

Locating CDN Edge Servers with HTTP Responses



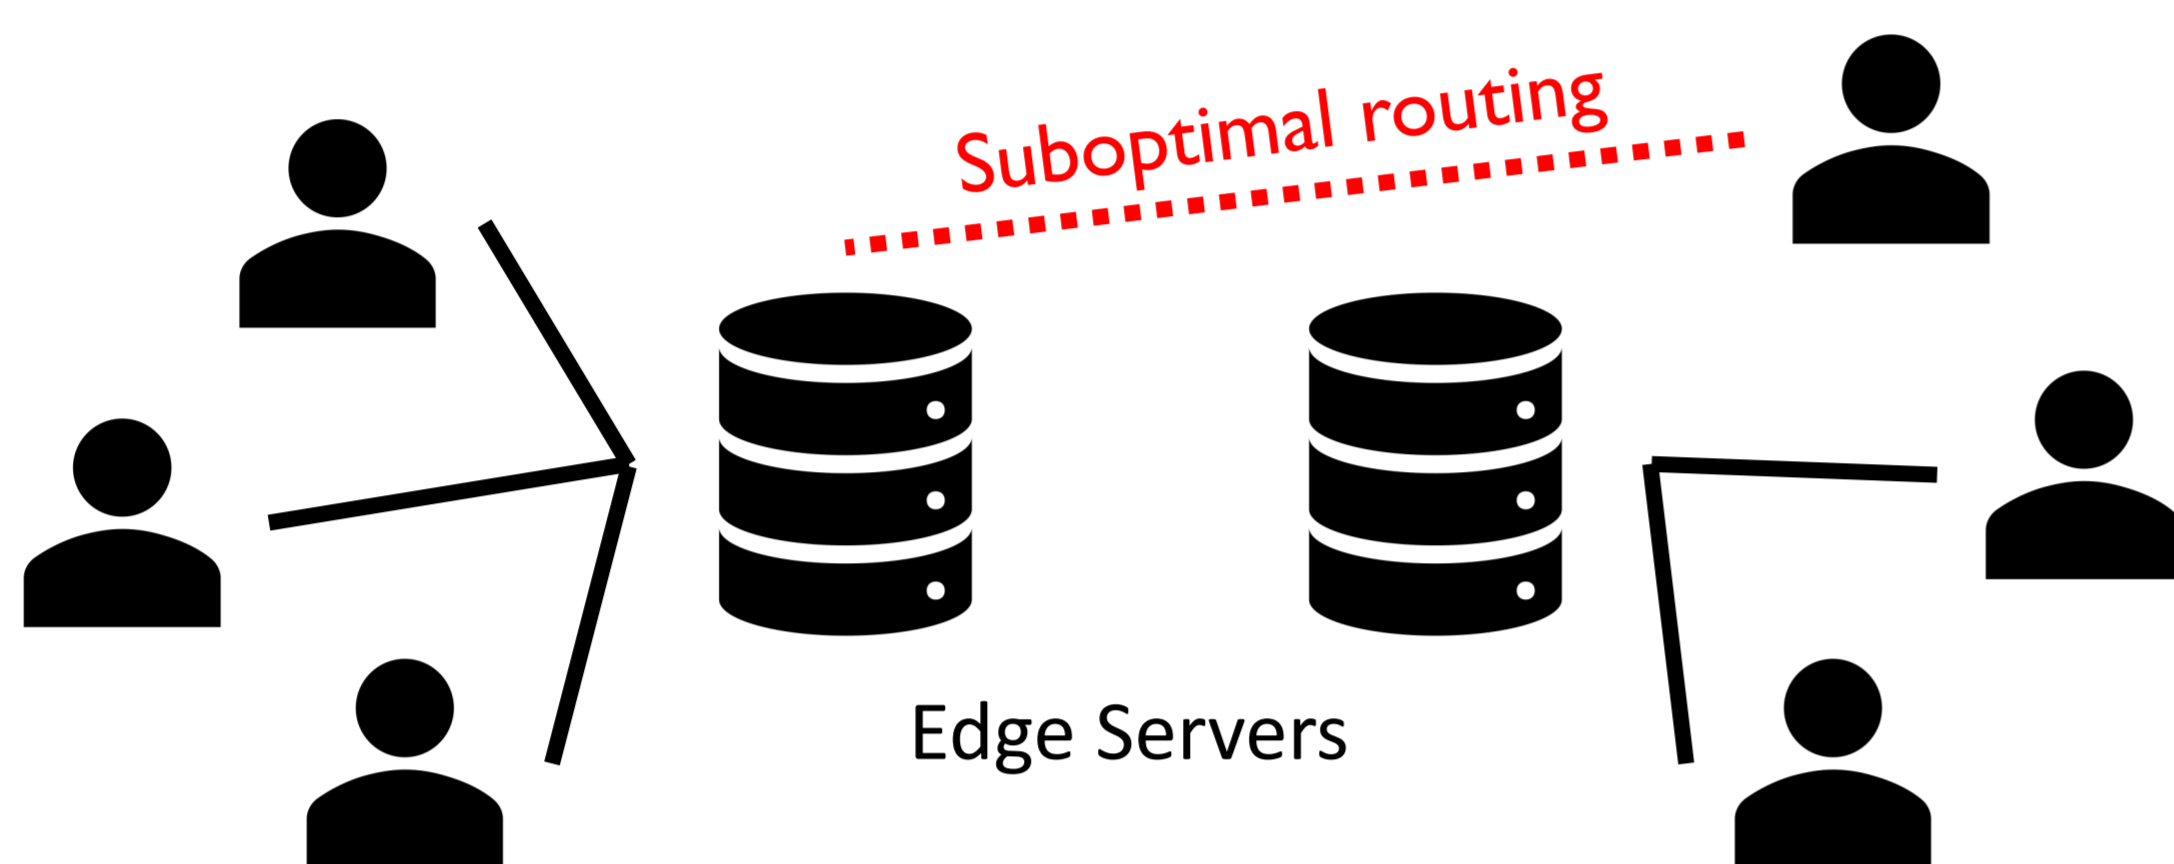
Run Huang, Mengying Zhou, Tiancheng Guo, Yang Chen

{runhuang19, myzhou19, tcguo20, chenyang}@fudan.edu.cn



Background and Motivation

- Content Delivery Networks (CDNs) reduce latency by providing resources to users from their nearest server.
- Many CDN providers rely on **IP Anycast** to find the "best" routing paths.
- Such routings may be suboptimal, leading to poor performance.



To monitor and understand routing inefficiencies, one'll need to know the **physical locations** of CDN edge servers.

Existing Solution

- ❖ GeolP databases are **inaccurate**
- ❖ Latency-based geolocation algorithms are **inconvenient**
- ❖ Using geohints in DNS names has **limited applicability**

Is it possible to have all their advantages without the drawbacks?

Our Idea

Many CDN providers would also embed geographic hints (geohints) in HTTP headers, e.g.,

CDN77: `x-77-pop: singapore2SG`

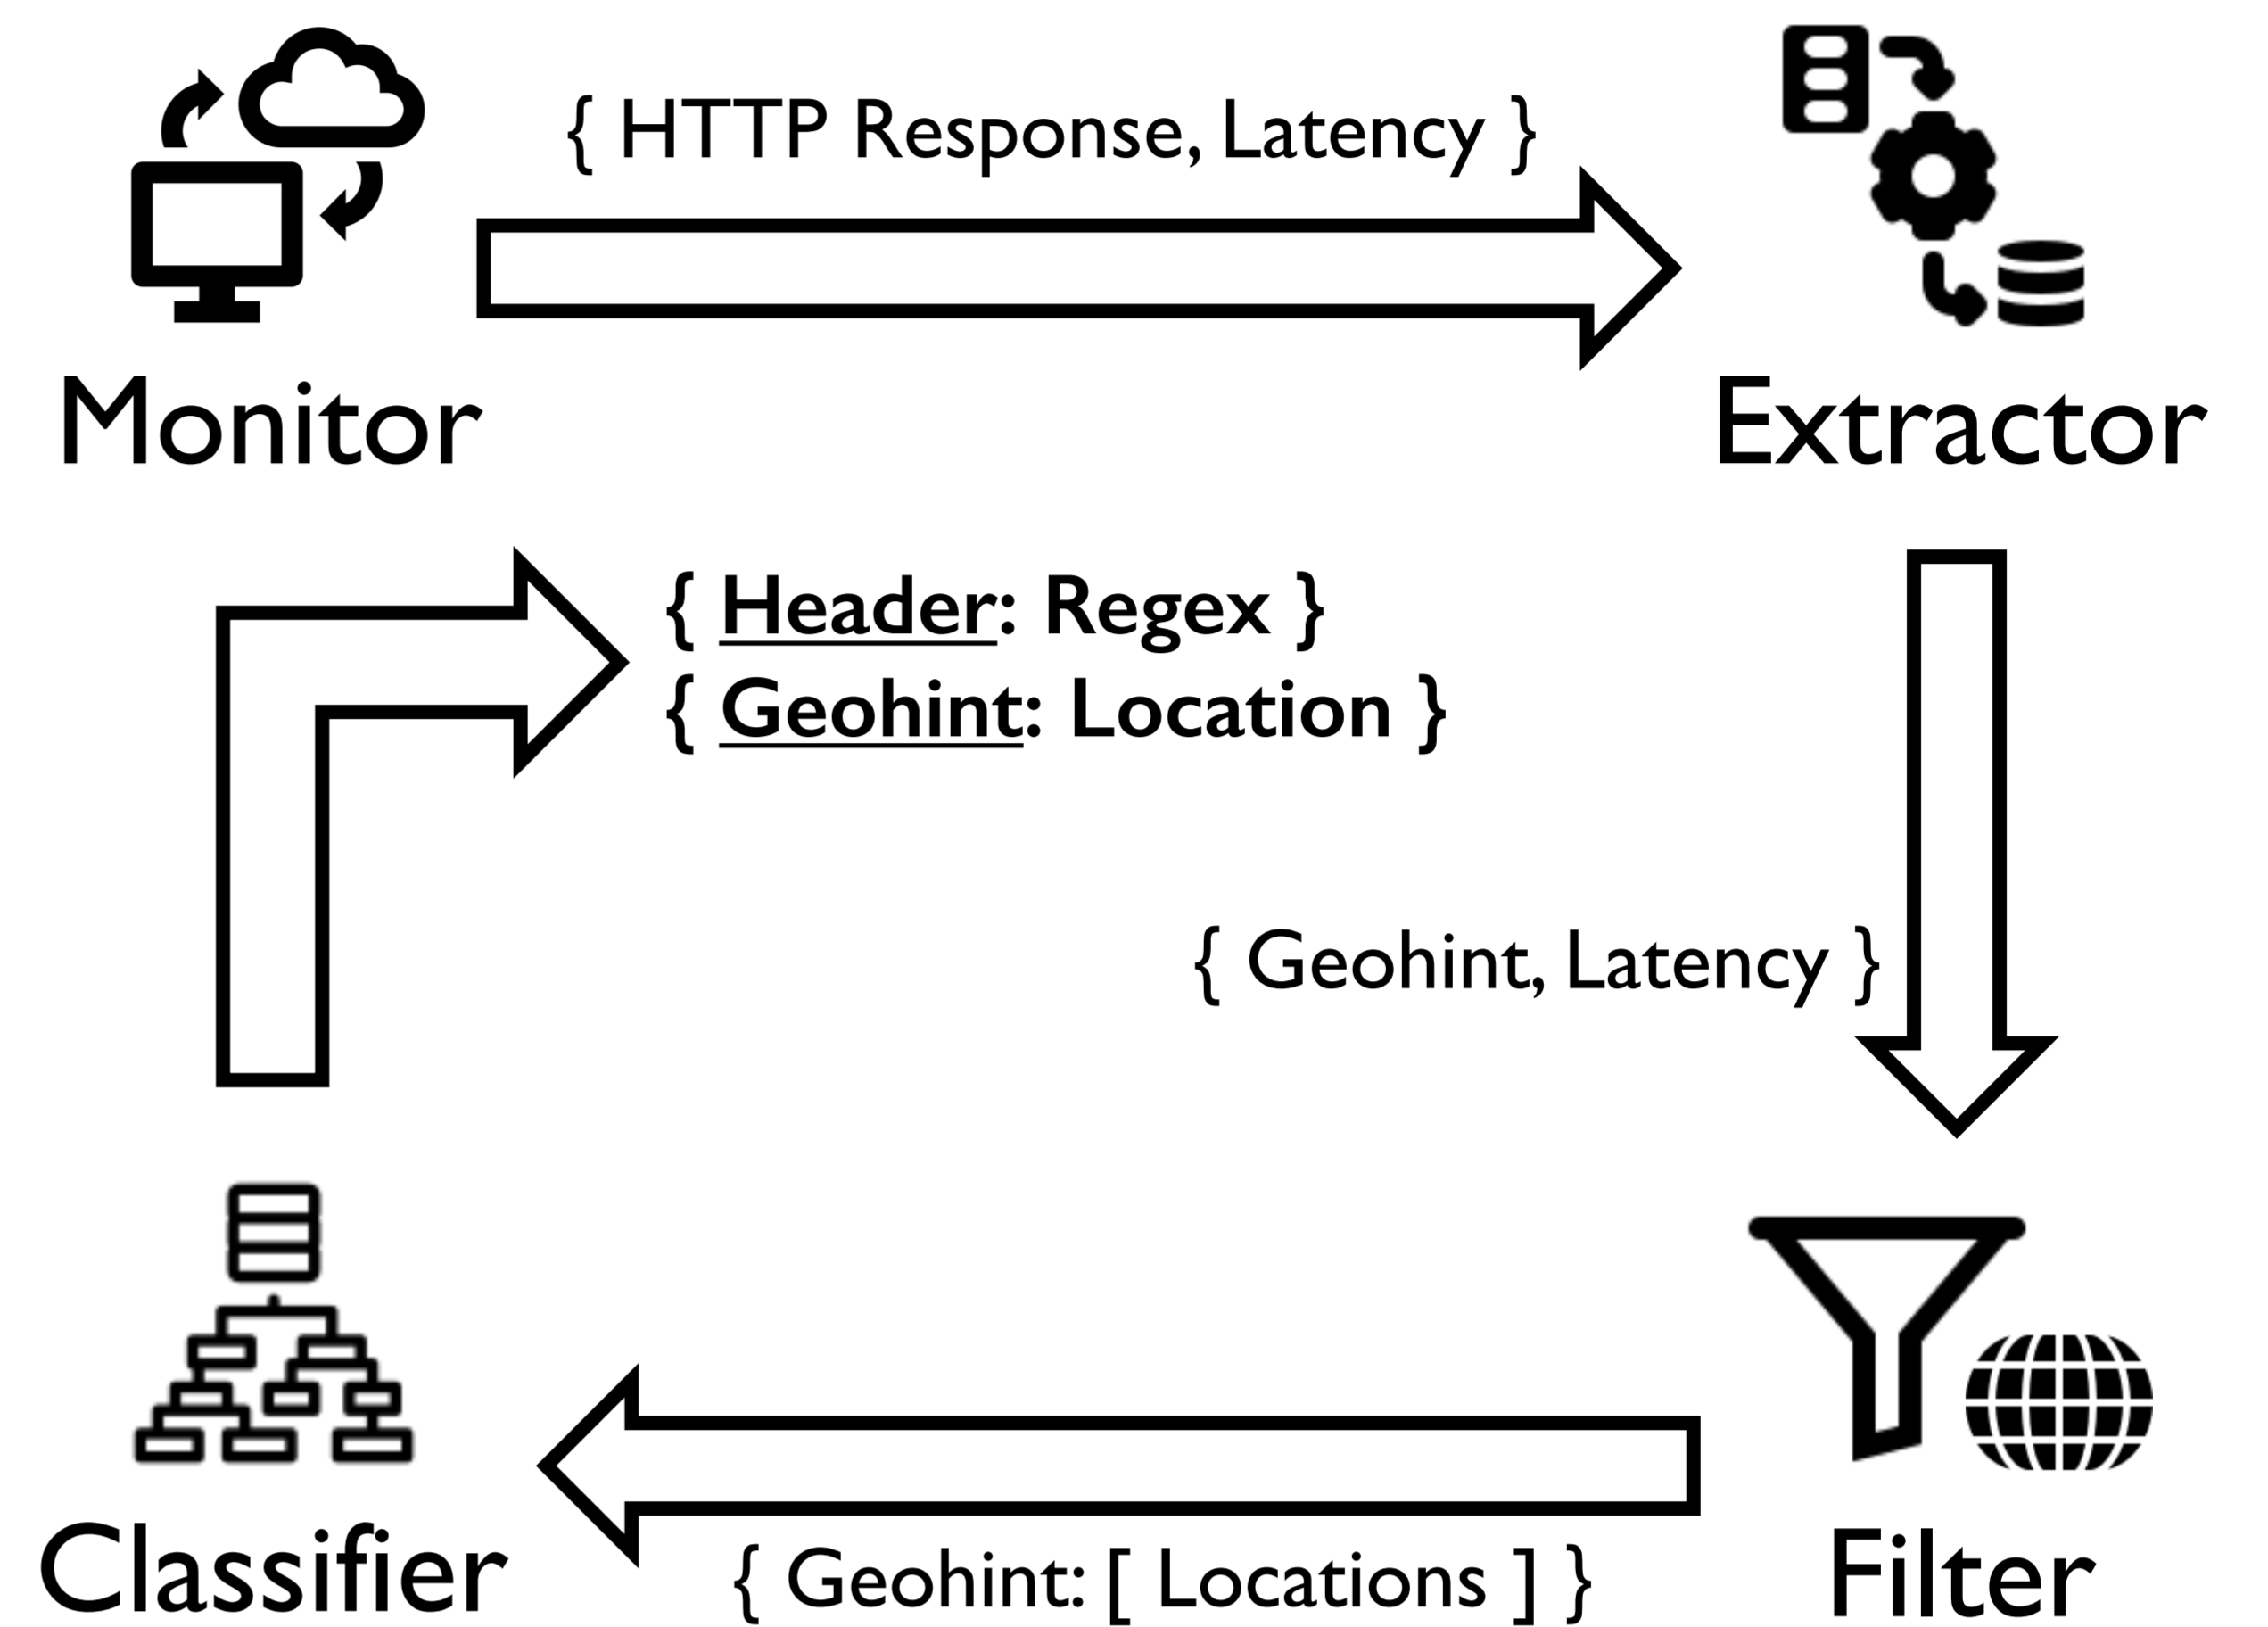
To use these geohints for server geolocation, we need to:

1. Identify and extract them
2. Find their corresponding locations

Figure 1: probe distribution



Methodology



- **Monitor**: Collects HTTP responses and latencies with 229 probes (Figure 1)
- **Extractor**: Identifies and extracts geohints using TF-IDF
- **Filter**: Detects Speed-of-Light violations to filter out ineligible geohints
- **Classifier**: Determines the corresponding location of each confirmed geohints

Results:

1. A **ruleset** for extracting geohints from specific headers
2. A **dataset** of the Geohint-to-City mappings

Evaluation

Metric:

Mean error distance between the predicted city and the actual location of each endpoint

- **Ground Truth**: Official Geohint-to-City mappings for 6 CDNs
- **Baseline**: The **best** of the results from Maxmind, IP2Location, and Hoiho for each endpoint

Our approach significantly outperforms the baseline (Table 1)

Table 1: Mean Error Distance

CDN	Mean Error (km)		CDN	Mean Error (km)	
	LocEdge	Baseline		LocEdge	Baseline
Amazon	25.41	1056.20	KeyCDN	26.65	571.26
CacheFly	20.74	2238.12	StackPath	30.71	2116.04
Fastly	62.34	1831.47	Tata	44.40	1086.03

Try web demo at <https://locedgex.web.app>