Space - Time Vehicle Tracking at the Edge of the Network

Zhuangdi Xu, Umakishore Ramachandran
Embedded Pervasive Lab, Georgia Tech

Background

- ➔ Cameras help to improve public/private safety due to its easy accessibility and low cost.
  - ◆ 2000+ cameras in Georgia Tech Campus.
- ➔ Reactively searching the camera streams after the occurrence of an event (e.g., a robbery) is unscalable.
  - ◆ Camera streams are recorded 24 x 7.

Space-Time Vehicle Tracking

- ➔ Track all vehicles over time and store their trajectories.
  - ◆ Answer queries from the stored trajectories.
- ➔ Proactive: video stream processing at ingestion time.
  - ◆ Circumvent time-intensive post-mortem video analytics
- ➔ Even low accurate result (i.e., more false positives) can help
  - ◆ Reduces the search space for more accurate analytics.

System Architecture

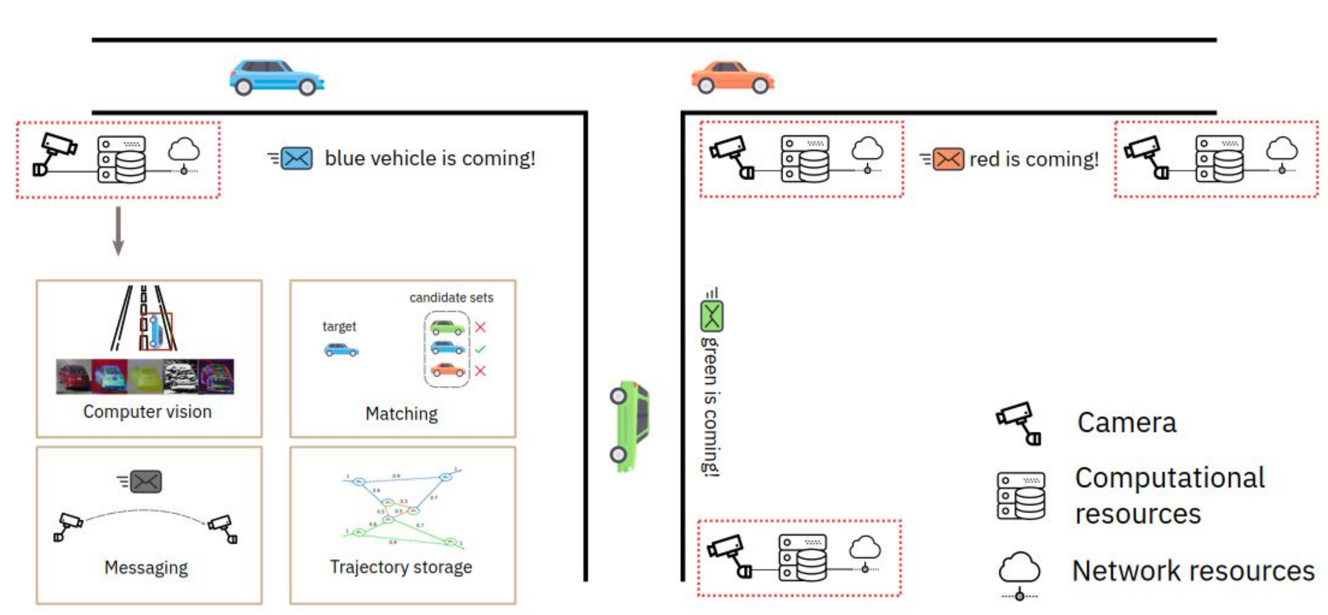

- ➔ Geographically distributed camera network
- ➔ Associated computational resources
- ➔ Well-connected network

Edge over Cloud

- ➔ Latency
  - ◆ Local nodes: ~2ms
  - ◆ Azure US East 2: ~50ms
- ➔ Bandwidth
  - ◆ Typical IP camera bandwidth: 2-24 Mbps[1]
  - ◆ Campus camera (1280 x 960) requires ~32 Mbps
- ➔ Administrative reasons
  - ◆ Edge => more controlled network
- ➔ Frame rate (for the Georgia Tech surveillance cameras)
  - ◆ 13-14 FPS on a local edge node
  - ◆ ~3 FPS on a cloud virtual machine

[1]: https://reolink.com/ip-camera-bandwidth-calculation/

Computer Vision

1). Detection

2). Tracking

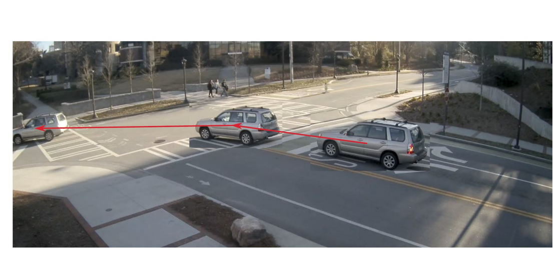

3). Event generation

```
{
    camera: ferst_hemphill,
    timestamp: 18:19-07/10/2019,
    features: {
        moving_direction: 90 (east),
        histogram: np.array([...])
    }
}
```

Messaging

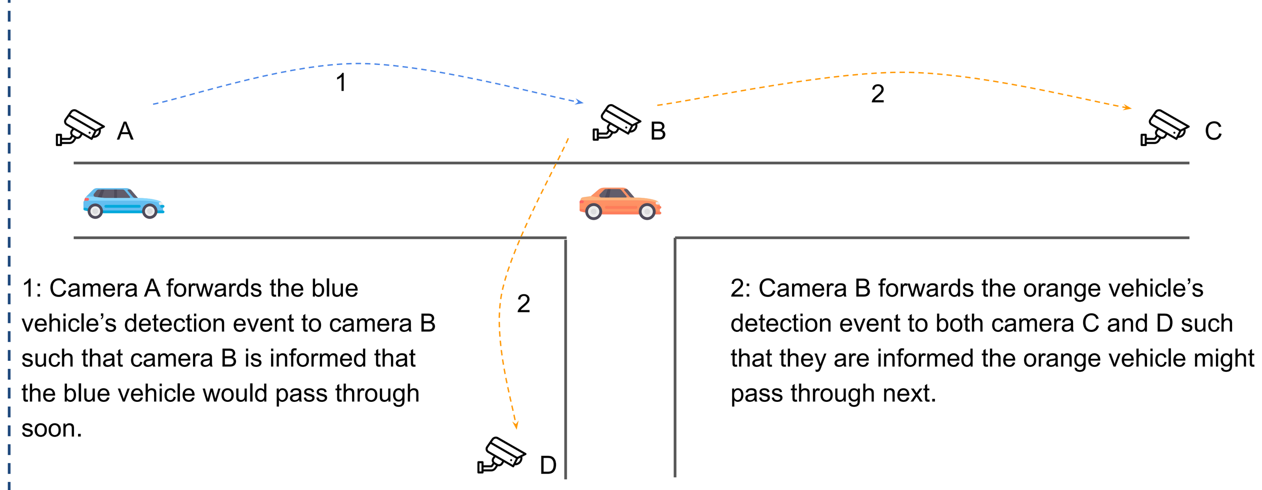

- ➔ Based on publish-subscribe.
  - ◆ topic = name of the camera
- ➔ Received events compose of candidate pool.
  - ◆ vehicles might pass through the camera in the future.

Trajectory Store

```
{
    camera: ferst_hemphill,
    timestamp: 18:19-07/10/2019,
    features: {
        moving_direction: 90 (east),
        histogram: np.array([...])
    },
    vertexId: V1
}
```

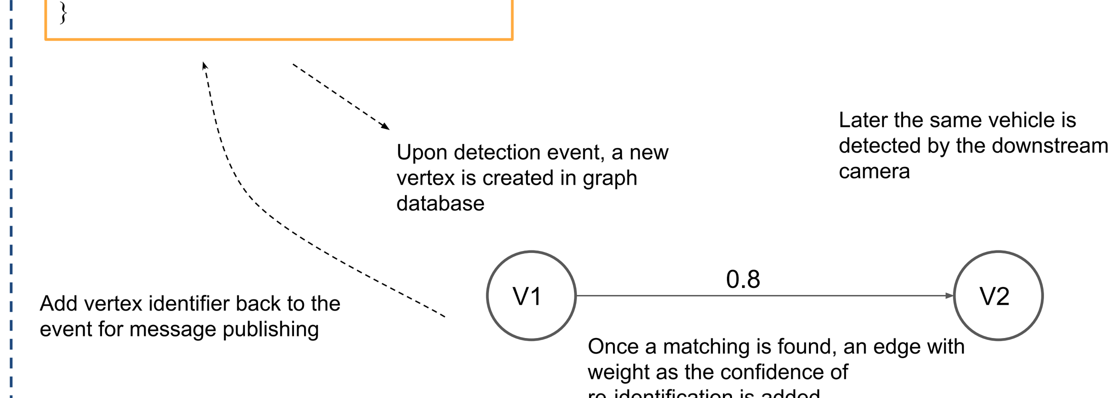

Trajectory of a vehicle is a path / collection of paths in the probabilistic graph.

Implementation

- ➔ Real camera streams from campus street cameras
- ➔ Cloud-version
  - ◆ Camera: Azure D4s v3(4 cores, 16G)
- ➔ Edge-version
  - ◆ Camera: 2 Raspberry Pi 3 B+s (1.4GHz 64-bit quad-core, 1G) and Coral EdgeTPU (USB accelerator)
- ➔ Demo:https://www.cc.gatech.edu/~zxu330/projects/STTR/index.html

Cloud-version

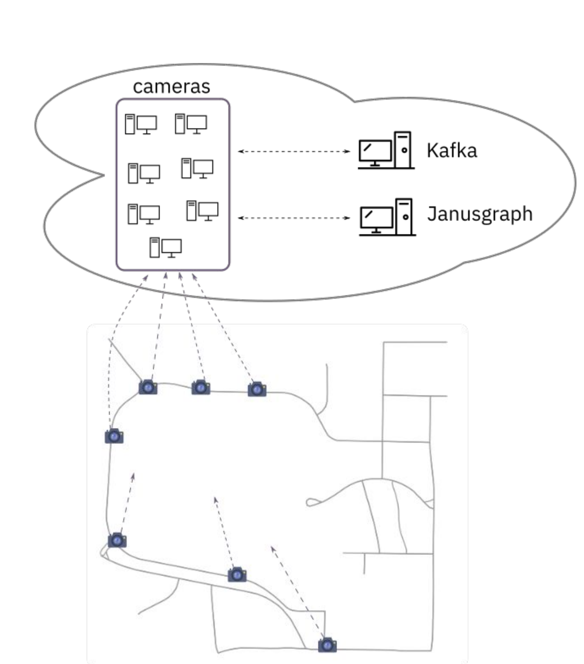

Edge-version

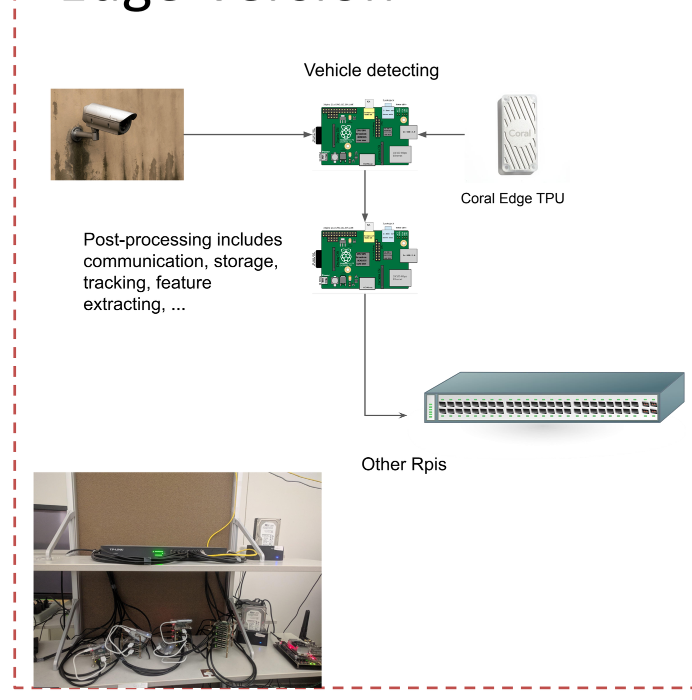

2019 NSF Cyber-Physical Systems Principal Investigators' Meeting
November 21-22, 2019 | Crystal City, Virginia

Award ID#: 1446801