ACM SenSys'24, Hangzhou, China.

Boosting Collaborative Vehicular Perception on the Edge with Vehicle-to-Vehicle Communication

Ruiyang Zhu, Xiao Zhu, Anlan Zhang[†], Xumiao Zhang, Jiachen Sun, Feng Qian[†], Hang Qiu[‡], Z. Morley Mao, Myungjin Lee*

University of Michigan [†]University of Southern California
[‡]University of California Riverside *Cisco Research

Presenter: Bo Chen, University of Illinois at Urbana-Champaign

Outline

- **Background and Motivation**
- System Design
- Evaluation Results
- Conclusion

Background: Collaborative Vehicular Perception

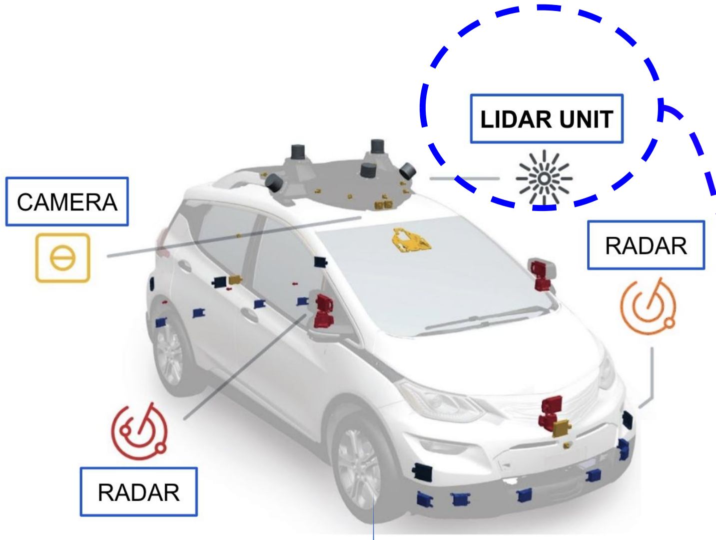

Limited range and occlusion

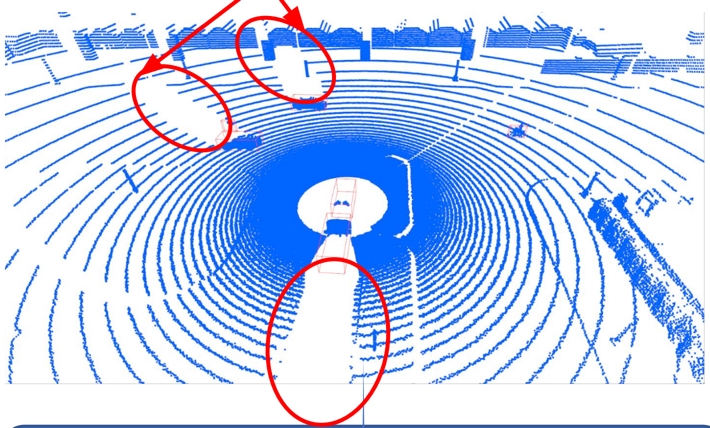

Extended Perception in occluded area

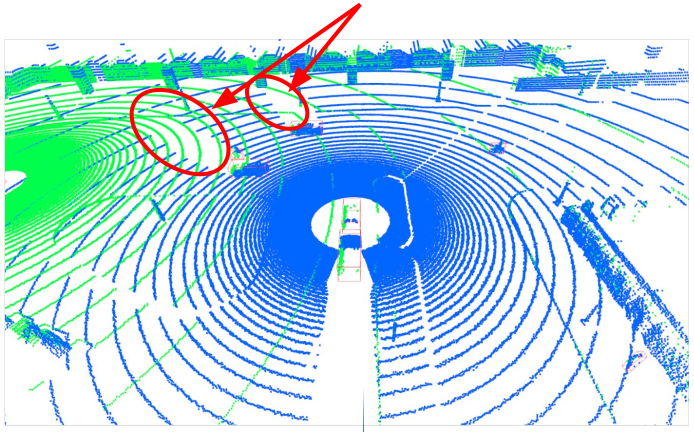

Sensors → Perception

Collaborative Vehicular Perception (**<u>CVP</u>**)

LiDAR 3D Point Clouds, 2D Camera Images, etc.

Drivable space detection, Object detection, etc.

Exchange sensor data <u>among CAVs</u> to extend sensing capabilities

Limitations of Existing V2V-only CVP

- **Lack of scalability** for larger number of vehicles [1, 2, 3]
  - AVR and Cooper [1, 2] focus only on a **2** vehicles scenario
  - Creates additional overheads by sharing ***N-1*** copies or use WiFi Broadcast mode [3], which creates congestion in V2V network
    - WiFi broadcast does not have MAC layer ACKs (no congestion control by default)[3]

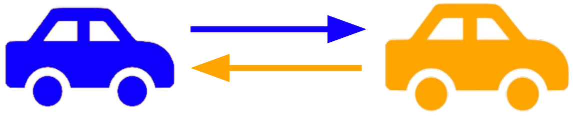

(1) 2 vehicles

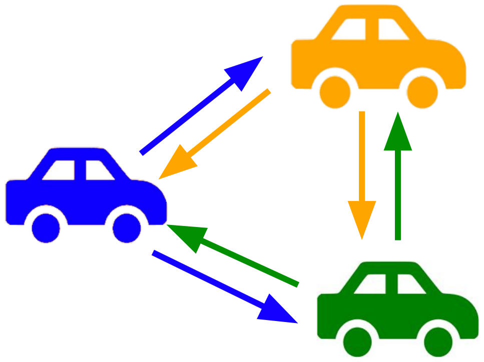

(2) >= 3 vehicles, N * (N-1) bandwidth sharing

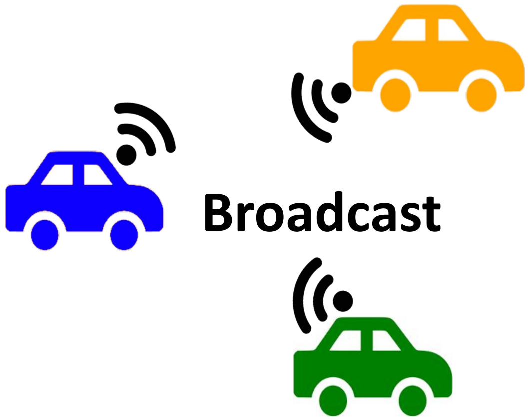

(3) >= 3 vehicles, through wireless broadcast

[1] AVR: Augmented vehicular reality. MobiSys 2018.
[2] Chen, Qi, et al. "Cooper: Cooperative perception for connected autonomous vehicles based on 3d point clouds." IEEE ICDCS, 2019.
[3] CarSpeak: A Content-Centric Network for Autonomous Driving. SIGCOMM 2012.

Limitation of Existing V2I-only CVP

- V2I (cellular) network conditions [1] can vary temporally and spatially
    - Ideally, C-V2X communication expects fast and stable network performance
    - Different carriers have different performance coverages

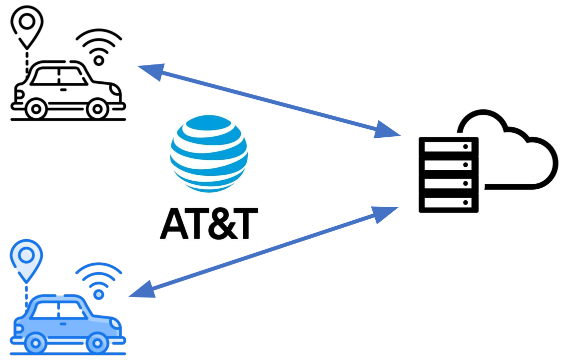

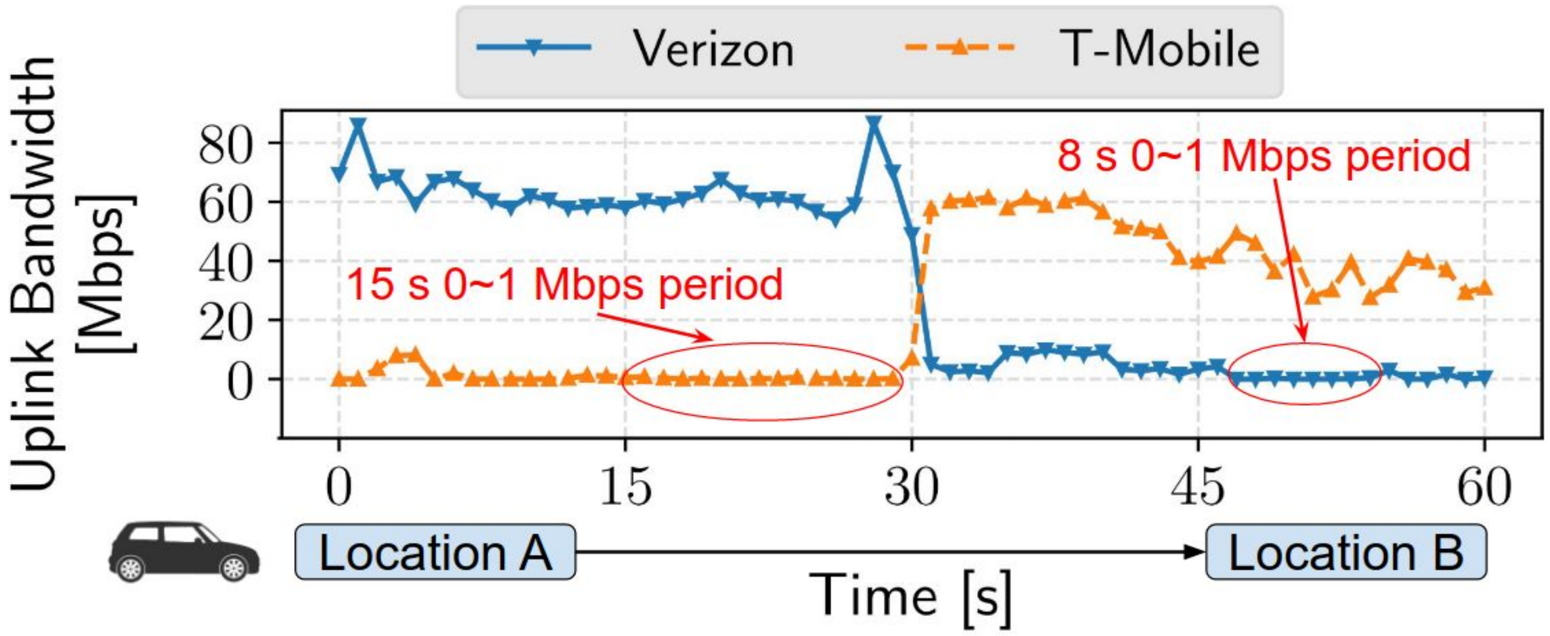

[1] EMP. MobiCom 2021.

Need for a hybrid V2V+V2I architecture

- Harbor: A **H**ybrid **ar**chitecture for colla**bor**ative vehicular perception that adaptively uses V2V and V2I connectivity
- **Key idea**: bridge V2I-disconnected vehicles (*helpees*) by strategically pairing them with V2I-connected vehicles (*helpers*) through V2V

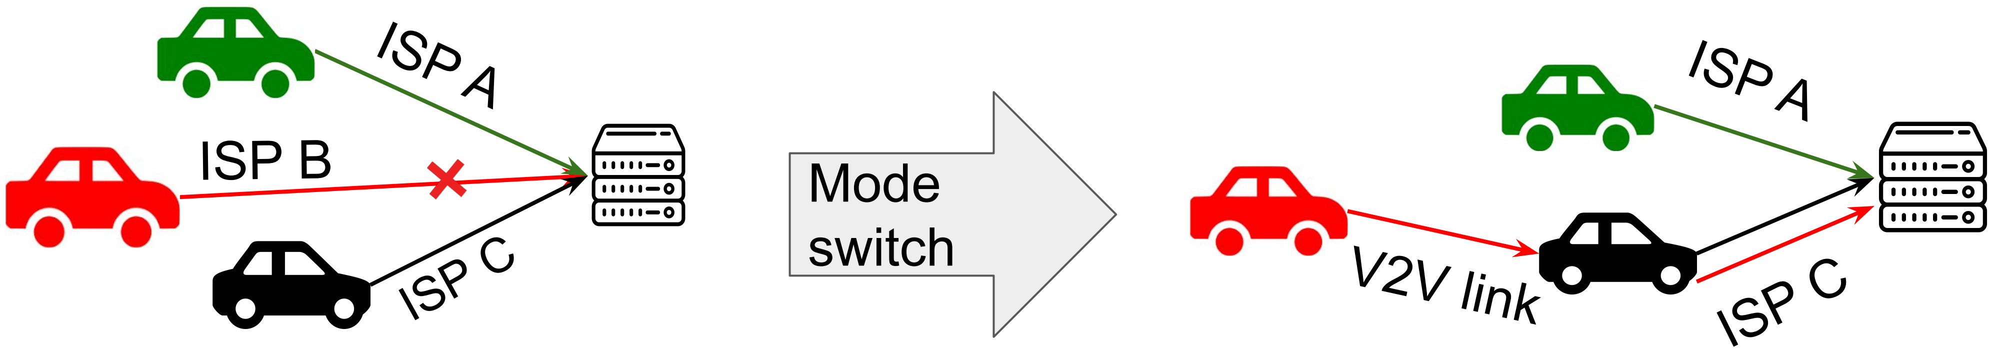

Pure V2I Mode

V2V+V2I Mode

Outline

- Background and Motivation
- **System Design**
- Evaluation Results
- Conclusion

Harbor design

- A hybrid architecture for collaborative vehicular sensing
  - Leverages both V2V and V2I network access
  - Flexibly manages V2V and V2I connections
- Strategic helper assignment
- Timely detection result delivery

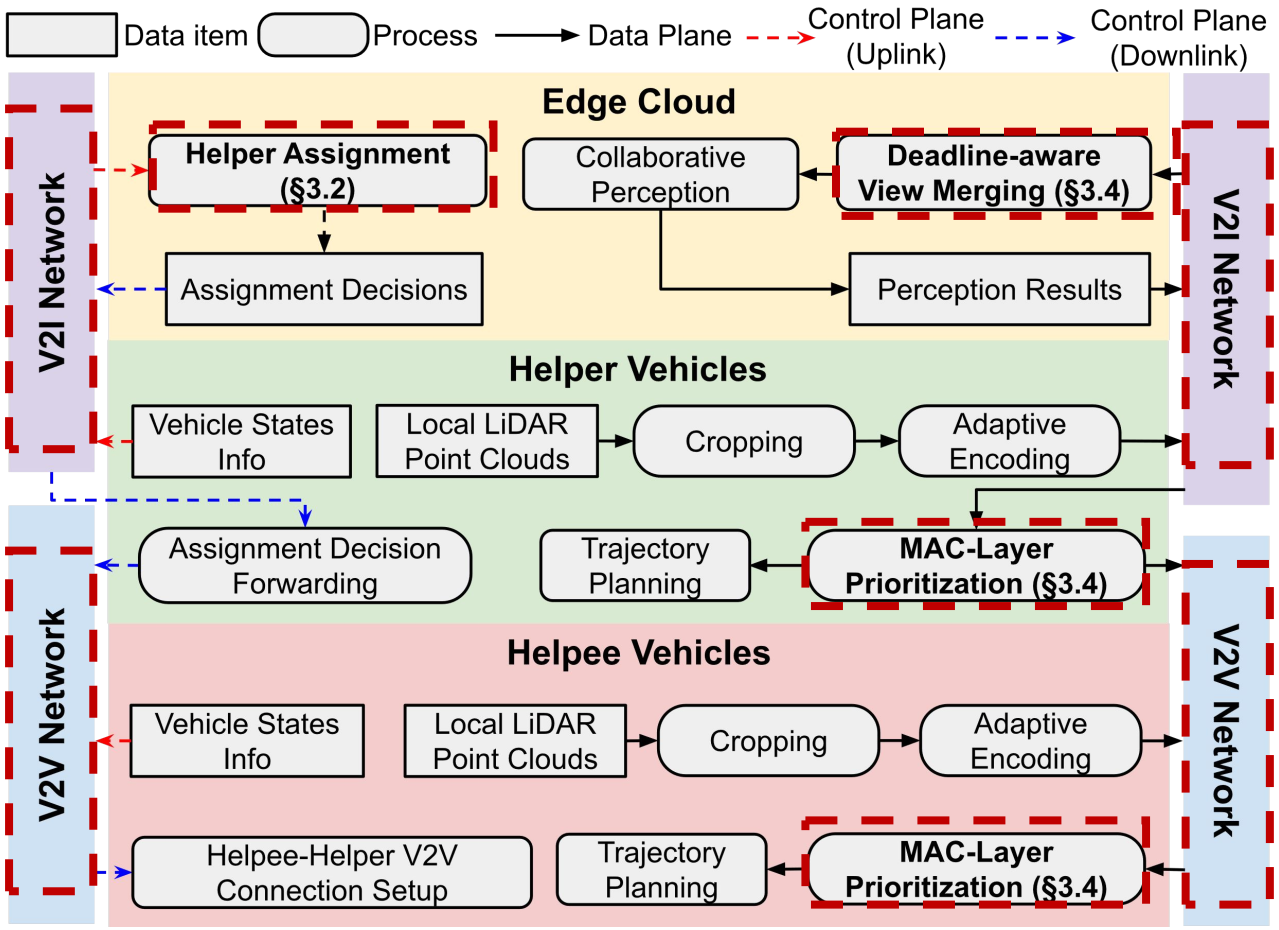

Strategic helper assignment - Goal

- Goal: Efficiently assign helpers to helpees to speedup sensor data upload
  - Different helpers have different V2I and V2V conditions
- End-to-end upload performance depends on network performance on both the V2V path and the V2I path
- Key idea: Identify performance impacting factors for V2V and V2I and combine these factors to score the assignments
  - **Direct measures**: V2V & V2I bandwidth
  - **Indirect measures**: factors that affects V2V & V2I bandwidth (e.g. distance, interference)

Identify performance impacting factors

- Both V2V and V2I paths' performance matter
  - V2I bandwidth can be measured from V2I path

    → Measured from ongoing data transfer to the edge server
  - However, V2V bandwidth is hard to quantify
    - Measured bandwidth by establishing connections? → Large overhead, can be inaccurate

    → Using indirect measures and heuristics
- Performance impacting factors
  - **V2I path: V2I bandwidth** and load of helper
  - **V2V path: Physical distance** between helper and helpee, **V2V network interference**

Strategic helper assignment - Overview

- Harbor jointly considers different factors and selects an assignment with the analytical model
  - A heuristic score function to merge effects caused by different factors
  - Normalize each factor to avoid a single factor becoming dominant
  - Run **periodically** to adapt to vehicle mobility & network state changes

Assignment $A = \{a_1, a_2, ..., a_n\}$

Assignment pair $a_k = (e_i, r_j)$

helpee (e_i) helper (r_j)

$$Score(a_k) = S_{dist}(a_k) + S_{intf}(a_k)$$

$$Score(A) = f(Score(a_1), ..., Score(a_m))$$

$$\textbf{maximize} \;\; Score(A) \;\; \textbf{s.t.} \;\; C_{bw}(a) \geq 0, \; \forall \, a \in A$$

Timely delivery of detection results

- It is vital to deliver remote detection results back to vehicles in time

- Key idea: **Application-layer** deadline awareness

- Application-layer deadline awareness
  - Server deadline for frame merging and detection
  - Vehicles use existing time synchronization methods (NTP) to sync with the edge

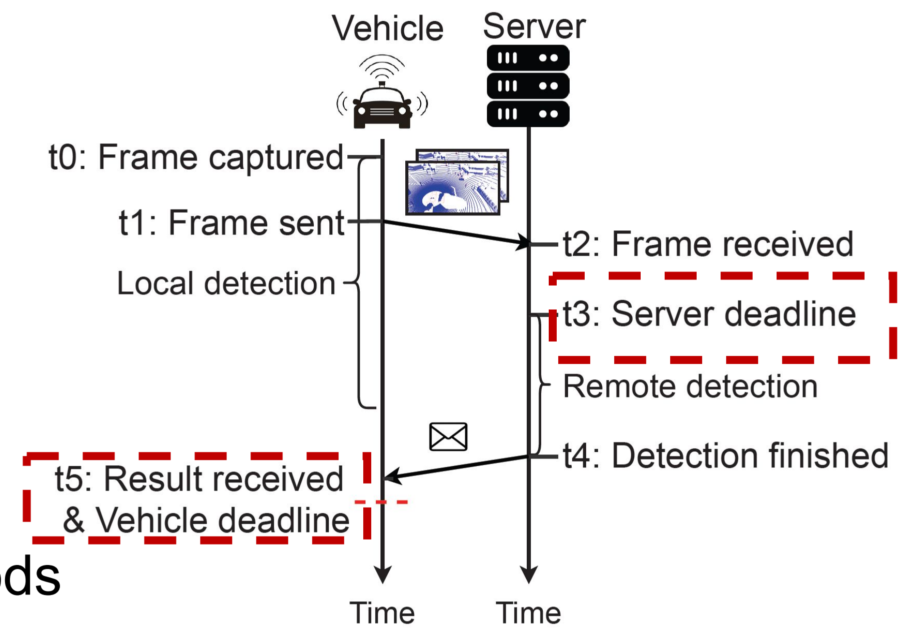

Timely delivery of results - MAC layer prioritization

- In V2V network, sensor data transmission (by helpee) and detection result delivery (from helper) contend for the shared wireless medium
- **Key Idea: Prioritize** small but latency sensitive traffic over bulk data transfer

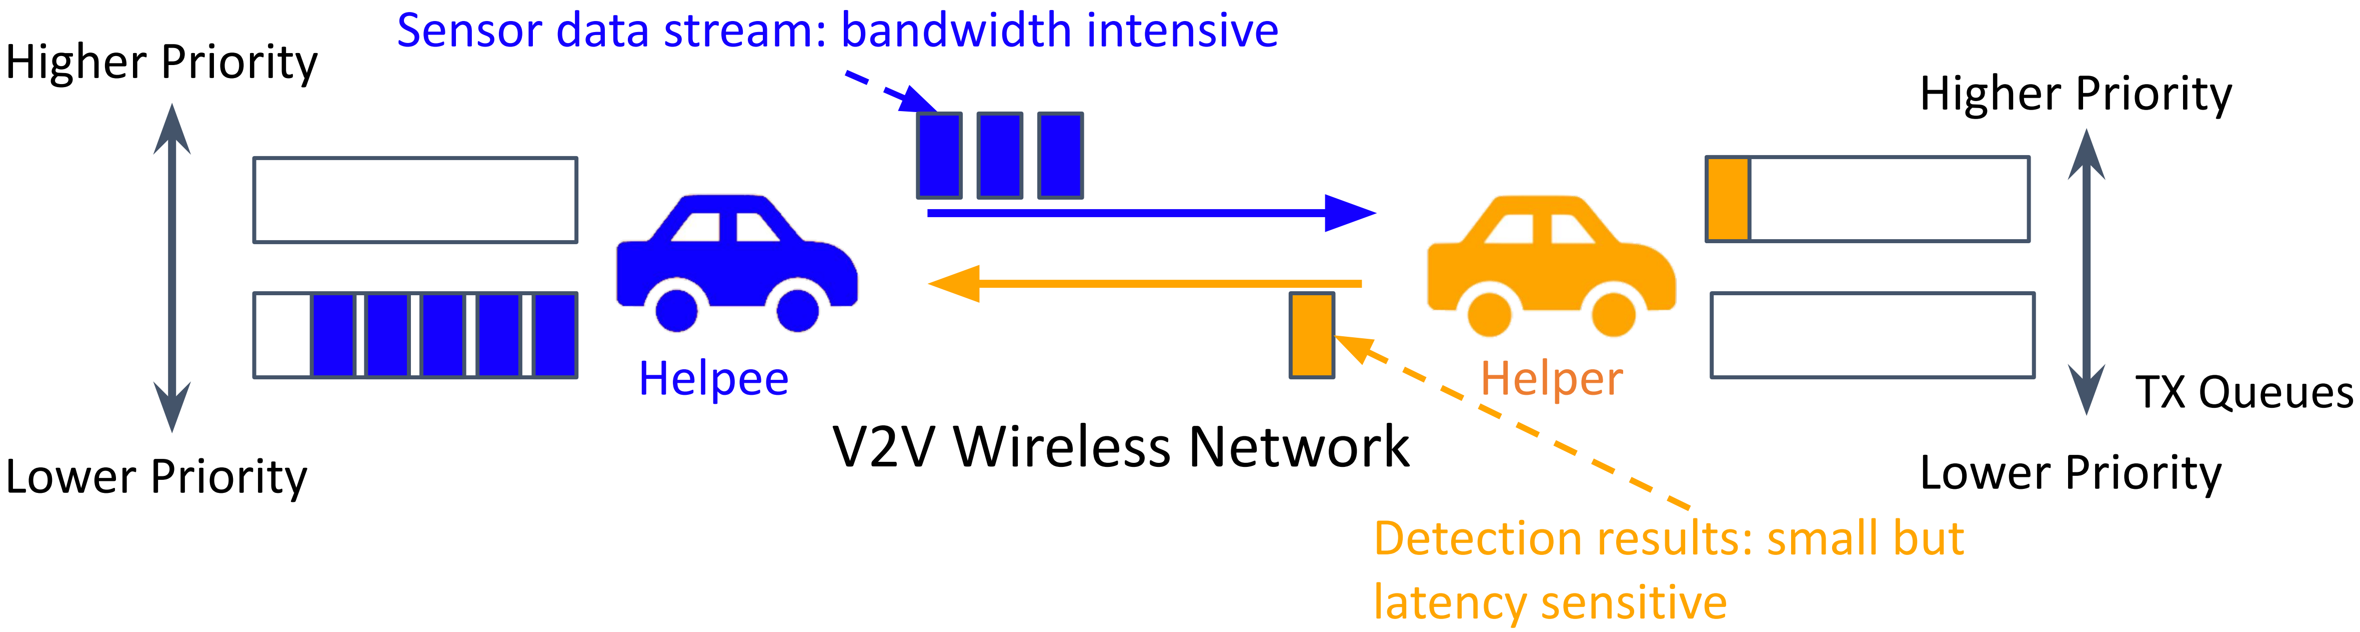

Outline

- Background and Motivation
- System Design
- **Evaluation Results**
- Conclusion

End-to-end evaluation - Emulation

- Evaluation Metrics
  - End-to-end detection latency (s) and detection accuracy
  - Compare with different baseline schemes

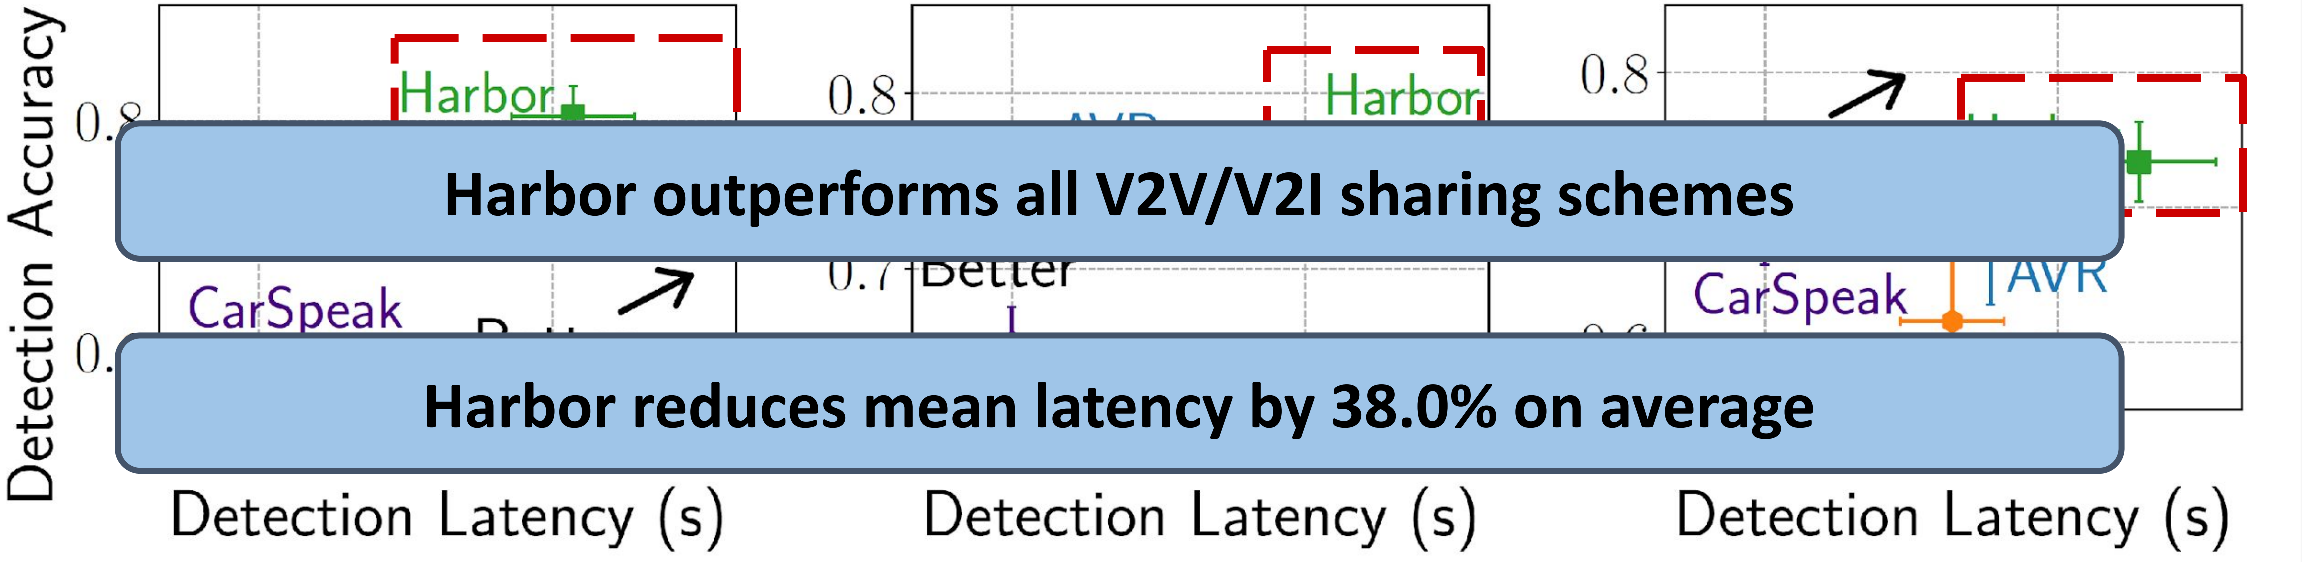

a) Better V2I Conditions b) Similar V2I and V2V Conditions c) Better V2V Conditions

End-to-end evaluation – Real-world Driving

- Harbor outperforms EMP, AVR and CarSpeak by reducing 18% - 37% of detection latency and improving 8.0% - 11.0% on accuracy.

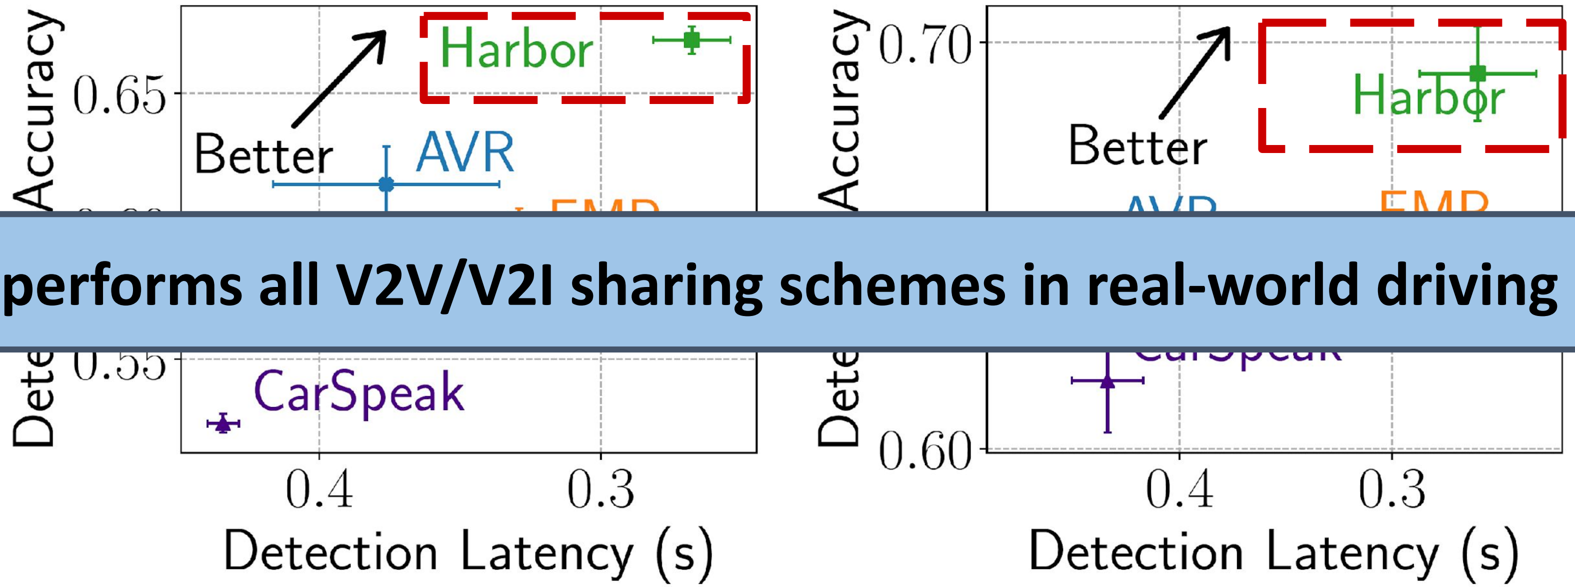

Harbor outperforms all V2V/V2I sharing schemes in real-world driving experiments

a) 3 Vehicles

b) 6 Vehicles

End-to-end Evaluation – Mcity Testbed

- Harbor achieves the best perception accuracy compared with EMP[1] and AVR[2]
- Harbor delivers better autonomous driving outcome by increasing the driving reaction time and reducing crash scenarios

Table 3: Collaborative perception accuracy under different collaboration schemes.

Traffic Scene	Object Detection Acc./Drivable Space Detection Acc.			
	Local-only	EMP	AVR	Harbor
Testbed-Overall	42.67/40.77%	69.84/57.89%	69.90/52.53%	**82.08/70.67%**
- Right turn	28.91/33.28%	59.60/53.84%	42.66/49.59%	**77.93/71.44%**
- Left turn A	34.60/44.57%	81.16/59.79%	79.49/57.25%	**83.55/63.87%**
- Left turn B	35.94/44.48%	79.50/58.12%	61.53/55.86%	**81.30/72.89%**
- Lane merge	71.25/40.78%	72.52/59.81%	80.46/47.43%	**87.57/74.51%**

Table 4: Additional driving reaction time (compared to Local-only) and outcome of different collaboration schemes.

Traffic Scene	Additional Reaction Time (s)/Driving Outcome		
	EMP	AVR	Harbor
Testbed-Overall	+ 0.76/1 crash	+ 0.58 /1 crash	**+ 1.26/0 crash**
- Right turn	+ 0.69/safe-pass	+ 0.41/crash	**+ 1.60/safe-pass**
- Left turn A	+ 1.05/safe-pass	+ 0.20/safe-pass	**+ 1.23/safe-pass**
- Left turn B	+ 0.73/crash	+ 1.13/safe-pass	**+ 1.42/safe-pass**
- Lane merge	+ 0.57/near-miss	+ 0.57/near-miss	**+ 0.79/safe-pass**

Notes: A near-miss occurs when the ego and target vehicle pass within 3 m of each other.

Outline

- Background and Motivation
- System Design
- Evaluation Results
- **Conclusion**

Conclusion

- Develop ***Harbor***, a hybrid system architecture for cooperative vehicular perception
- Harbor strategically assigns helpers to helpees
  - Harnesses performance impacting factors from different network layers
- Harbor performs timely detection result delivery
  - App-layer deadline awareness and MAC-layer message prioritization
- Harbor outperforms V2V and V2I collaboration solutions
  - Reduces up to *57.1%* in end-to-end latency and improves up to *12%* in detection accuracy
  - Result in significantly fewer collisions under dangerous driving scenarios

Thank You!

Our Team