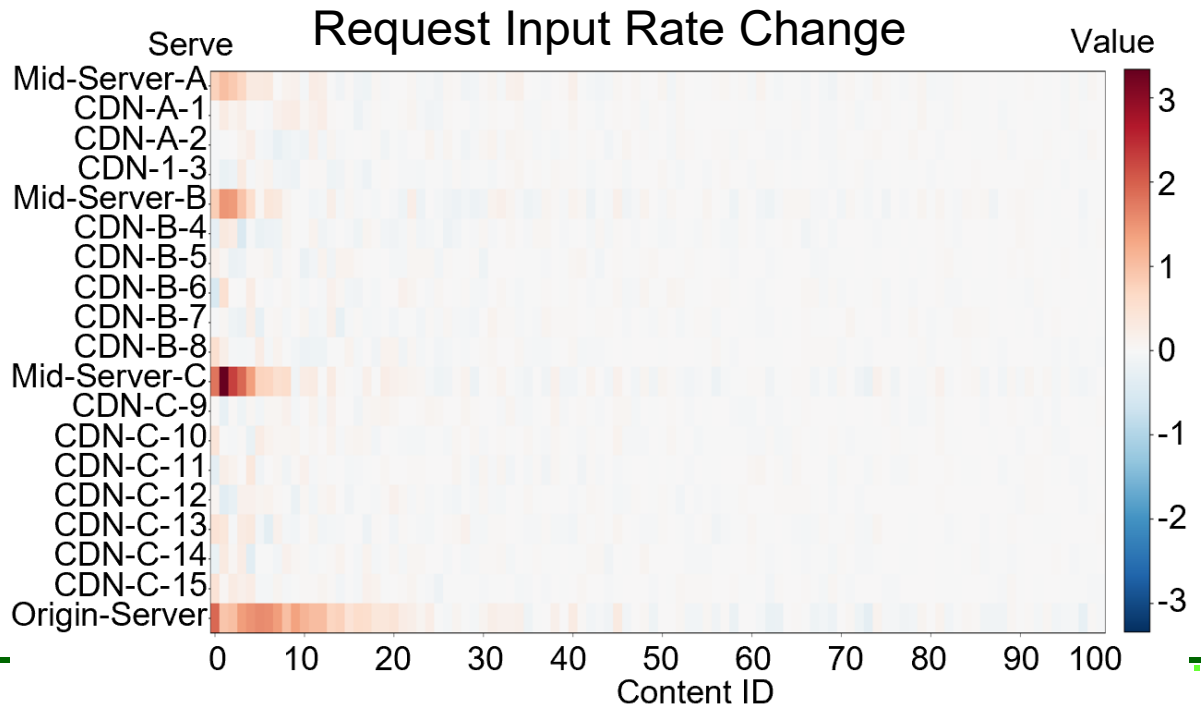

Detailed Explanation

- Why did the request rate at mid-edge servers for popular contents increase when CPA attack occurred?
  - On edge server, popular contents' hit ratio decrease but request rate increase a little
  - It makes more popular contents arrive mid-server

Detailed Explanation

- Why did the hit ratio at mid-edge servers for popular contents decrease even when the request rate increased?
  - Zipf parameter increase, popularity become more concentrative and hit ratio increase
  - When attack occurs, request rate disperse, like Zipf parameter decrease, popularity become more dispersive and hit ratio decrease

Conclusion

- Vulnerability of Hierarchical CDNs
  - Multi-tier caching architectures are highly fragile against adaptive, distribution-based CPA, even at low capacities
- The Power of Dispersion over Volume
  - The optimal attack strategy is not bandwidth exhaustion, but maximizing traffic dispersion
  - Flattening the Zipf distribution systematically dismantles caching logic via Cache Thrashing