Automobile Sensing and Intelligent Transportation Systems 2

CSE 40437/60437-Spring 2015

Prof. Dong Wang

Papers

Paper 3: Kumar, Swarun, et al. "Carspeak: a content-centric network for autonomous driving." ACM SIGCOMM Computer Communication Review 42.4 (2012): 259-270.

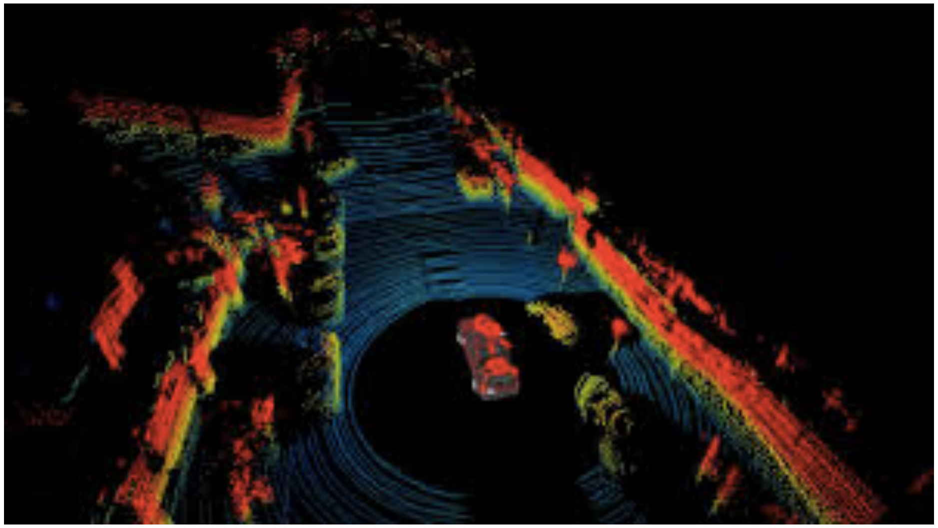

Much Interest in Autonomous Vehicles

Google's Autonomous Car

- Benefits include lower traffic congestion, higher fuel efficiency, improved productivity
- Projected to save **$100B/yr** in US alone [WPI'07]

Robots for Disaster Areas (Fukushima)

“Expect them on the road by 2020”
- General Motors

- Nevada and California legalized testing autonomous vehicles. Florida expected to follow.

- Autonomous Vehicles tested on Europe’s roads

Challenge: Safely Detecting Hidden Objects

- Sensors on a car see only line of sight objects

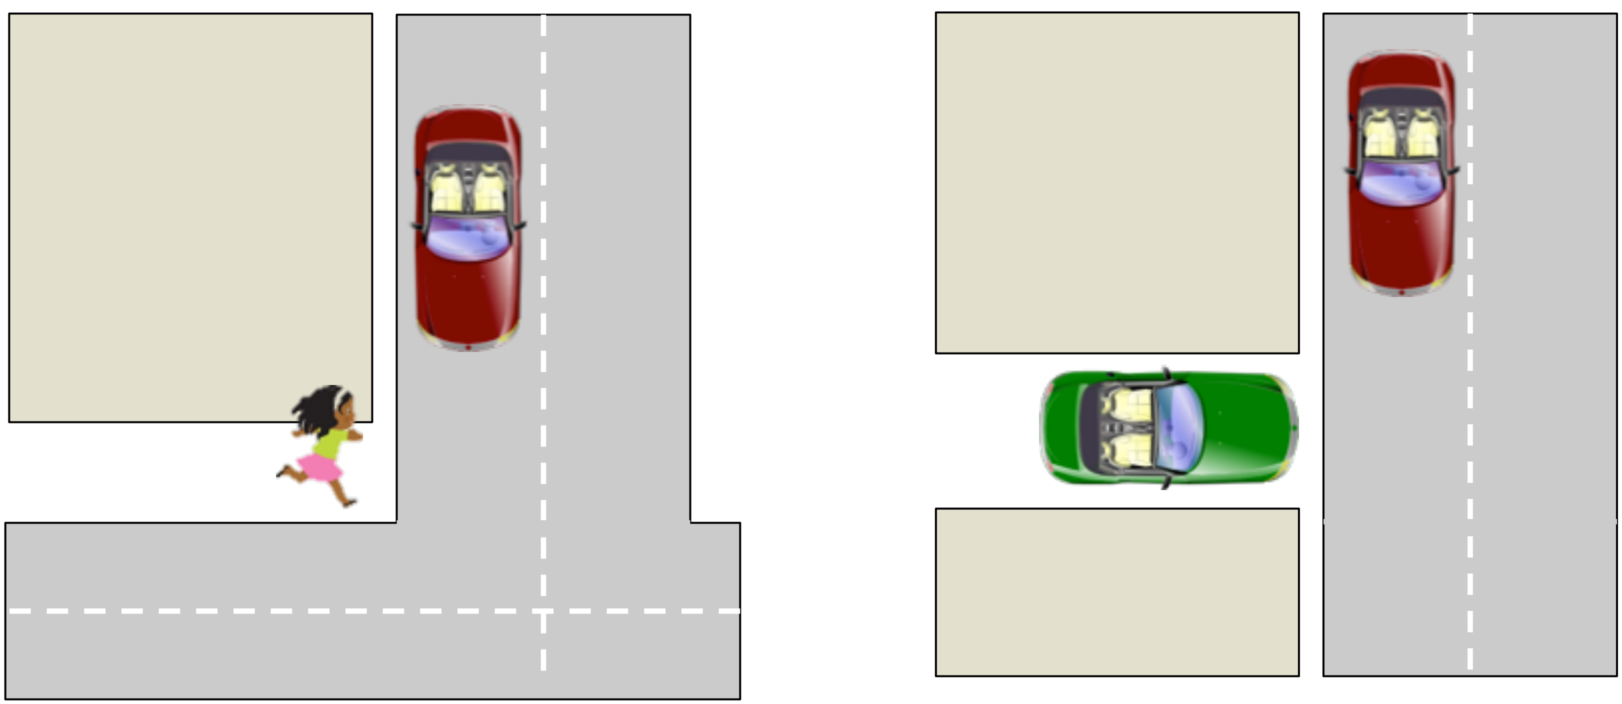

Challenge: Safely Detecting Hidden Objects

- Sensors on a car see only line of sight objects

- Hidden objects affect autonomous cars

  - "Google's autonomous car requires occasional human intervention to prevent accident"

  - "Future of autonomous driving depends on detecting hidden objects & blind spots" - DARPA Challenge [JFR08]

How can autonomous vehicles detect hidden objects?

Communication

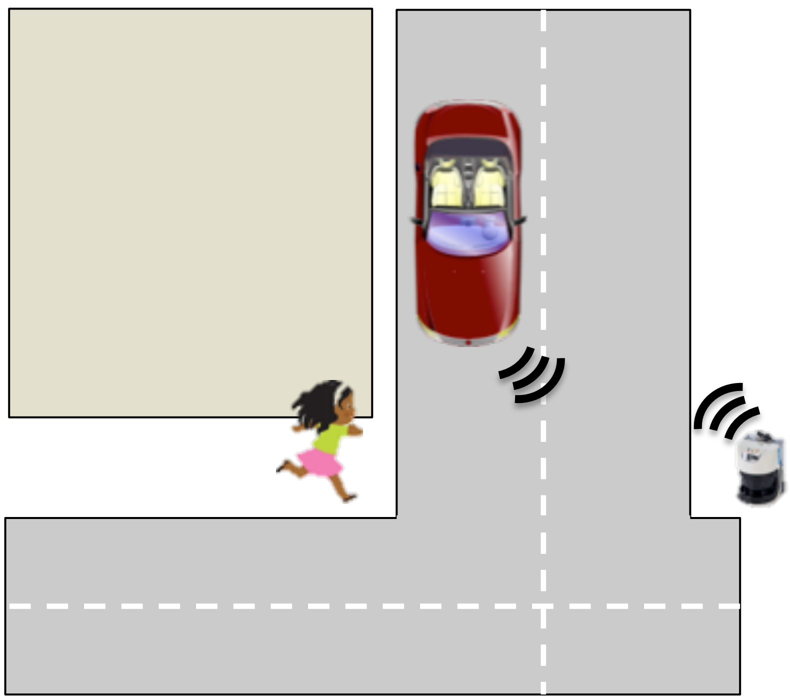

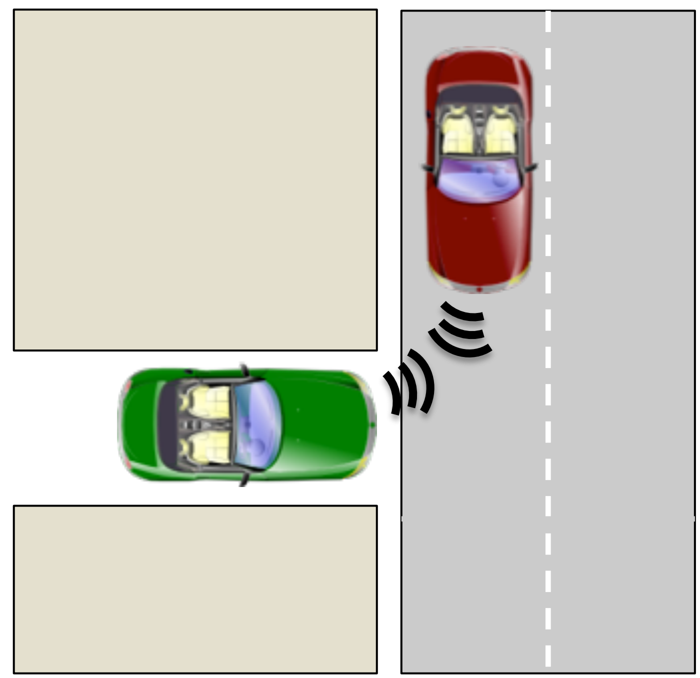

How can autonomous vehicles detect hidden objects?

Communication

- Expand field of view beyond line of sight
- Also valuable for human drivers – can react faster to objects they couldn't see

Simply use past work on VANETs?

VANETs(Vehicular Ad-hoc networks) typically oblivious to application

- Efficient routing
- Reliable message delivery
- … (etc)

But, not integrated with specific applications

Autonomous Driving needs Tight Integration with Communication

- Data is huge (Gb/s) and time critical
  - →Communication should focus on information most critical to the application

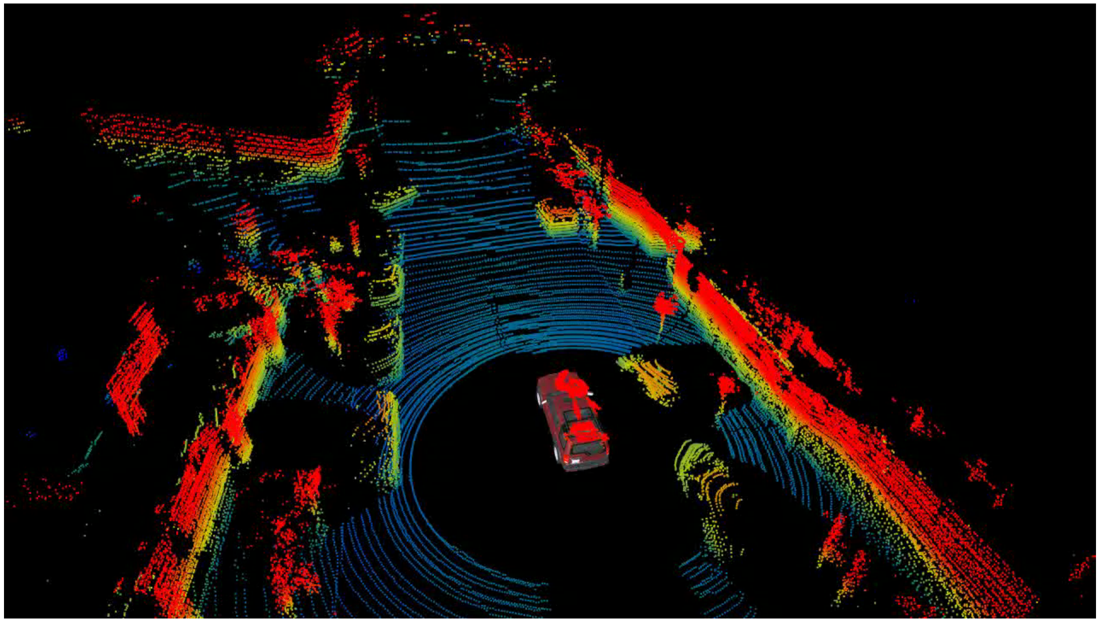

Autonomous Driving needs Tight Integration with Communication

- Data is huge (Gb/s) and time critical

  → Communication should focus on information most critical to the application

- Don’t know who has the desired content
  - In typical networks, you know your destination
  - Instead, autonomous car seeks sensor data from part of the road, e.g. an intersection
  - It doesn’t know which car has this information
  - Focus on content as opposed to accessing a particular destination

CarSpeak

- A communication system integrated with path planning and navigation in autonomous vehicles
- Has a content centric design
  - content, i.e. parts of road, is a first class citizen
- New MAC design where content, not senders, contend for the medium
- Implemented & evaluated on real autonomous vehicles

1. What is “content” and how do we represent it?

2. How do we disseminate this content?

1. What is “content” and how do we represent it?

2. How do we disseminate this content?

What is "content" and how do we represent it?

- Content is sensor data from cubes in the environment

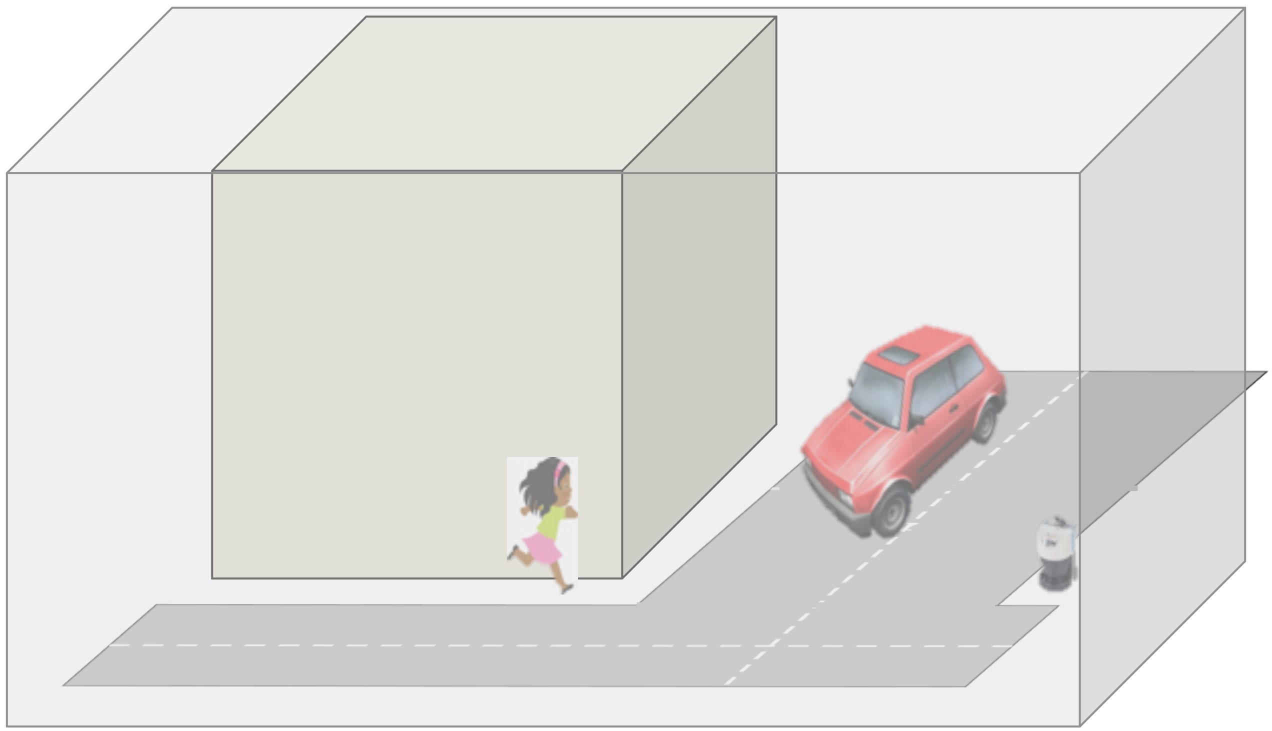

How do we represent these cubes in environment?

Ideally…

- Obtain low resolution view of environment
- Zoom in for higher resolution view of a smaller part of environment

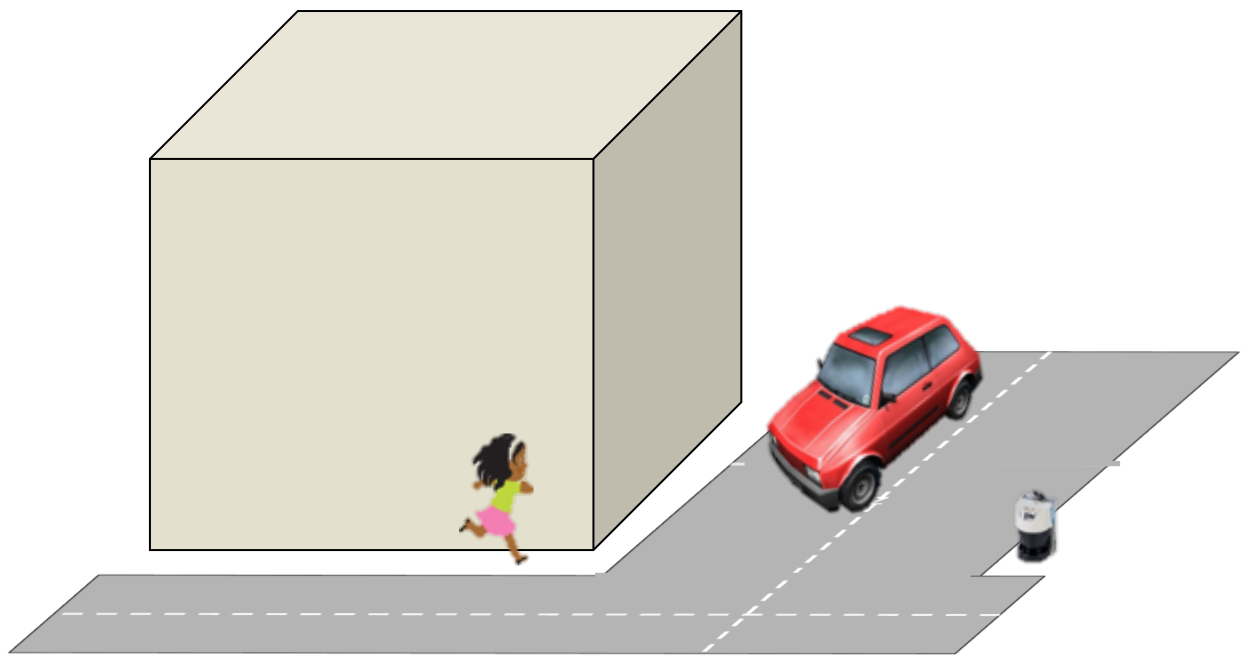

Ideally…

- Obtain low resolution view of environment
- Zoom in for higher resolution view of a smaller part of environment

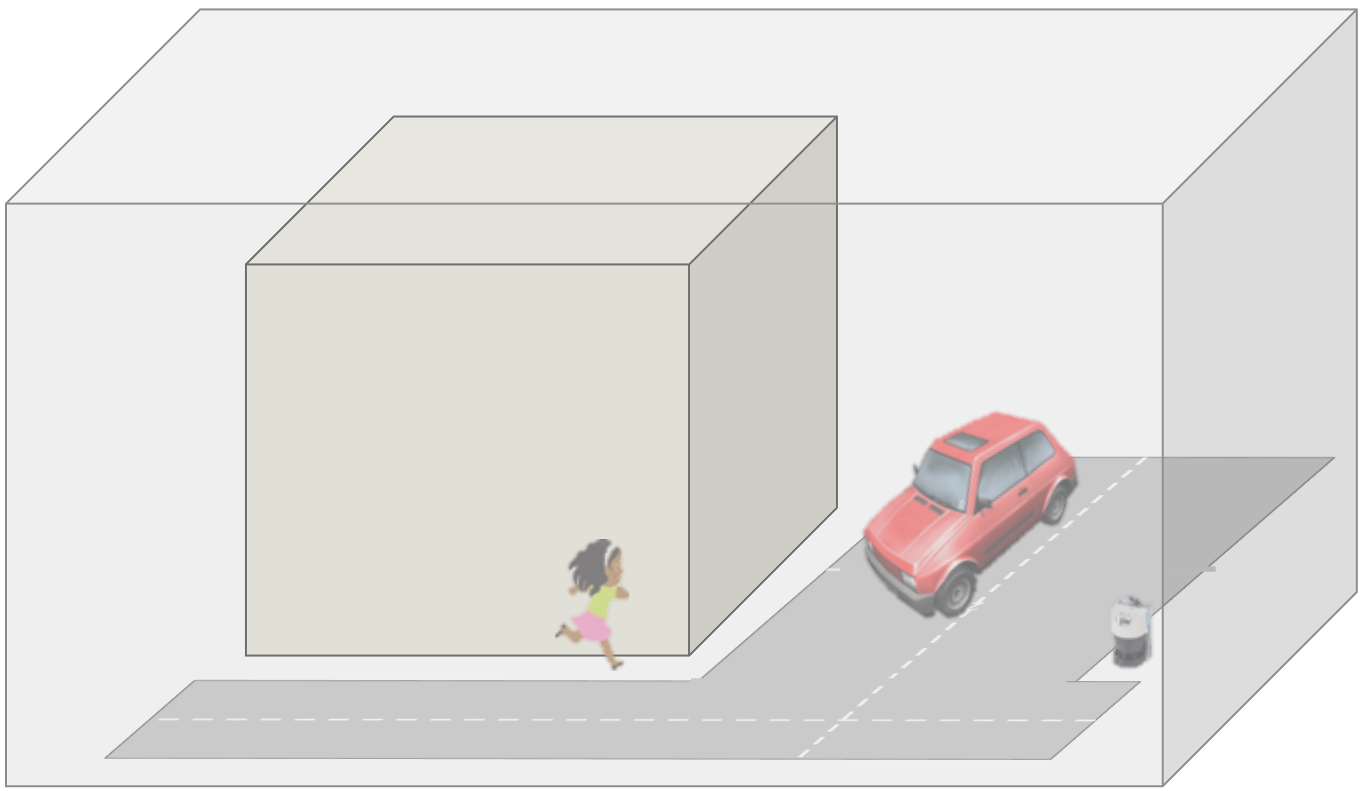

Ideally…

- Obtain low resolution view of environment
- Zoom in for higher resolution view of a smaller part of environment

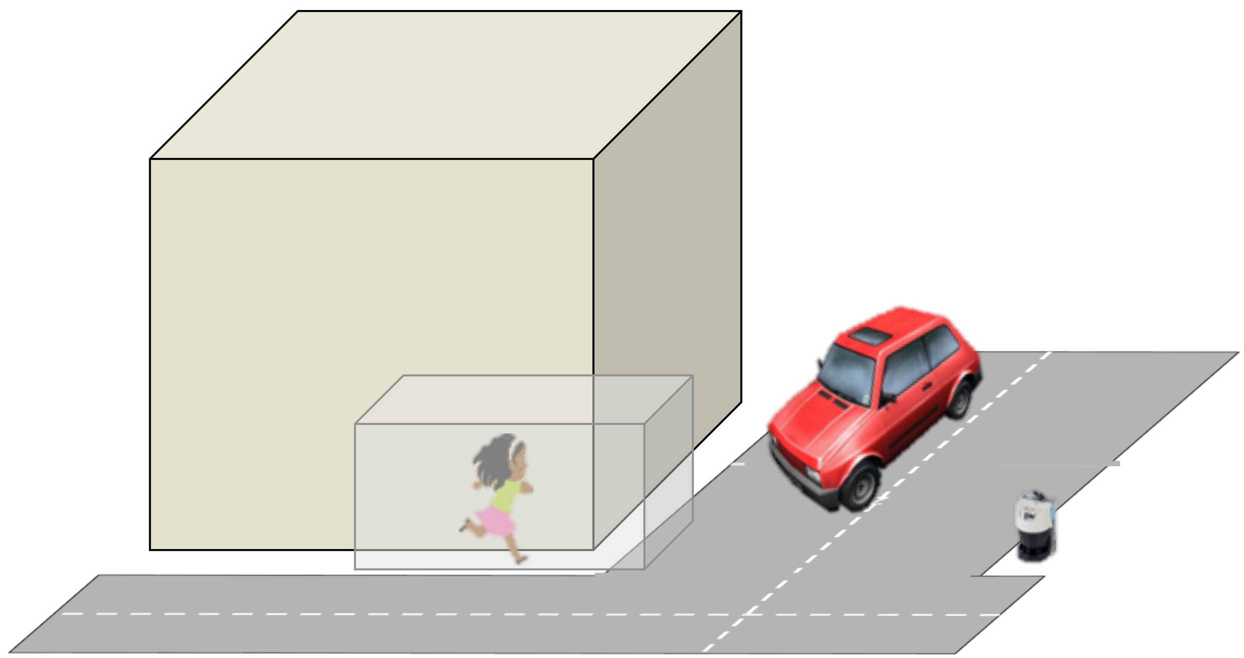

Ideally…

- Obtain low resolution view of environment
- Zoom in for higher resolution view of a smaller part of environment

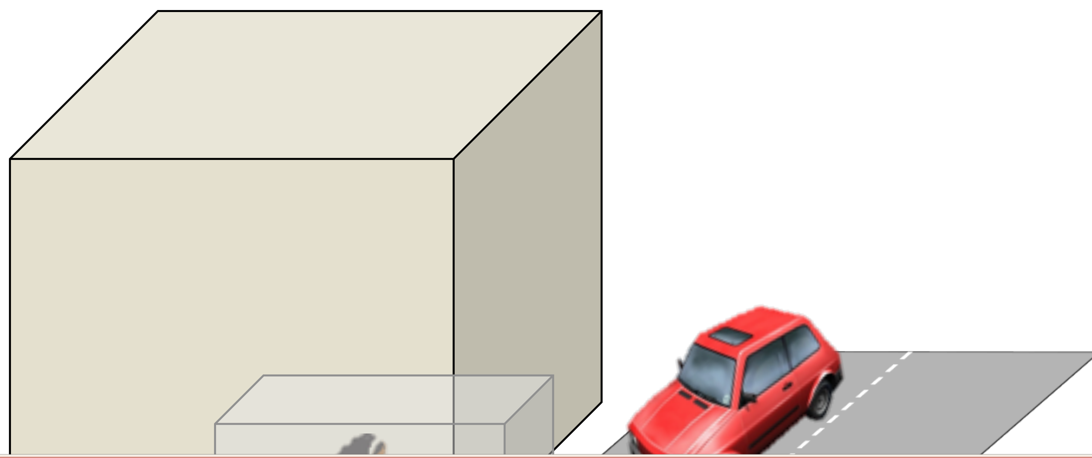

Need recursive representation that makes best use of available wireless bandwidth

Representing Content in CarSpeak

- Consider large cube encompassing environment

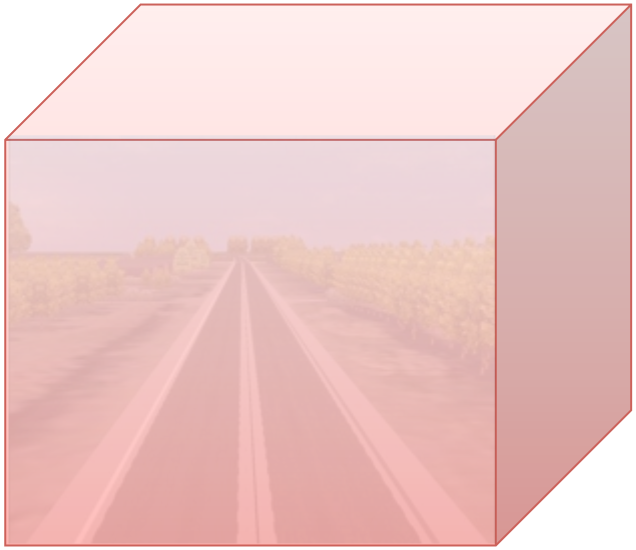

Representing Content in CarSpeak

- Consider large cube encompassing environment
- Recursively divide into 8 smaller cubes

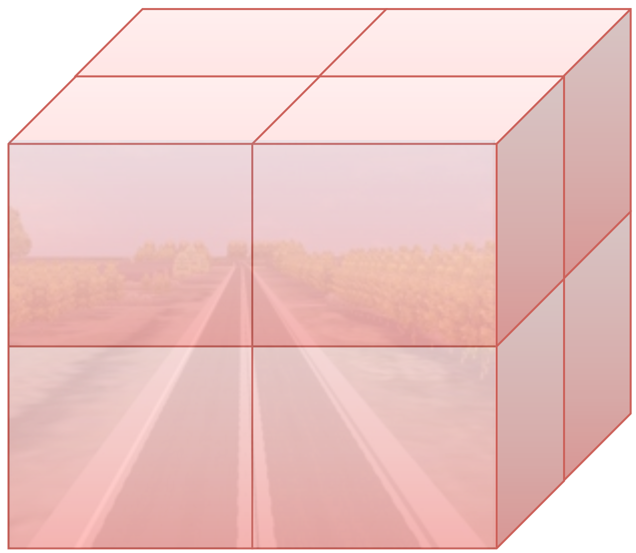

Representing Content in CarSpeak

- Consider large cube encompassing environment
- Recursively divide into 8 smaller cubes

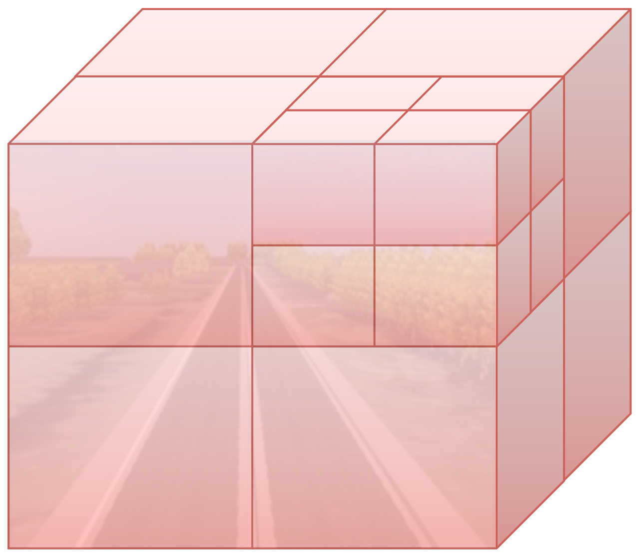

Representing Content in CarSpeak

- Consider large cube encompassing environment
- Recursively divide into 8 smaller cubes

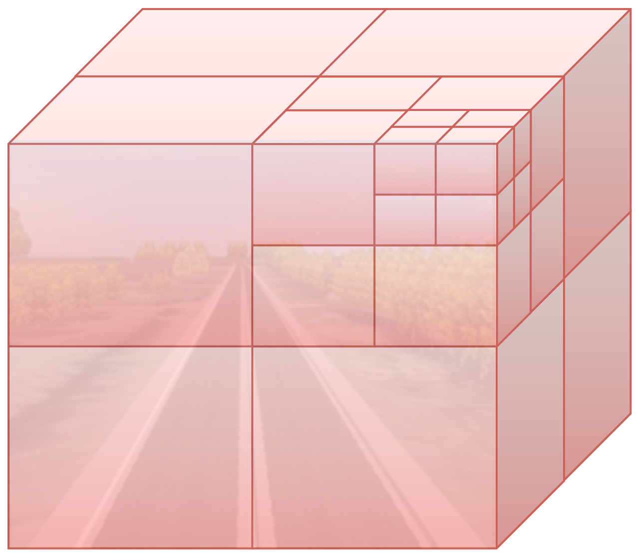

Representing Content in CarSpeak

- Consider large cube encompassing environment
- Recursively divide into 8 smaller cubes

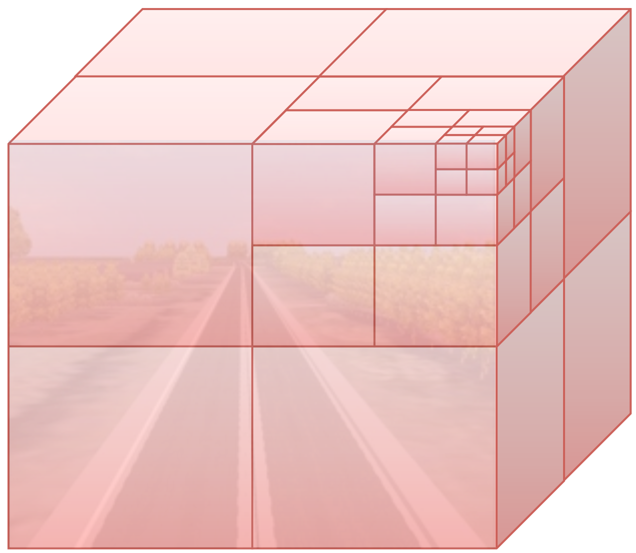

Representing Content in CarSpeak

- Consider large cube encompassing environment
- Recursively divide into 8 smaller cubes

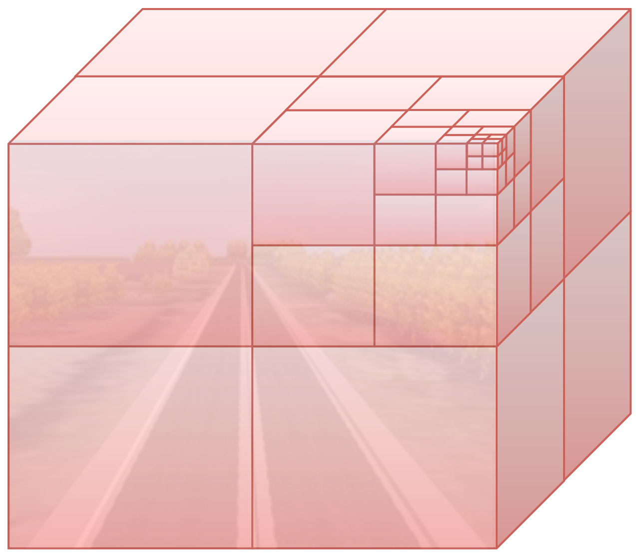

Representing Content in CarSpeak

- Consider large cube encompassing environment
- Recursively divide into 8 smaller cubes

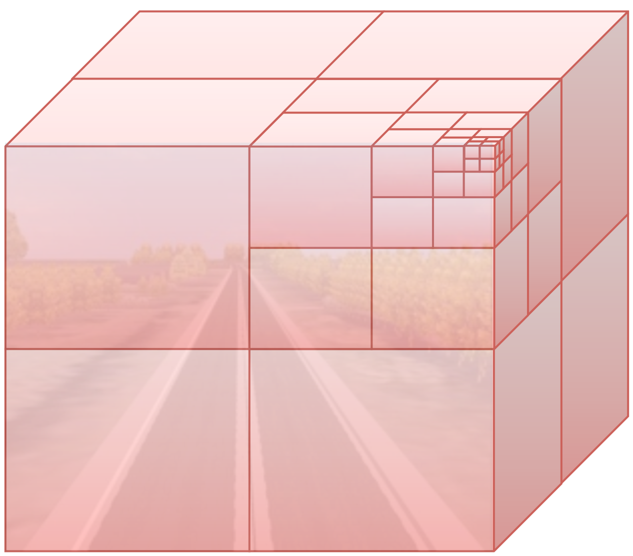

Cubes

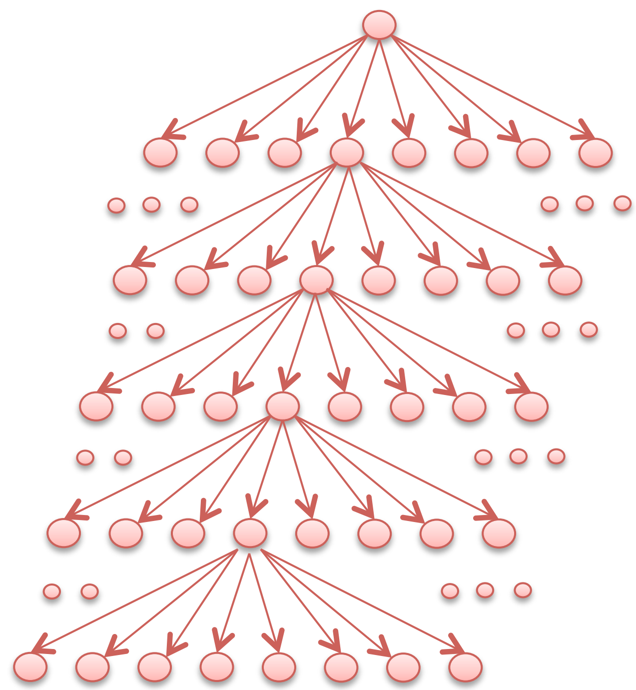

Tree Representation

Representing Content in CarSpeak

- Consider large cube encompassing environment
- Recursively divide into 8 smaller cubes

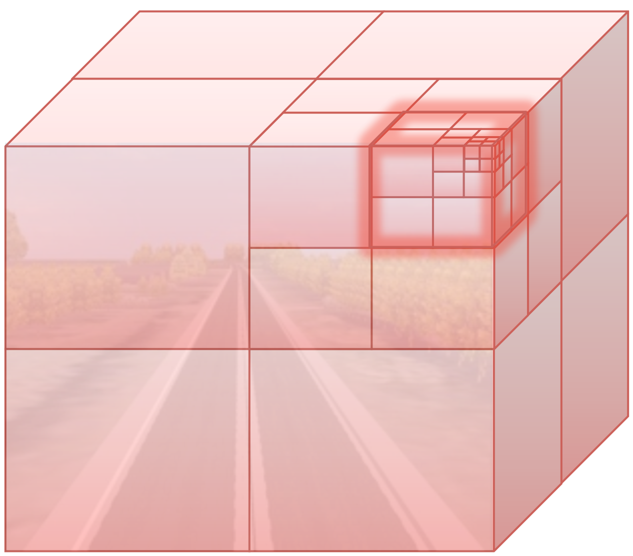

Cubes

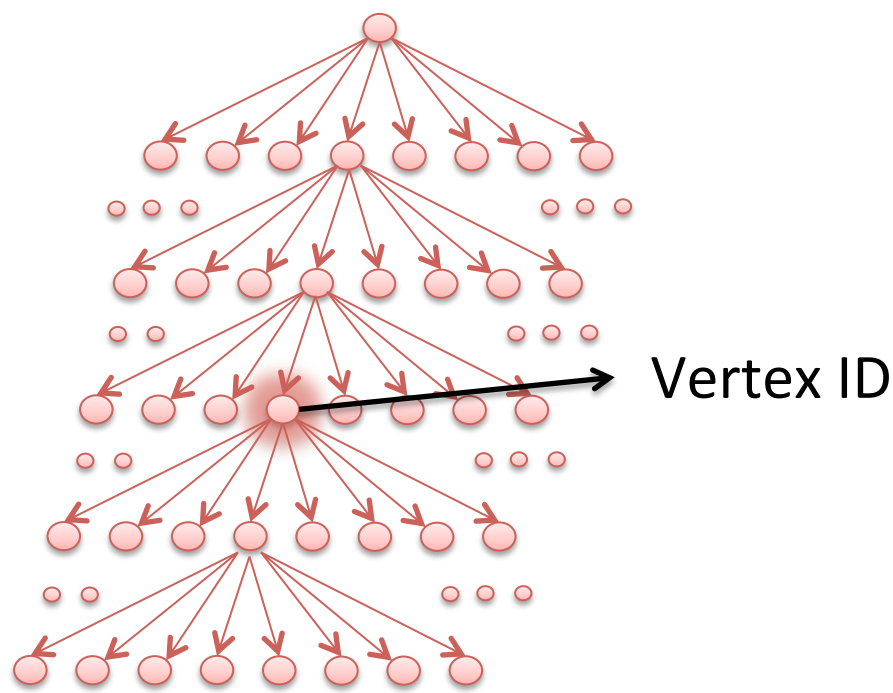

Tree Representation

Representing Content in CarSpeak

- Consider large cube encompassing environment
- Recursively divide into 8 smaller cubes
- Car needs cube at resolution (vertex ID, depth)

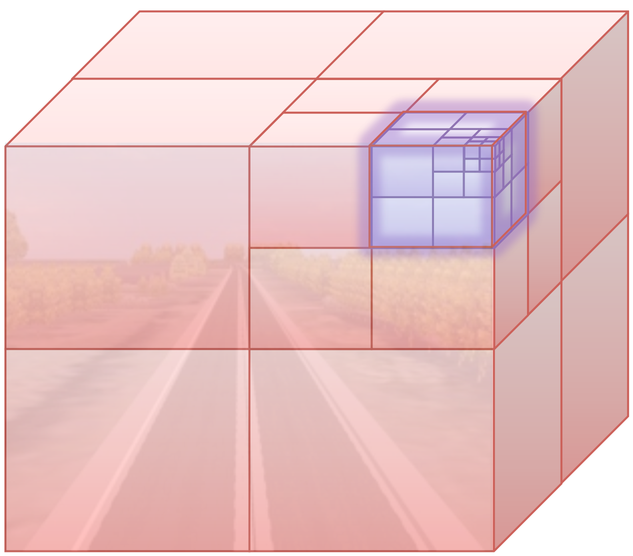

Cubes

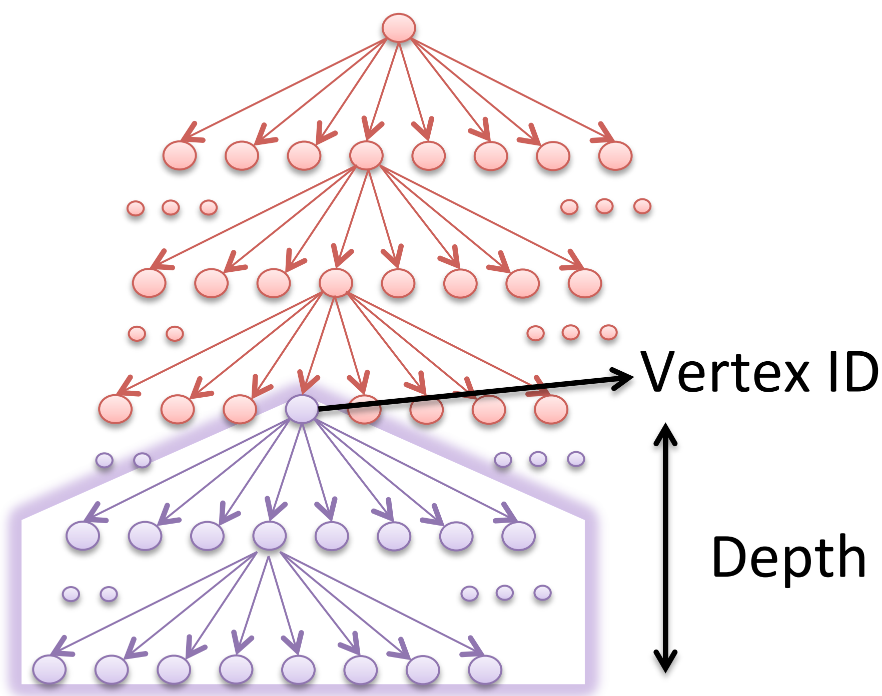

Tree Representation

What Info does Autonomous Car Need?

- Looks for obstacle free paths to destination
- Needs to know which parts of environment:
  - Are empty and safe to pass through
  - Are occupied and unsafe to pass through

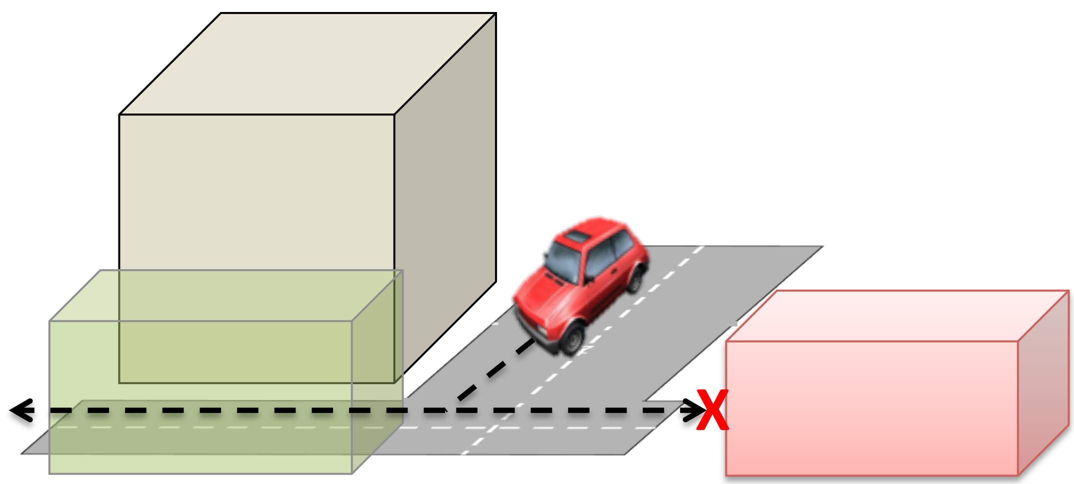

What Info does Autonomous Car Need?

- Each cube has one bit: Empty (0) or Occupied (1)
- If cube is empty

  →all cubes inside are empty

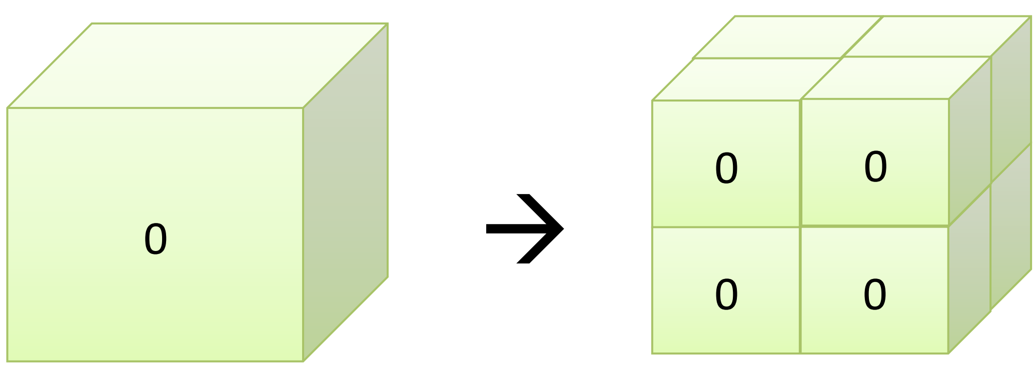

What Info does Autonomous Car Need?

- Each cube has one bit: Empty (0) or Occupied (1)
- If cube is empty
  → all cubes inside are empty
- If cube is occupied
  → at least one cube inside is occupied

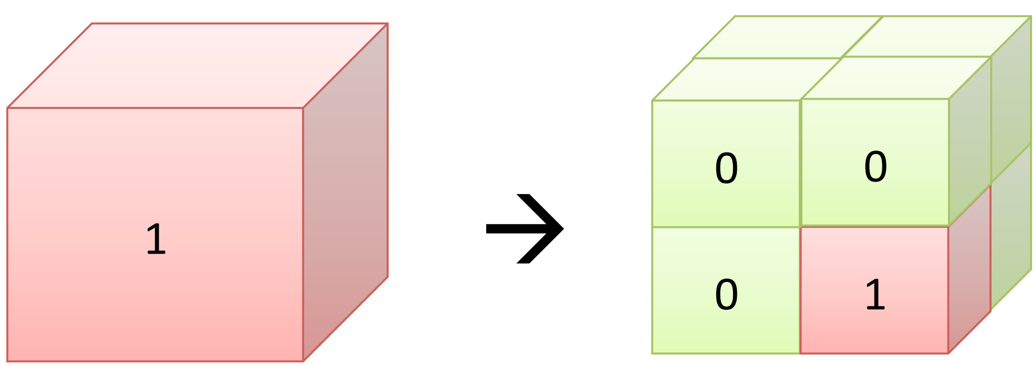

Compactly Representing Data

- Level 1 has 8 bits where 0-empty, 1-occupied

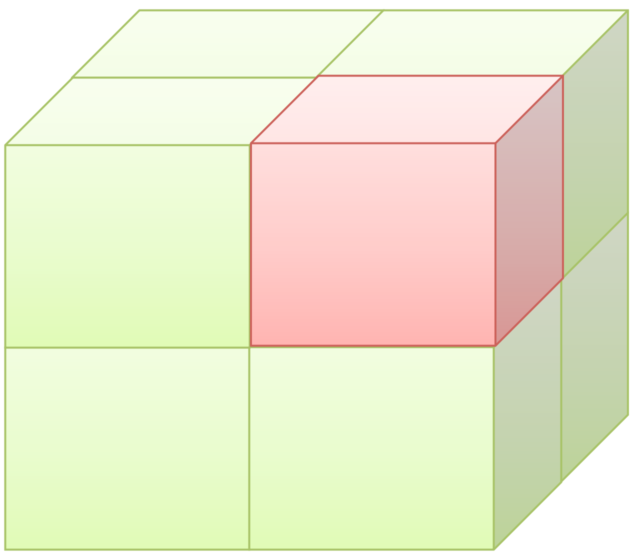

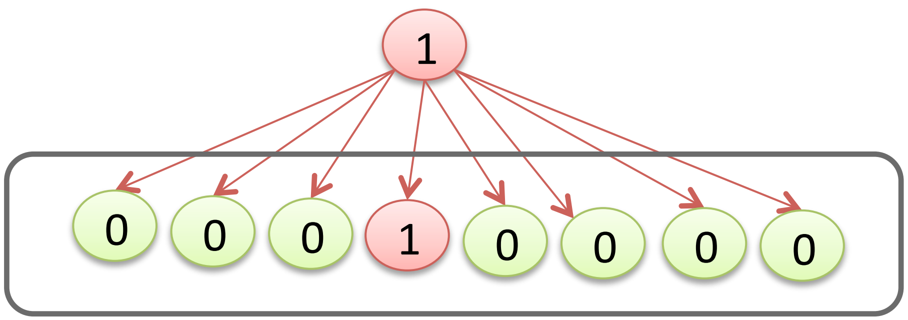

Compactly Representing Data

- Level 1 has 8 bits where 0-empty, 1-occupied
- None of 0 nodes need to be expanded

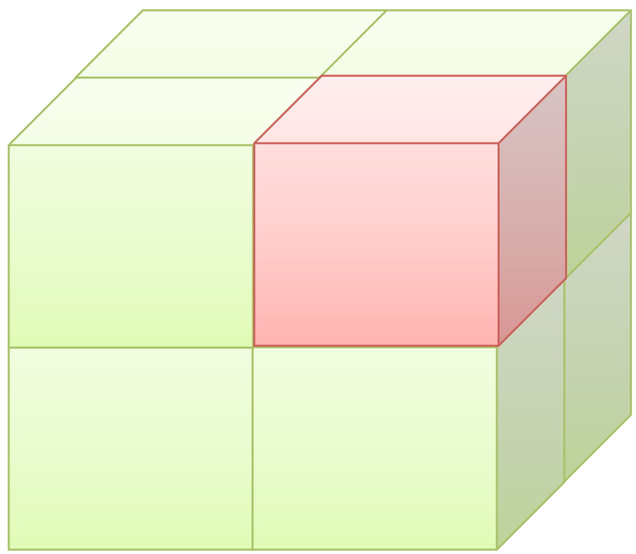

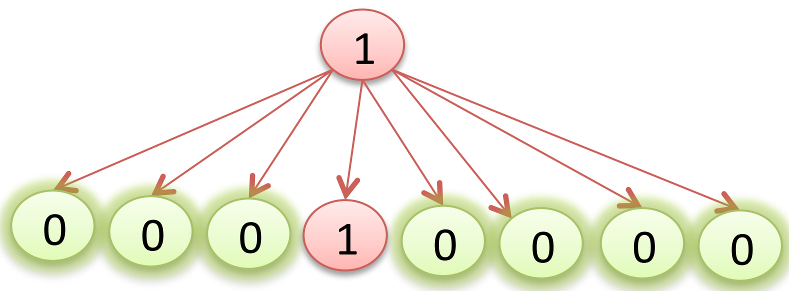

Compactly Representing Data

- Level 1 has 8 bits where 0-empty, 1-occupied
- None of 0 nodes need to be expanded
- Expand 1 node to see inside at more resolution

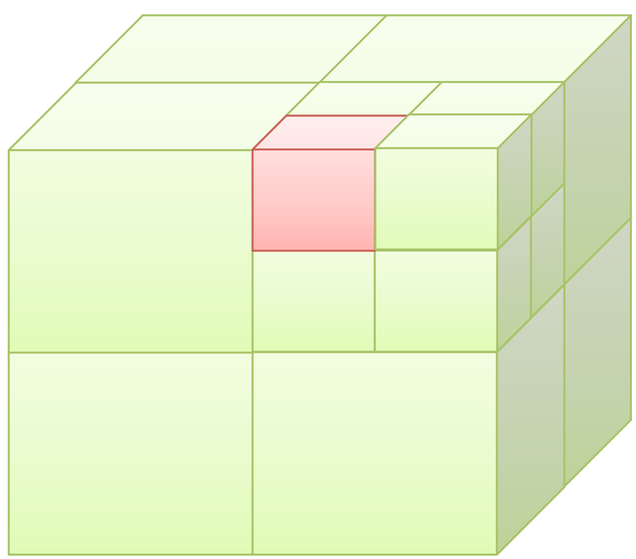

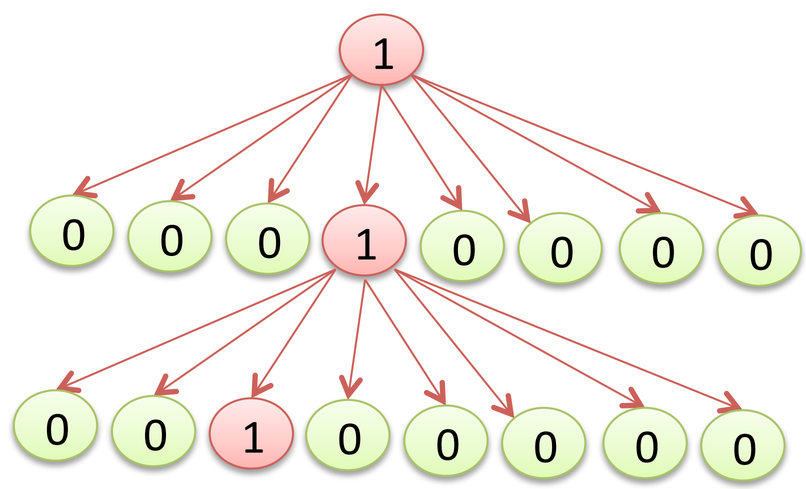

Compactly Representing Data

- Level 1 has 8 bits where 0-empty, 1-occupied
- None of 0 nodes need to be expanded
- Expand 1 node to see inside at more resolution

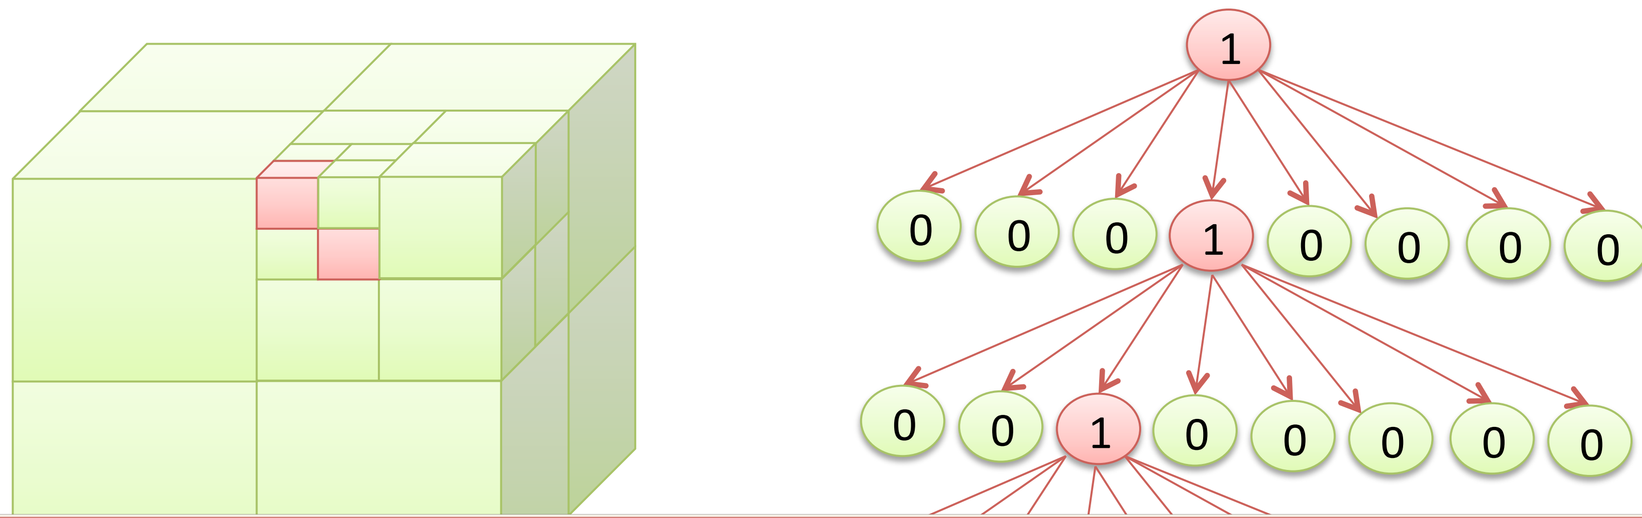

Tree representation compresses data efficiently

1. What is “content” and how do we represent it?

2. How do we disseminate this content?

How do we disseminate this content?

Q: Can we ask each car to broadcast the data it has been collected?

Autonomous cars collect huge amount of data

→ Cannot flood medium with all their data

A Request-Response Approach

A Request-Response Approach

- Car requests only data of interest
  - E.g. at blind spots, intersections, etc.
- Cars which sense the data, may respond

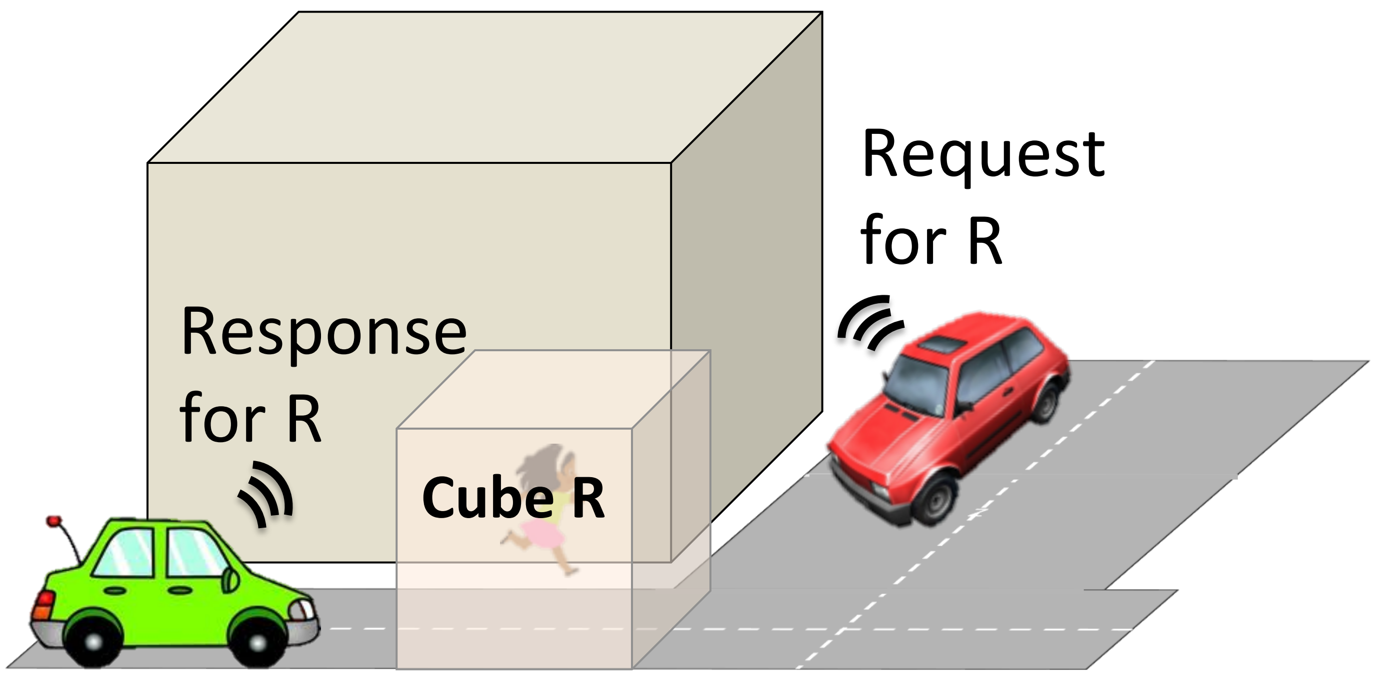

Challenge 1: Who Should Respond?

Challenge 1: Who Should Respond?

Naïve solution 1: Simply let all cars respond

Q: What is the problem of doing this?

→ A lot of redundant data

→ Congested wireless medium

Q: What are other possible solutions?

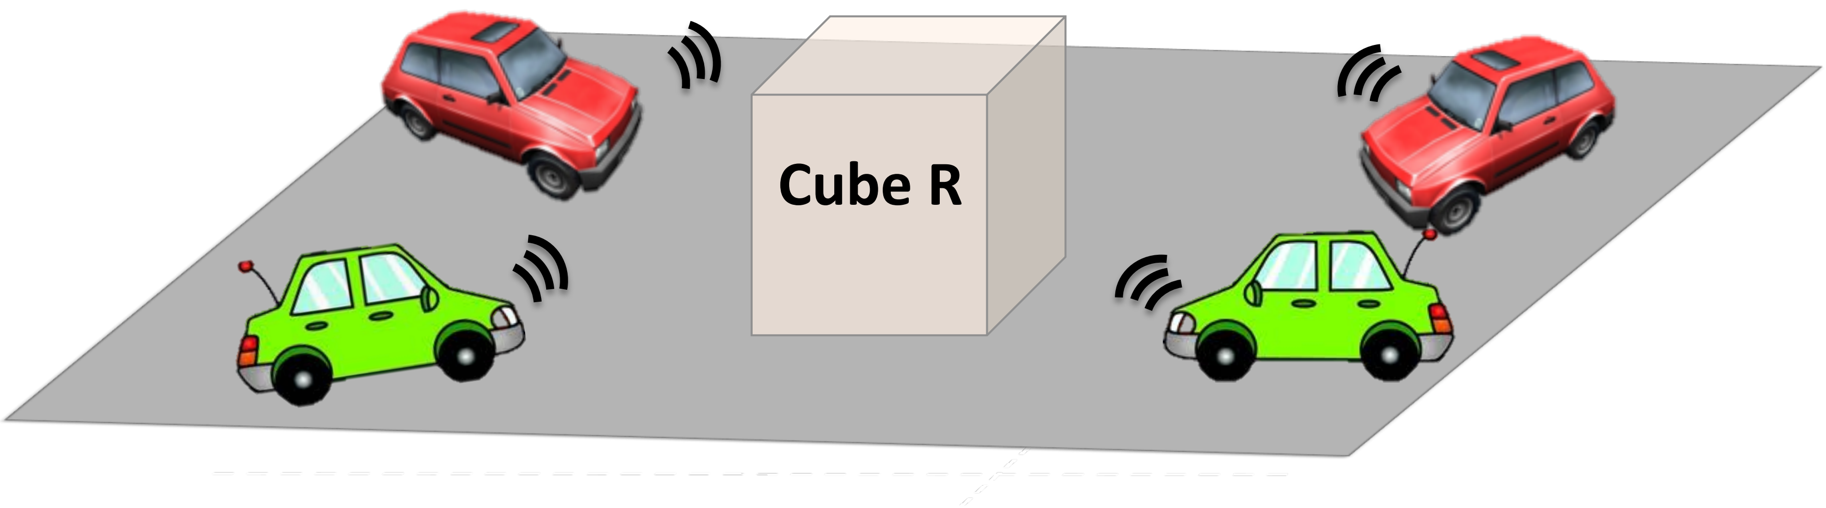

Challenge 1: Who Should Respond?

Naïve solution 2: One car respond; others who hear it suppress their response

Need to balance data diversity with data overlap

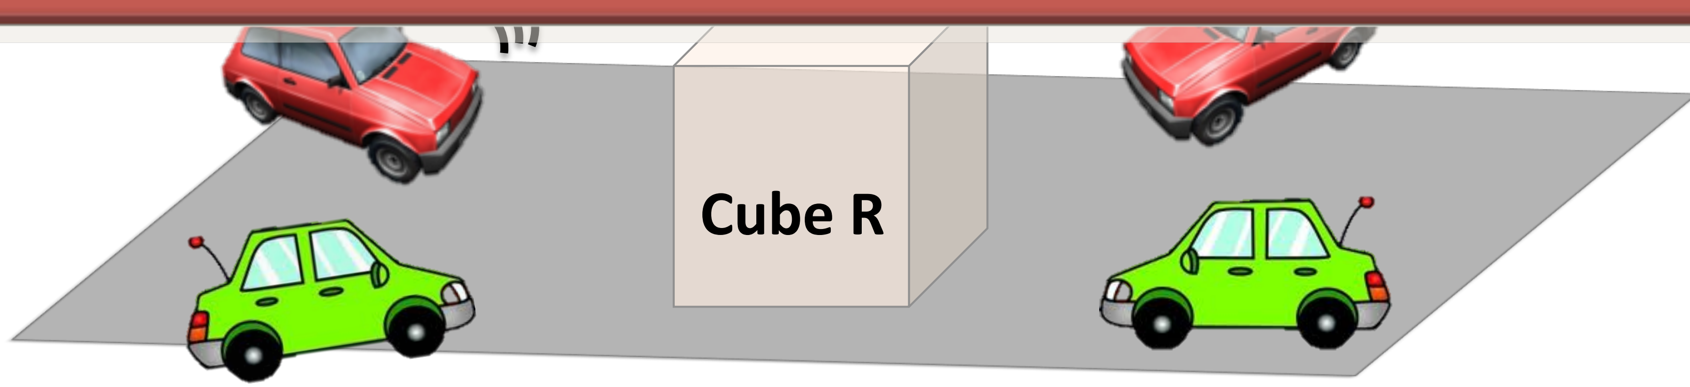

Solution: Random Walks

- Content of the cube (i.e., its subtree) is divided into packets
- Each car uses a different random walk to transmit packets

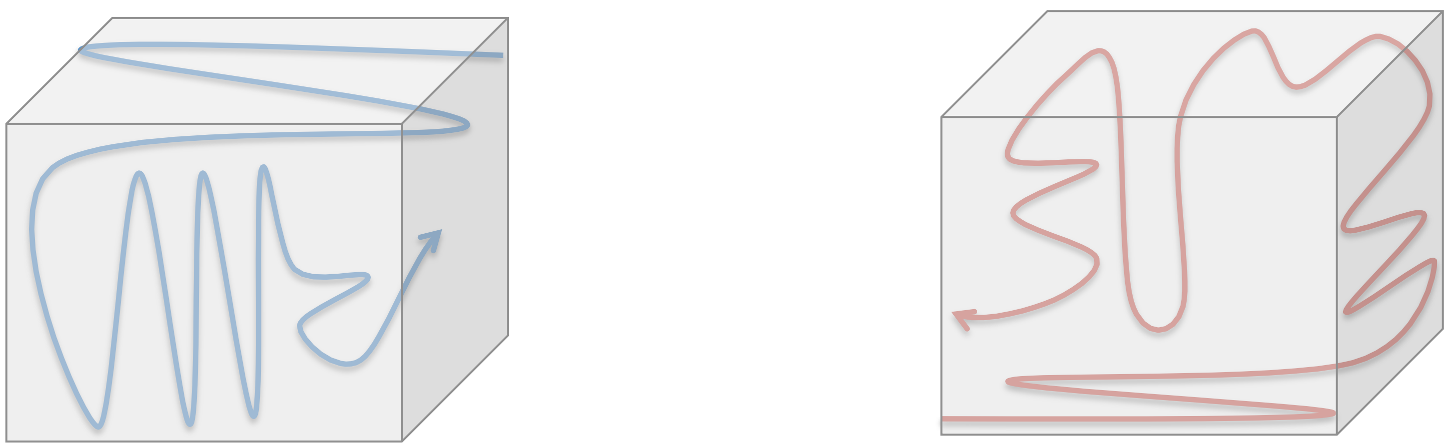

If one car transmits → Eventually finishes walk

If multiple cars transmit → Overlap is minimum

Challenge 2: Ensure medium is shared fairly by requested content

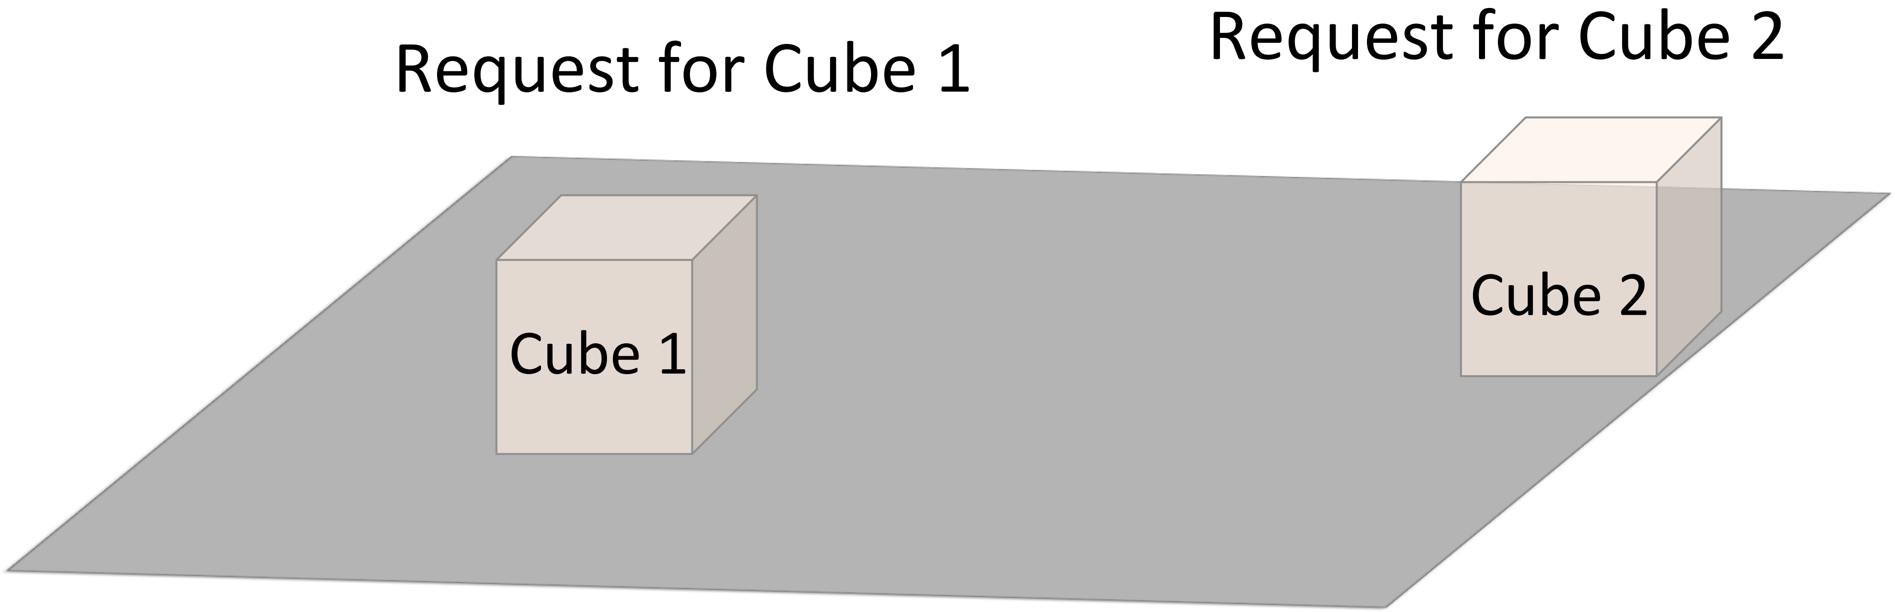

Challenge 2: Ensure medium is shared fairly by requested content

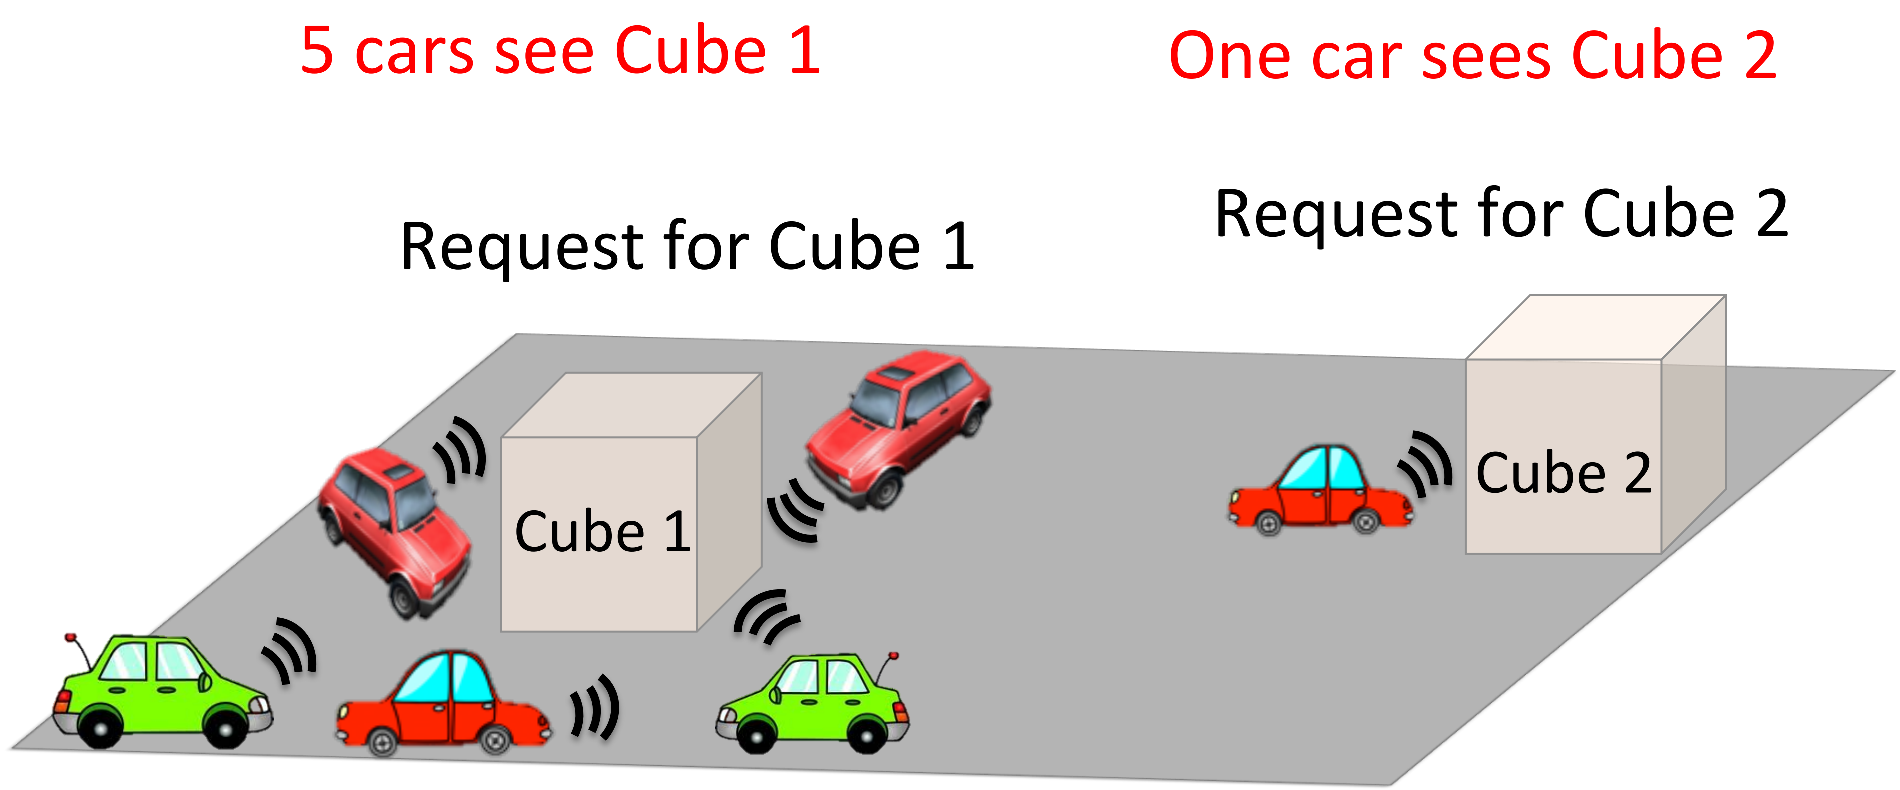

Challenge 2: Ensure medium is shared fairly by requested content

802.11 shares medium between senders
→ Cube 1's share = 5 x Cube 2's share

Ideally, we want a MAC where
→ Cube 1 share = Cube 2's share

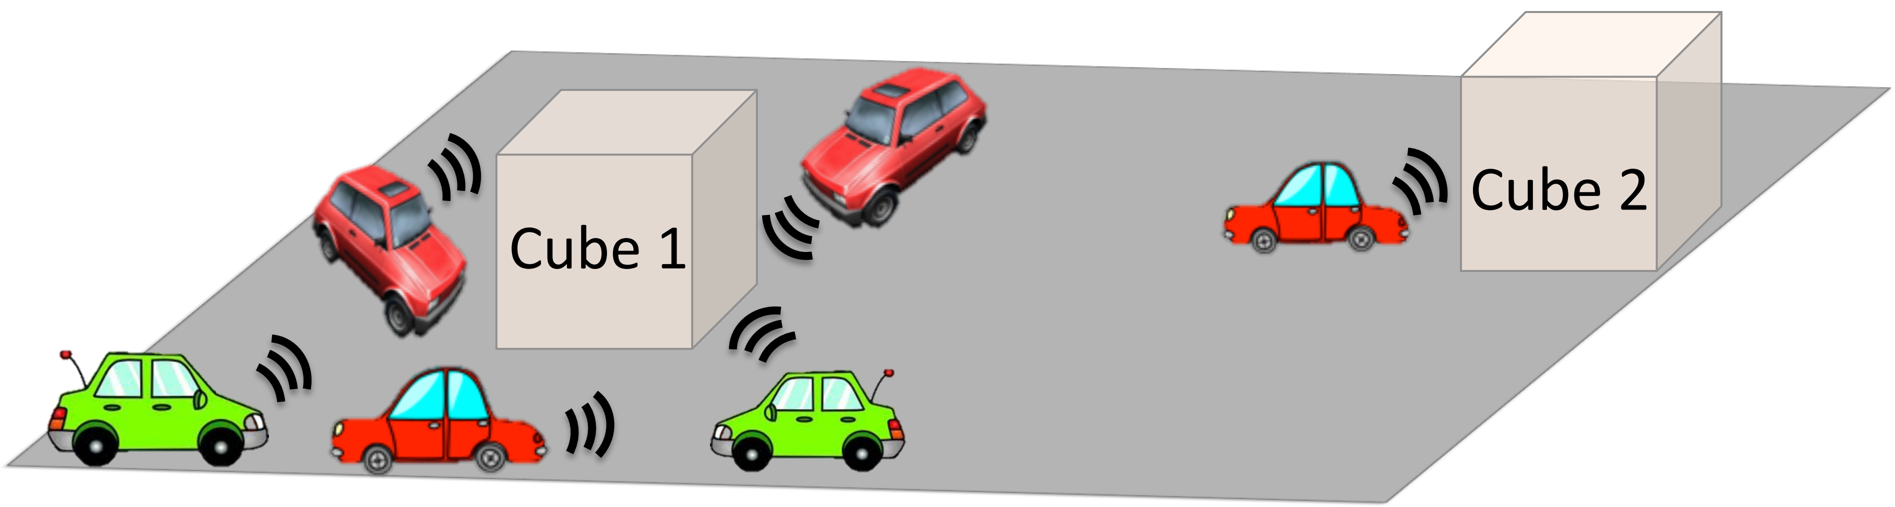

Solution: Replace sender-contention by content-contention

Ideally, we want a MAC where

→ Cube 1 share = Cube 2's share

Q: How would you design a MAC protocol that ensures the equal medium share between cubes (instead of cars)?

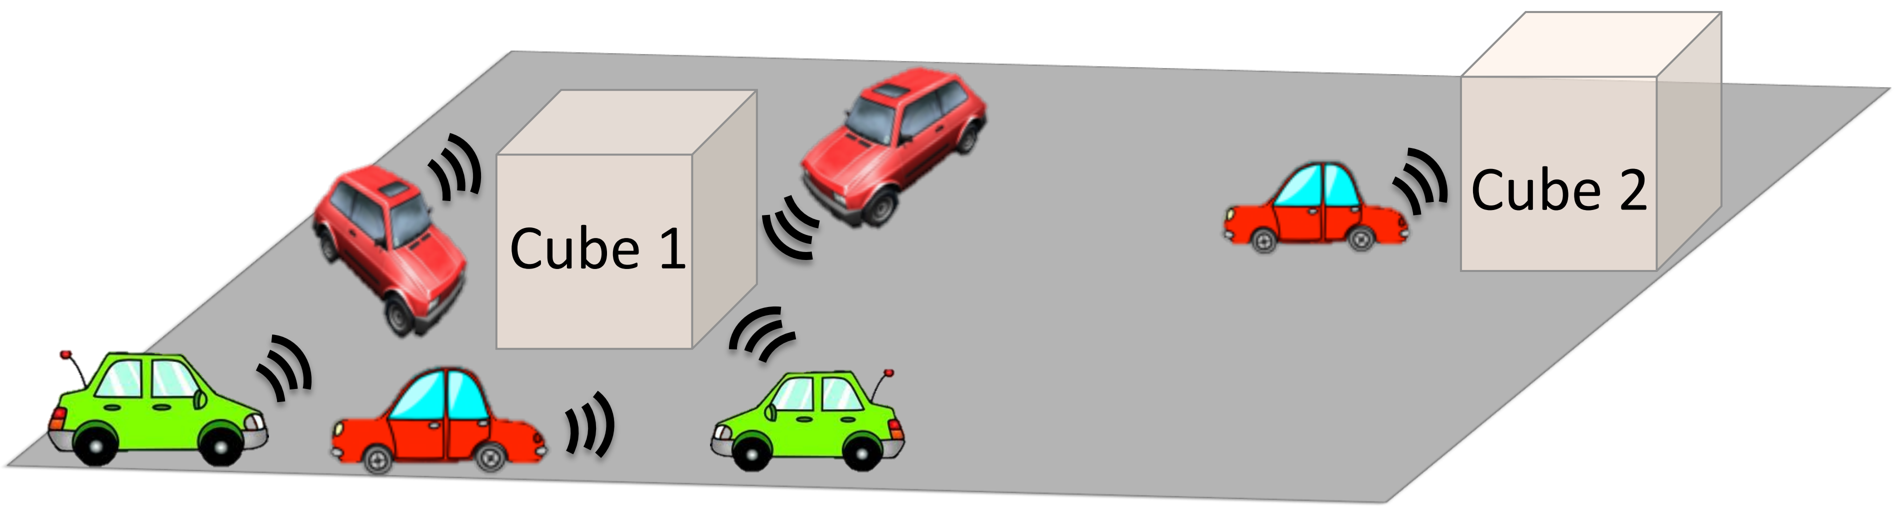

<u>Solution 2</u>: Replace sender-contention by content-contention

- Instead of senders, cubes contend for the medium
  → Requested cubes get equal share of medium

- But cubes are virtual entities
  → Cars viewing a cube contend on its behalf

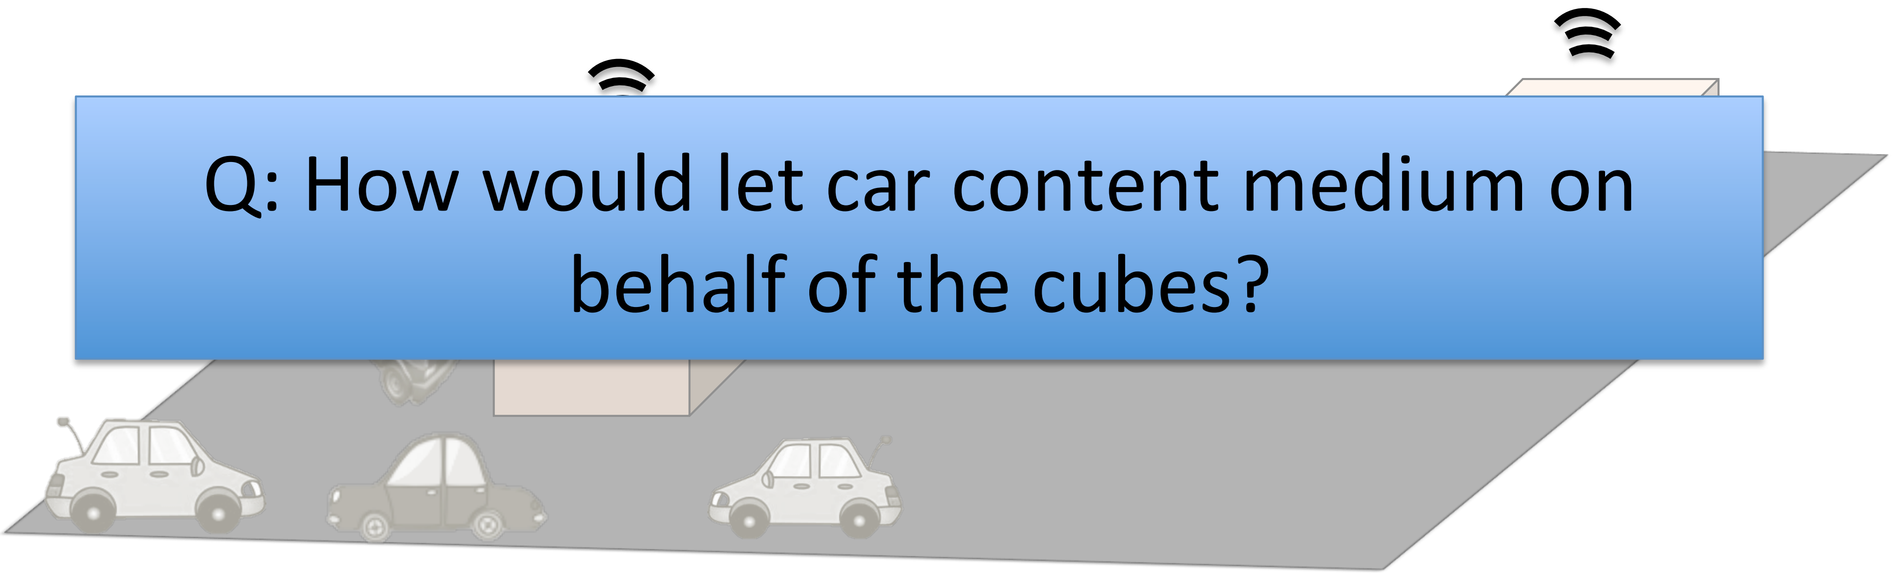

Q: How would let car content medium on behalf of the cubes?

Content-Based Contention

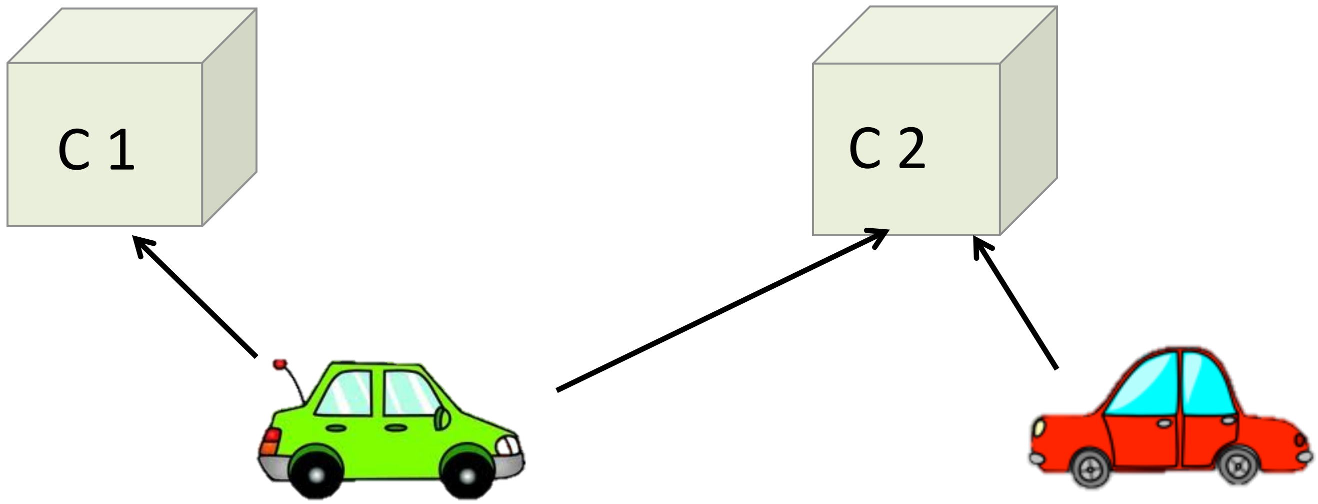

Each cube should get a share of 1/2

Green car share of the medium 3/4

red car share of the medium 1/4

Q: But how can a car compute its medium share?

Content-Based Contention

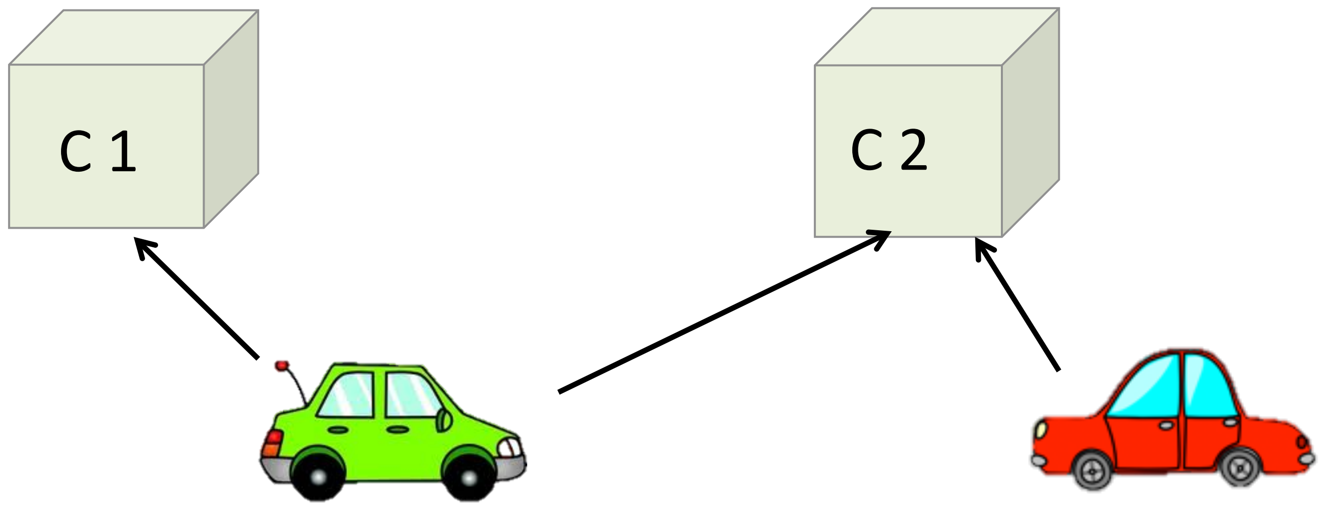

Car listens to how many cars respond for a particular cube

Share-per-cube = 1 / # cars responding

Car's total share = (Σ share-per-cube) / # requested cubes

Set contention window using car's share on off-the-shelf cards

Content-Based Contention

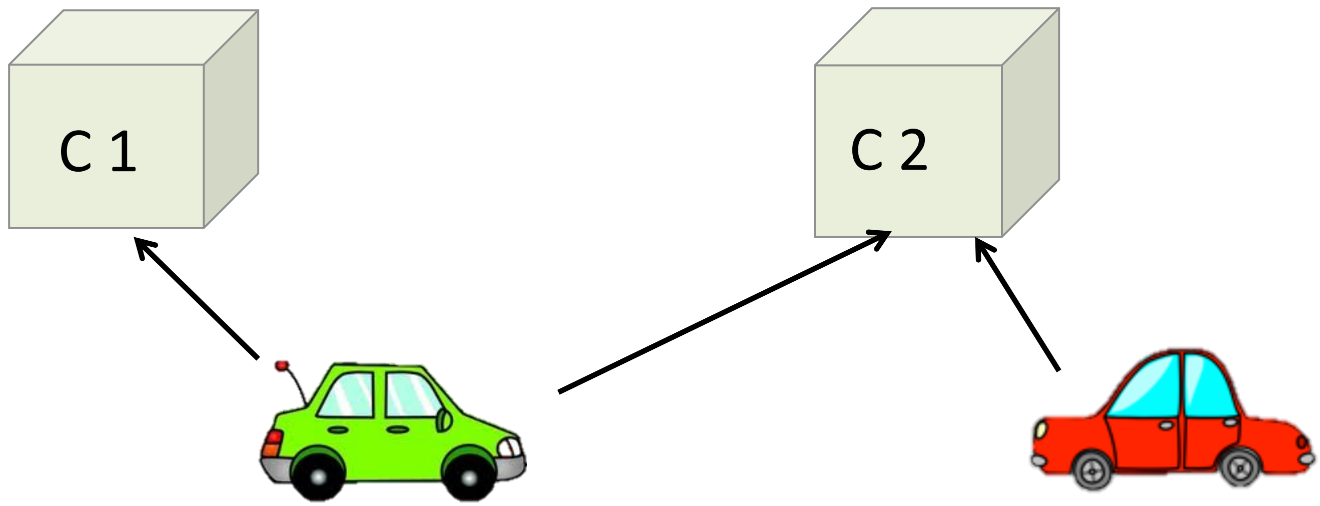

Cube 1: Share-per-cube = 1; Cube 2: Share-per-cube=1/2

Green Car's total share = (1+1/2) / 2=3/4
Red Car's total share = 1/2/2=1/4

Empirical Results

CarSpeak Implementation

- Implemented in Robot OS (ROS)
- Integrated with MIT's Path Planner from DARPA challenge
- MAC implemented in the ath9k driver for Atheros WiFi cards

Compared Schemes

- CarSpeak
- 802.11 – Request & Response

Experiments

- Indoor Testbed with Robots
- Outdoor Testbed with Autonomous Car

Indoor Testbed

- 10 Roomba robots with Kinect
- Navigate environment with obstacles

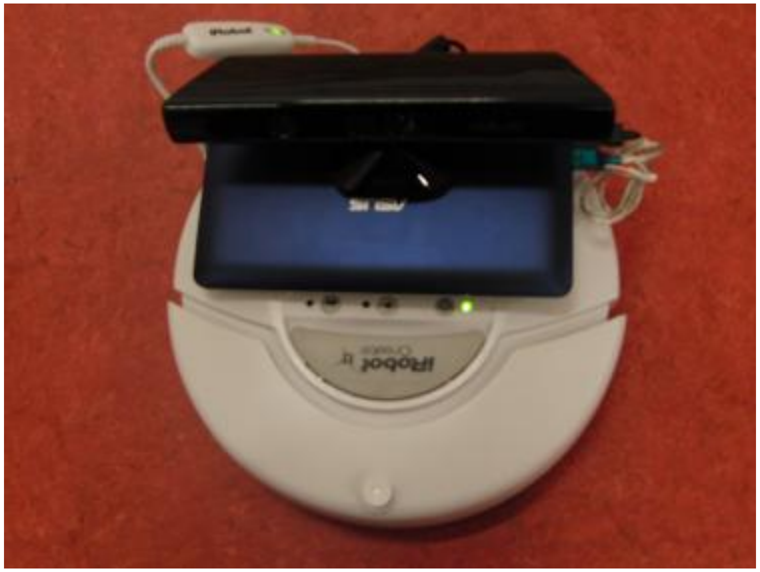

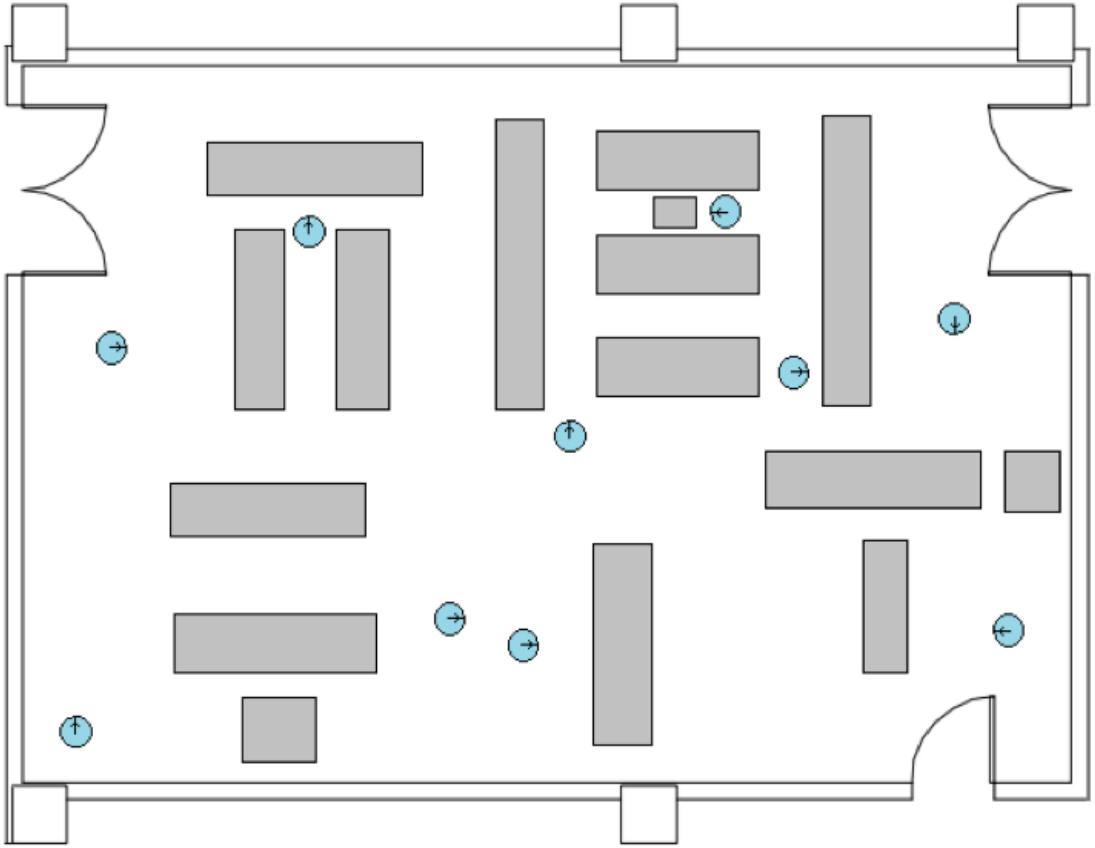

Can CarSpeak's MAC assign fair share to content?

2 requested cubes, 10 moving robots

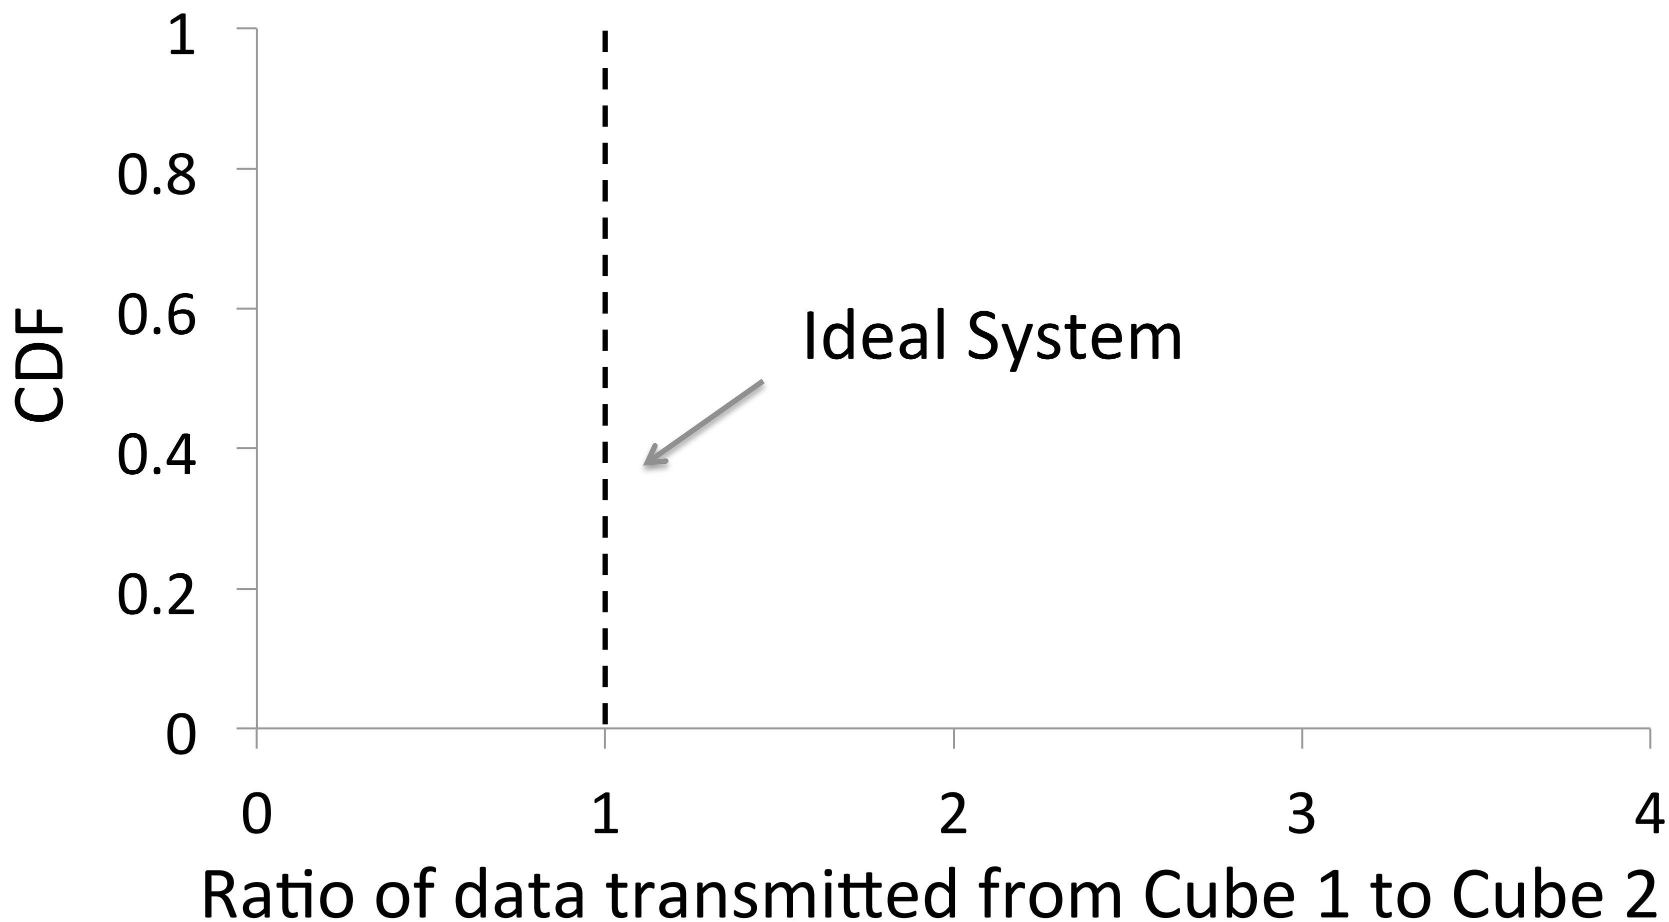

Can CarSpeak's MAC assign fair share to content?

2 requested cubes, 10 moving robots

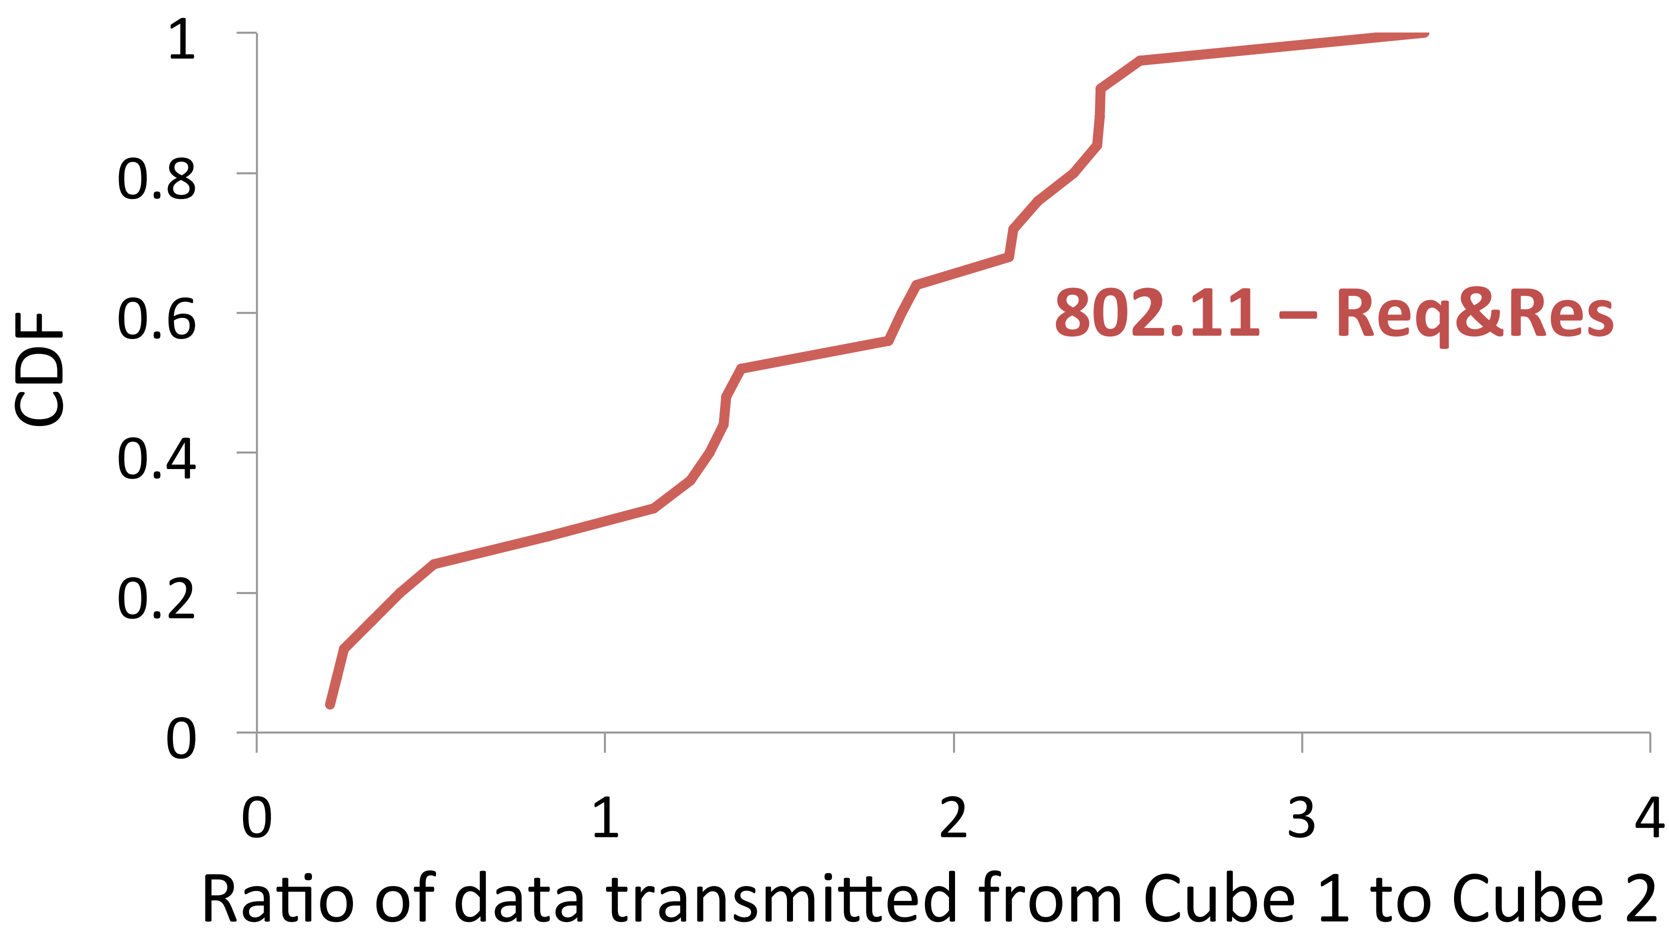

Can CarSpeak's MAC assign fair share to content?

2 requested cubes, 10 moving robots

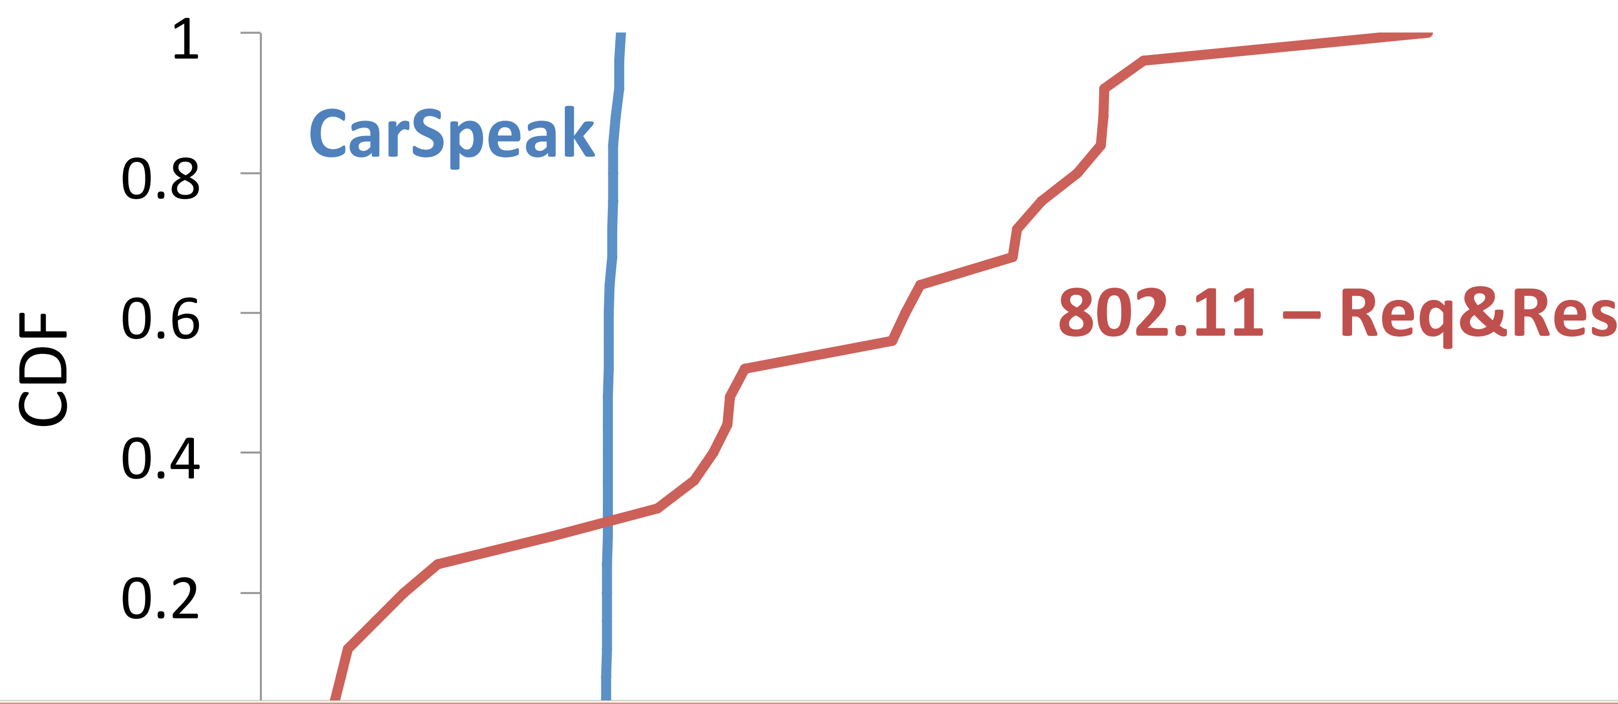

CarSpeak's content-centric MAC divides the medium fairly between requested content

Does CarSpeak Improve Reaction to Objects in Blind Spots?

A drives on road while B pulls out of occluded driveway

Timely communication can help avoid collisions

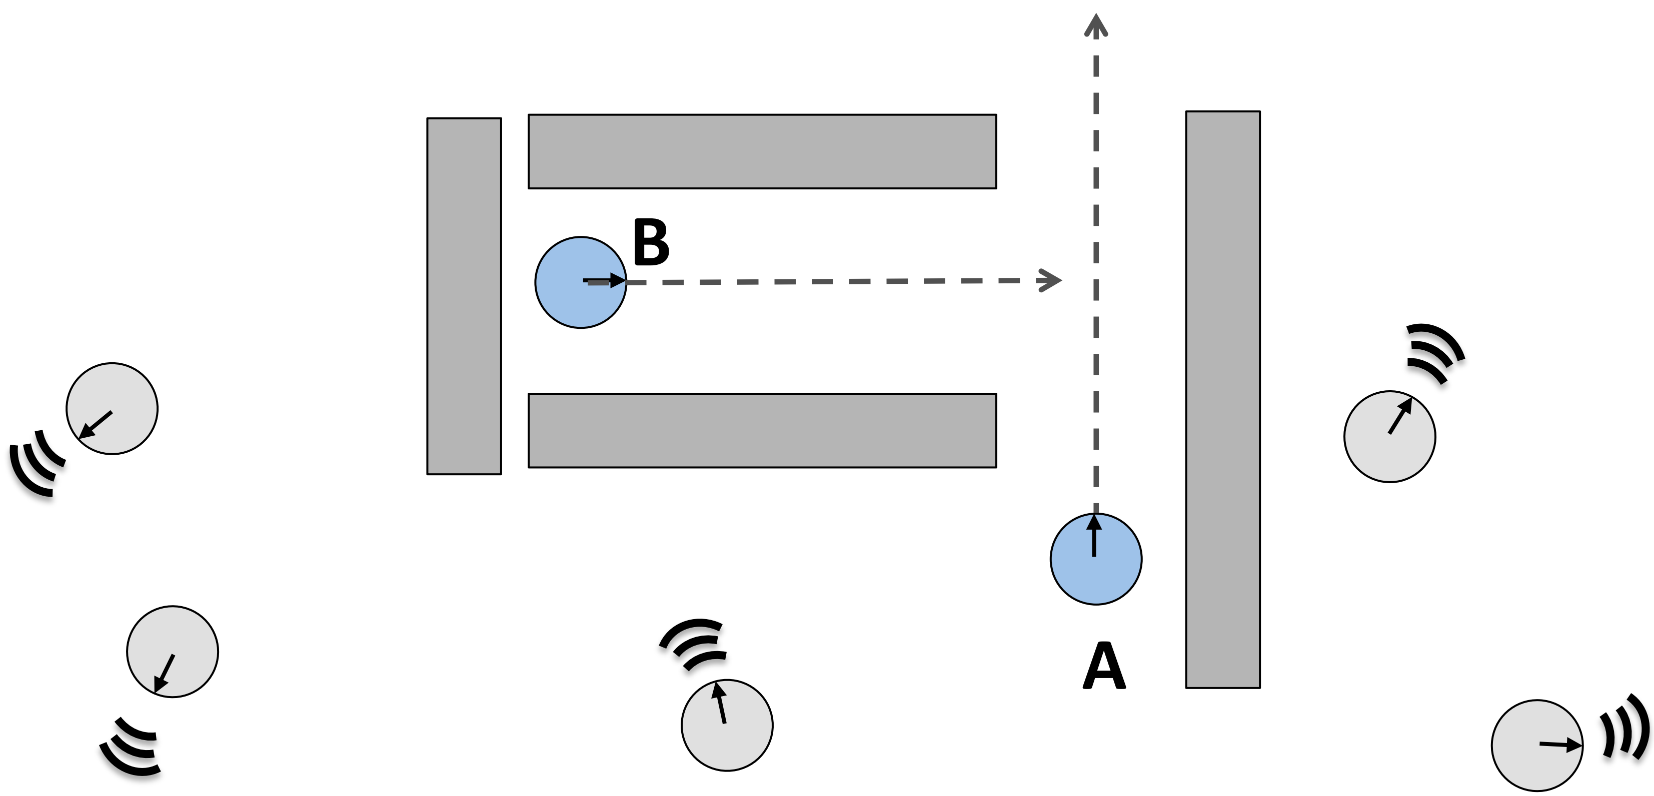

Reaction to Objects in Blind Spots

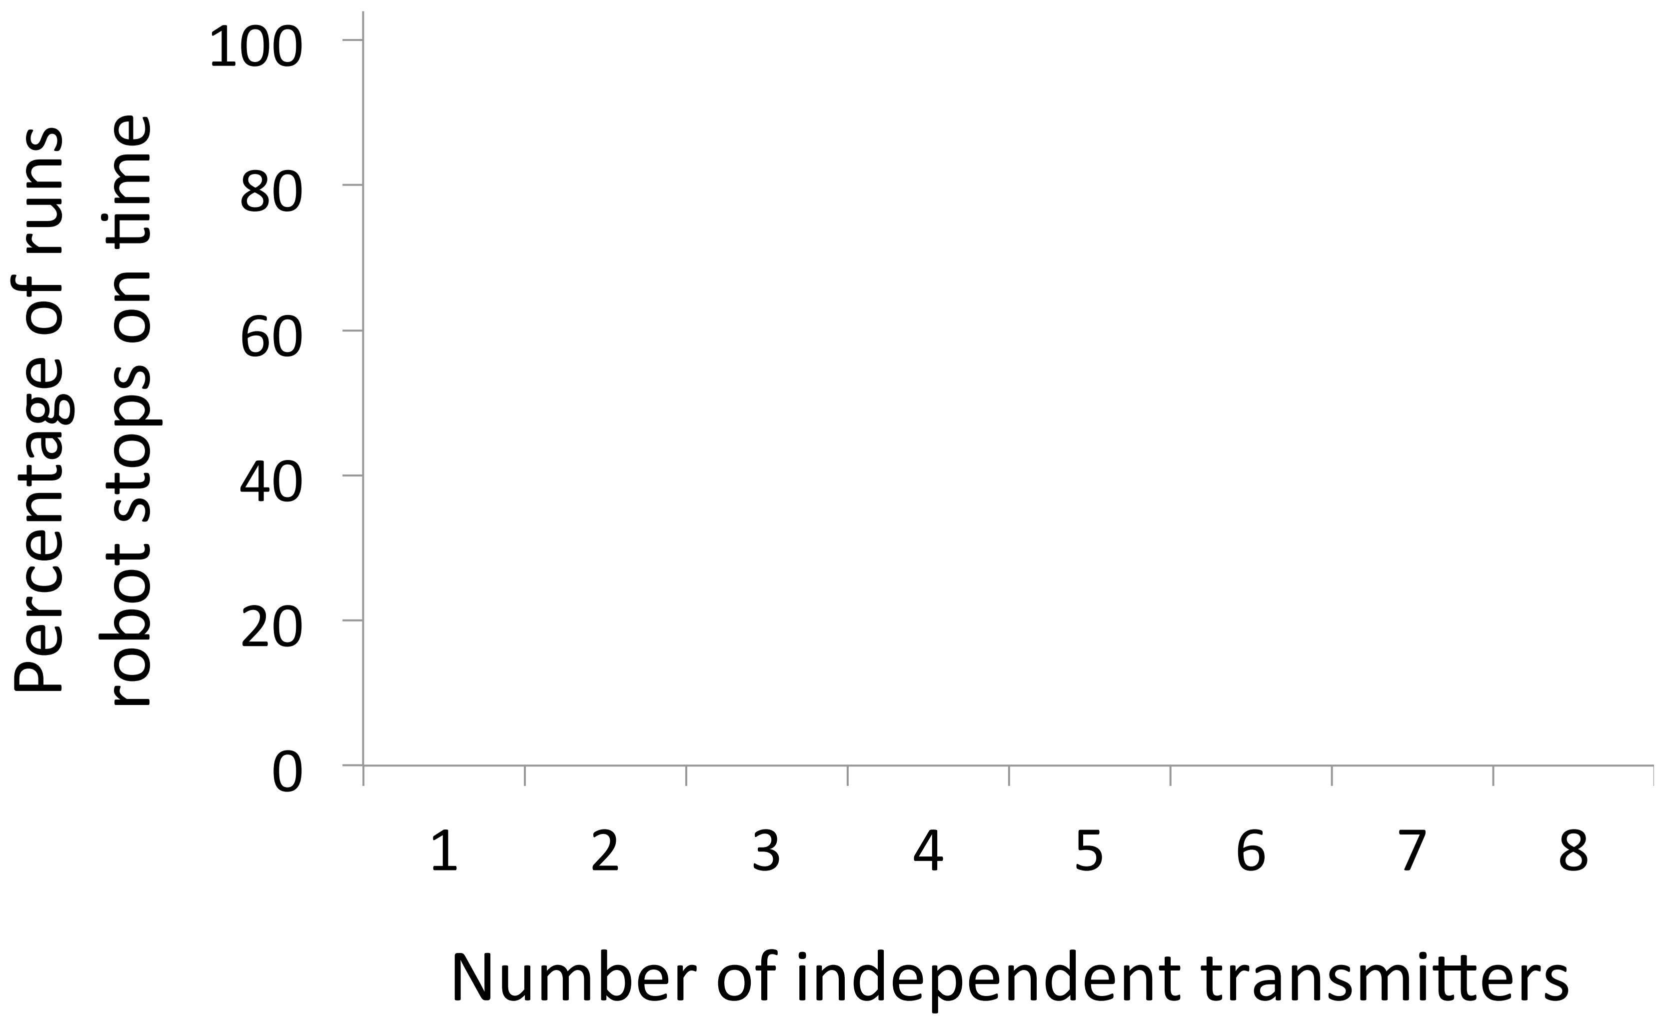

Reaction to Objects in Blind Spots

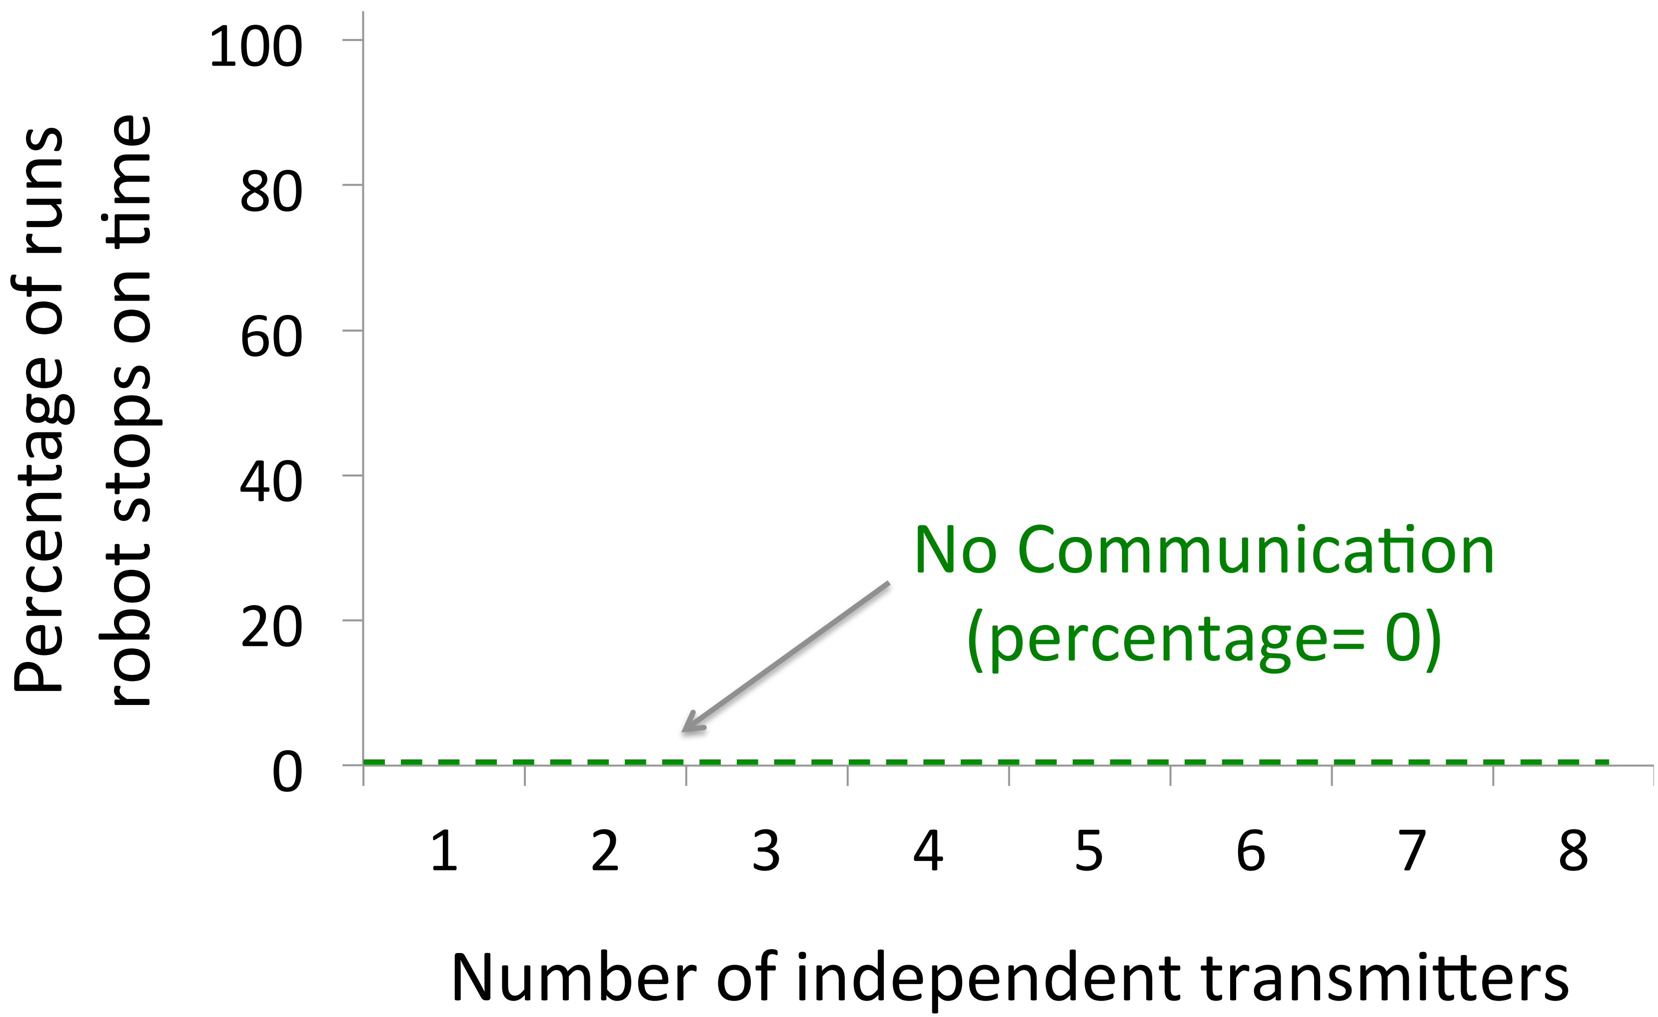

Reaction to Objects in Blind Spots

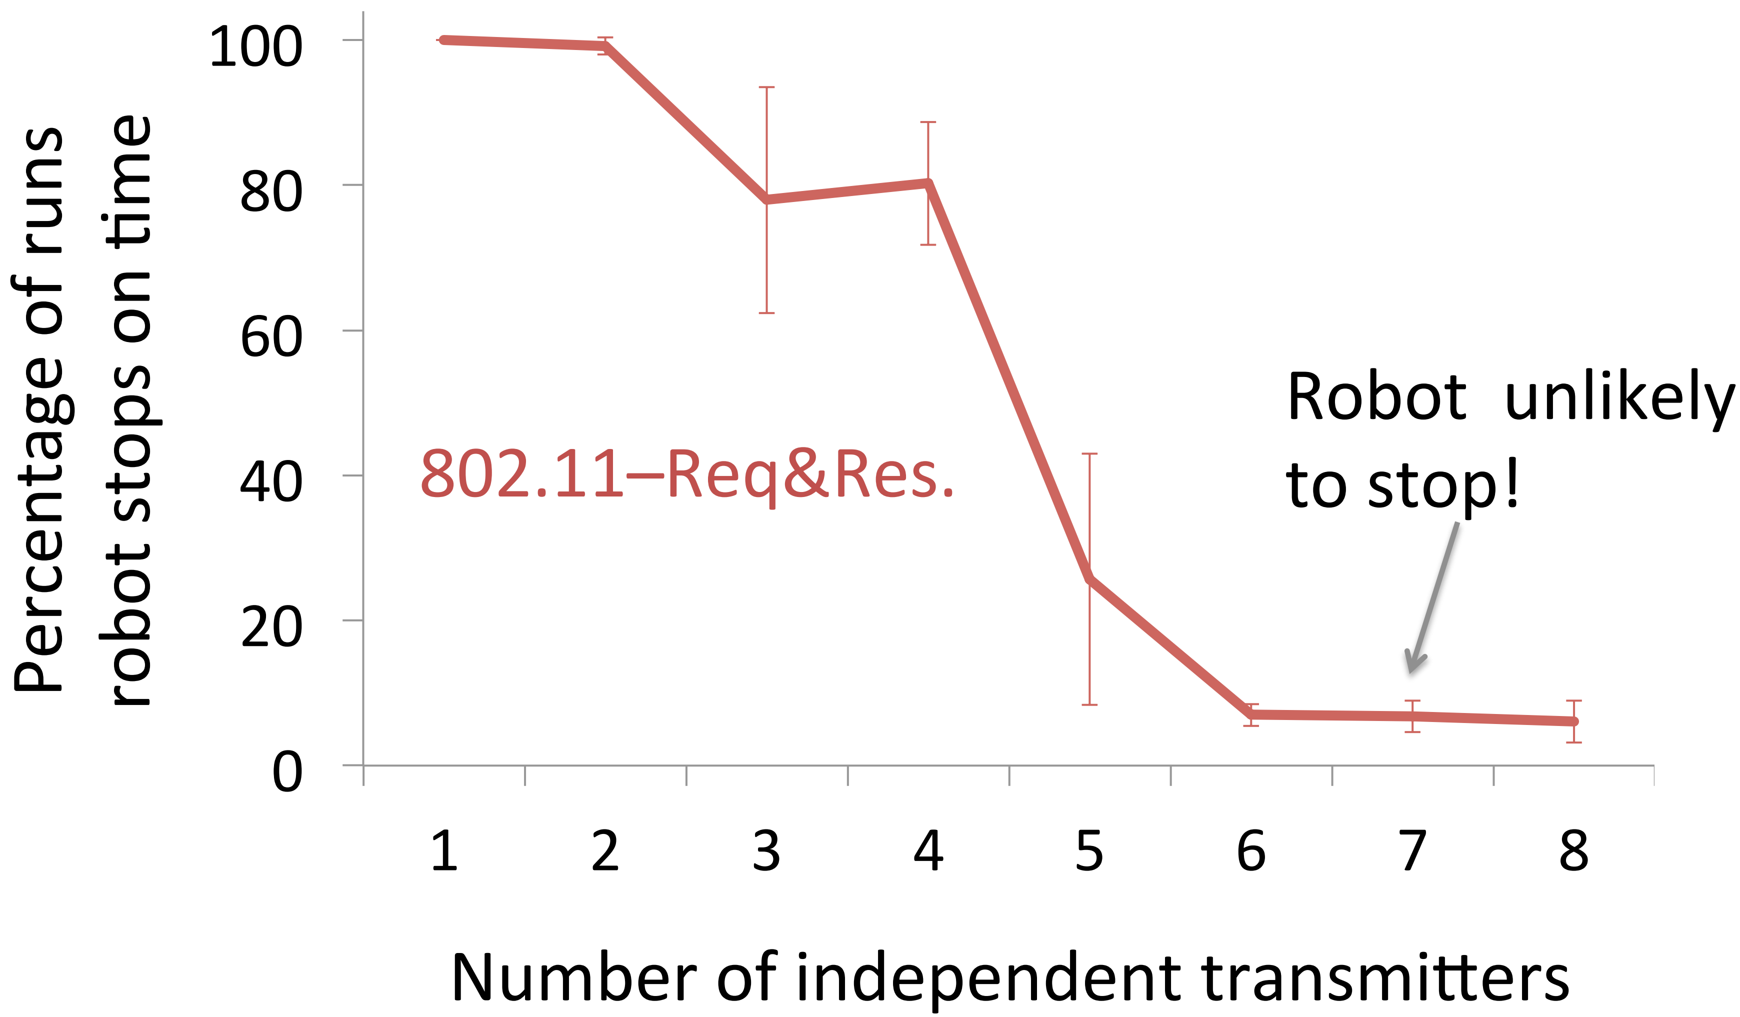

Reaction to Objects in Blind Spots

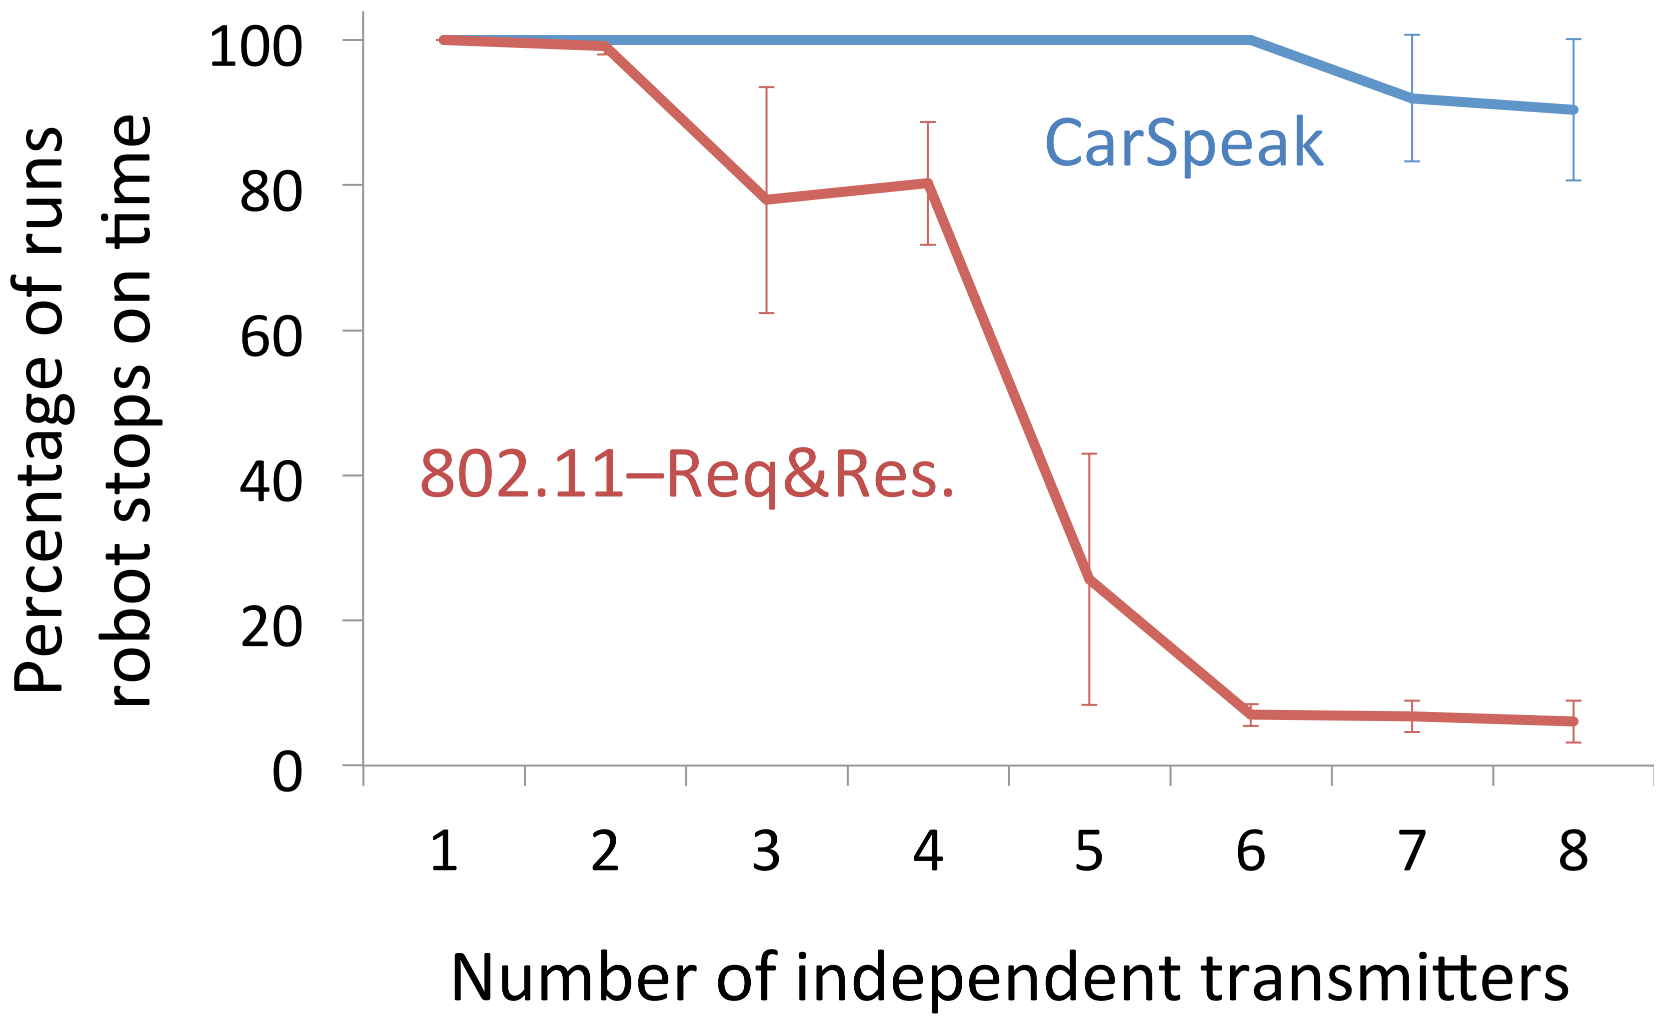

Reaction to Objects in Blind Spots

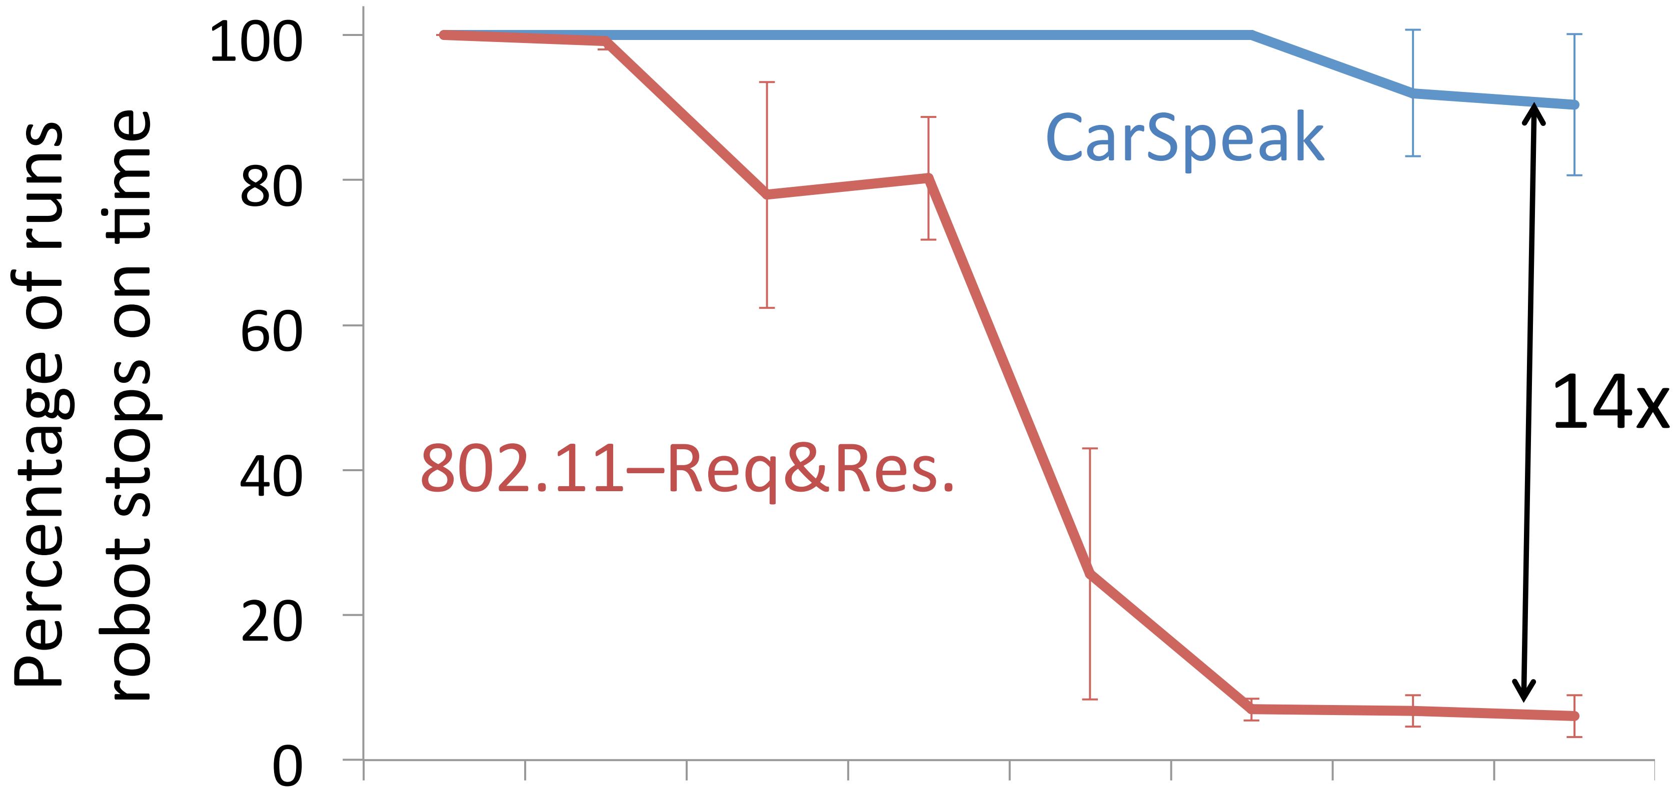

CarSpeak enables vehicles to better deal with hidden objects in blind spots

Robustness to Packet Loss

- Degrades gracefully with packet loss

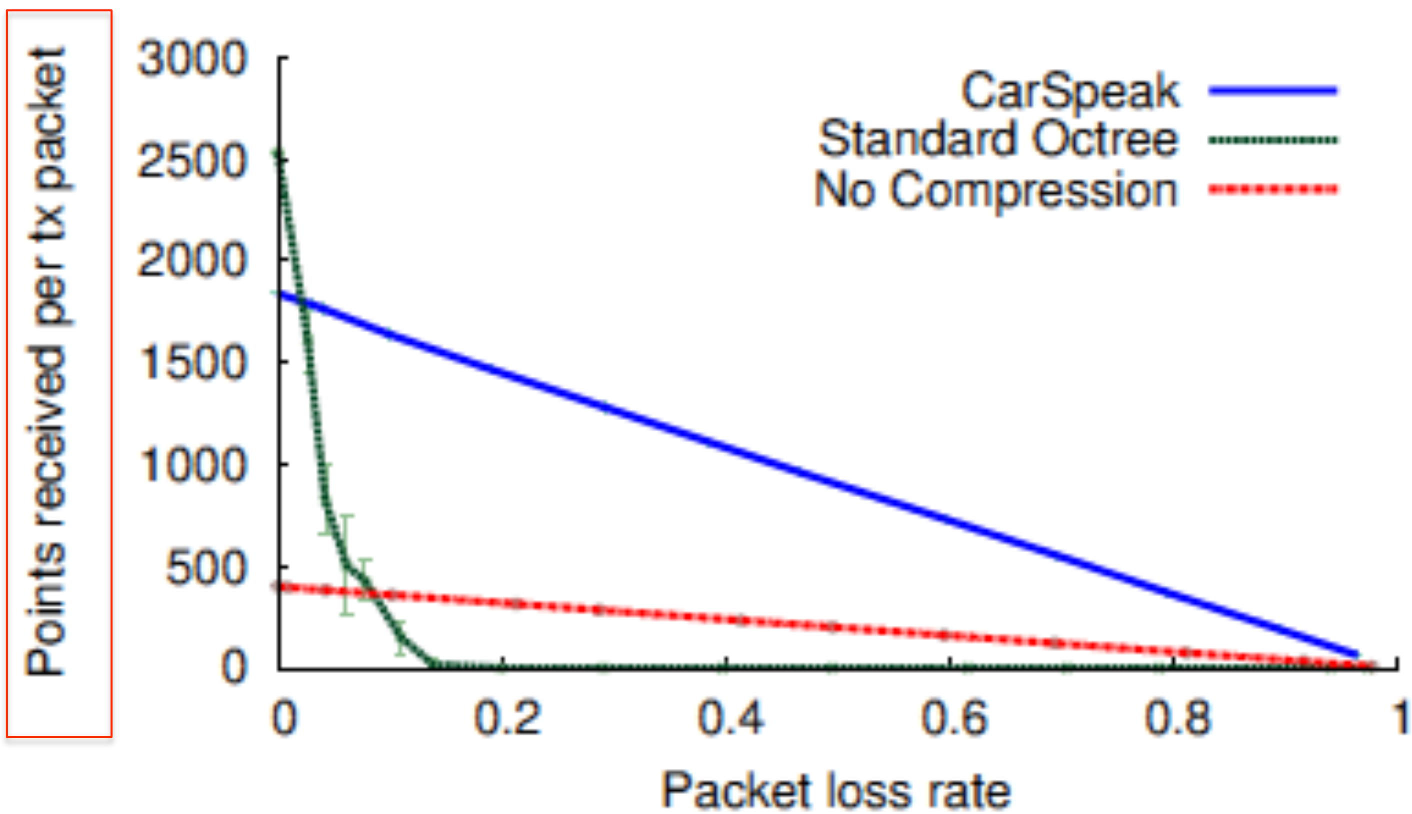

Planning Efficiency

- Better planning due to reduced losses

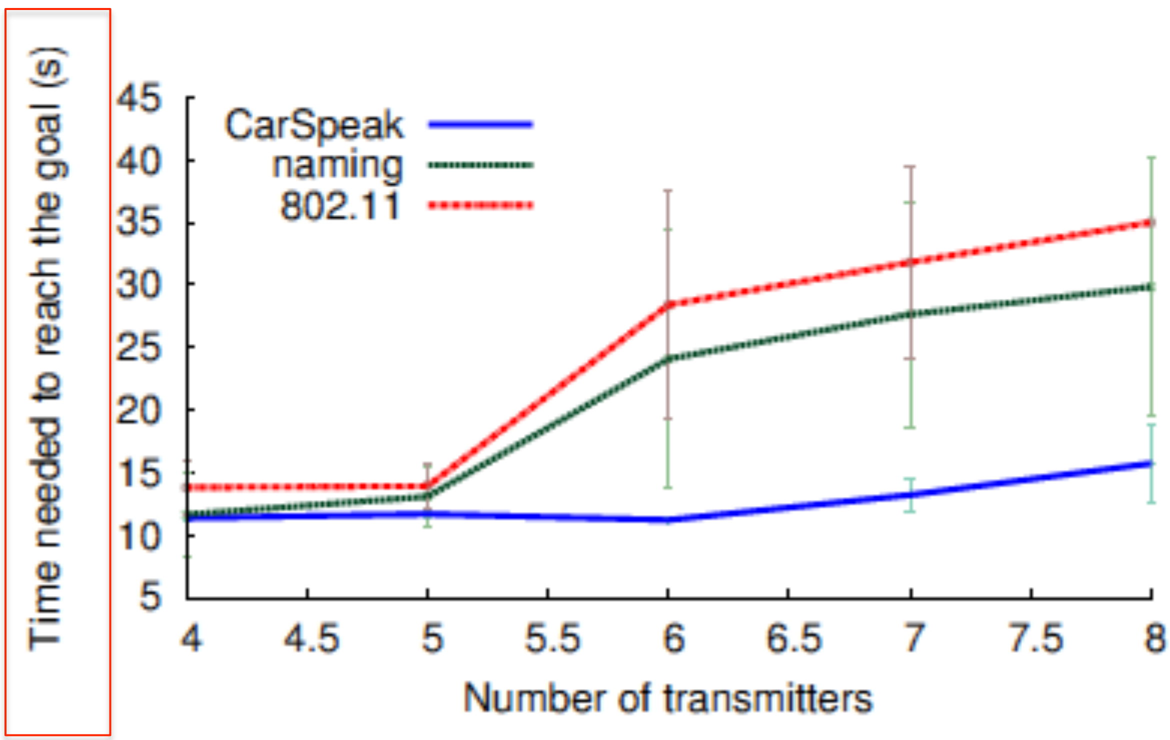

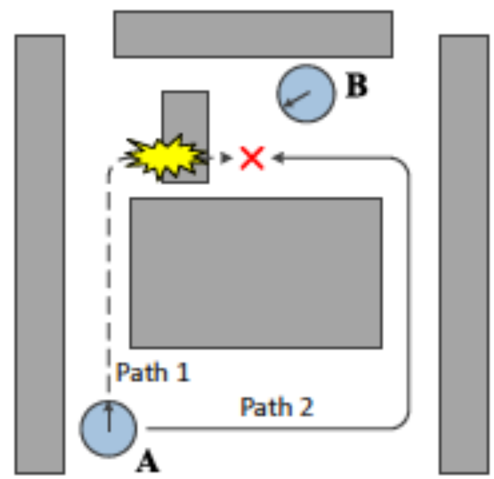

Outdoor Testbed

- Instrumented Yamaha car with laser sensors
- Pedestrian crosswalk in campus-like environment

Detecting a Pedestrian in Blind Spot

- Pedestrians emerge from lobby
- Lobby is a blind spot
- Infrastructure sensors, some can send view of lobby

Outdoor Results

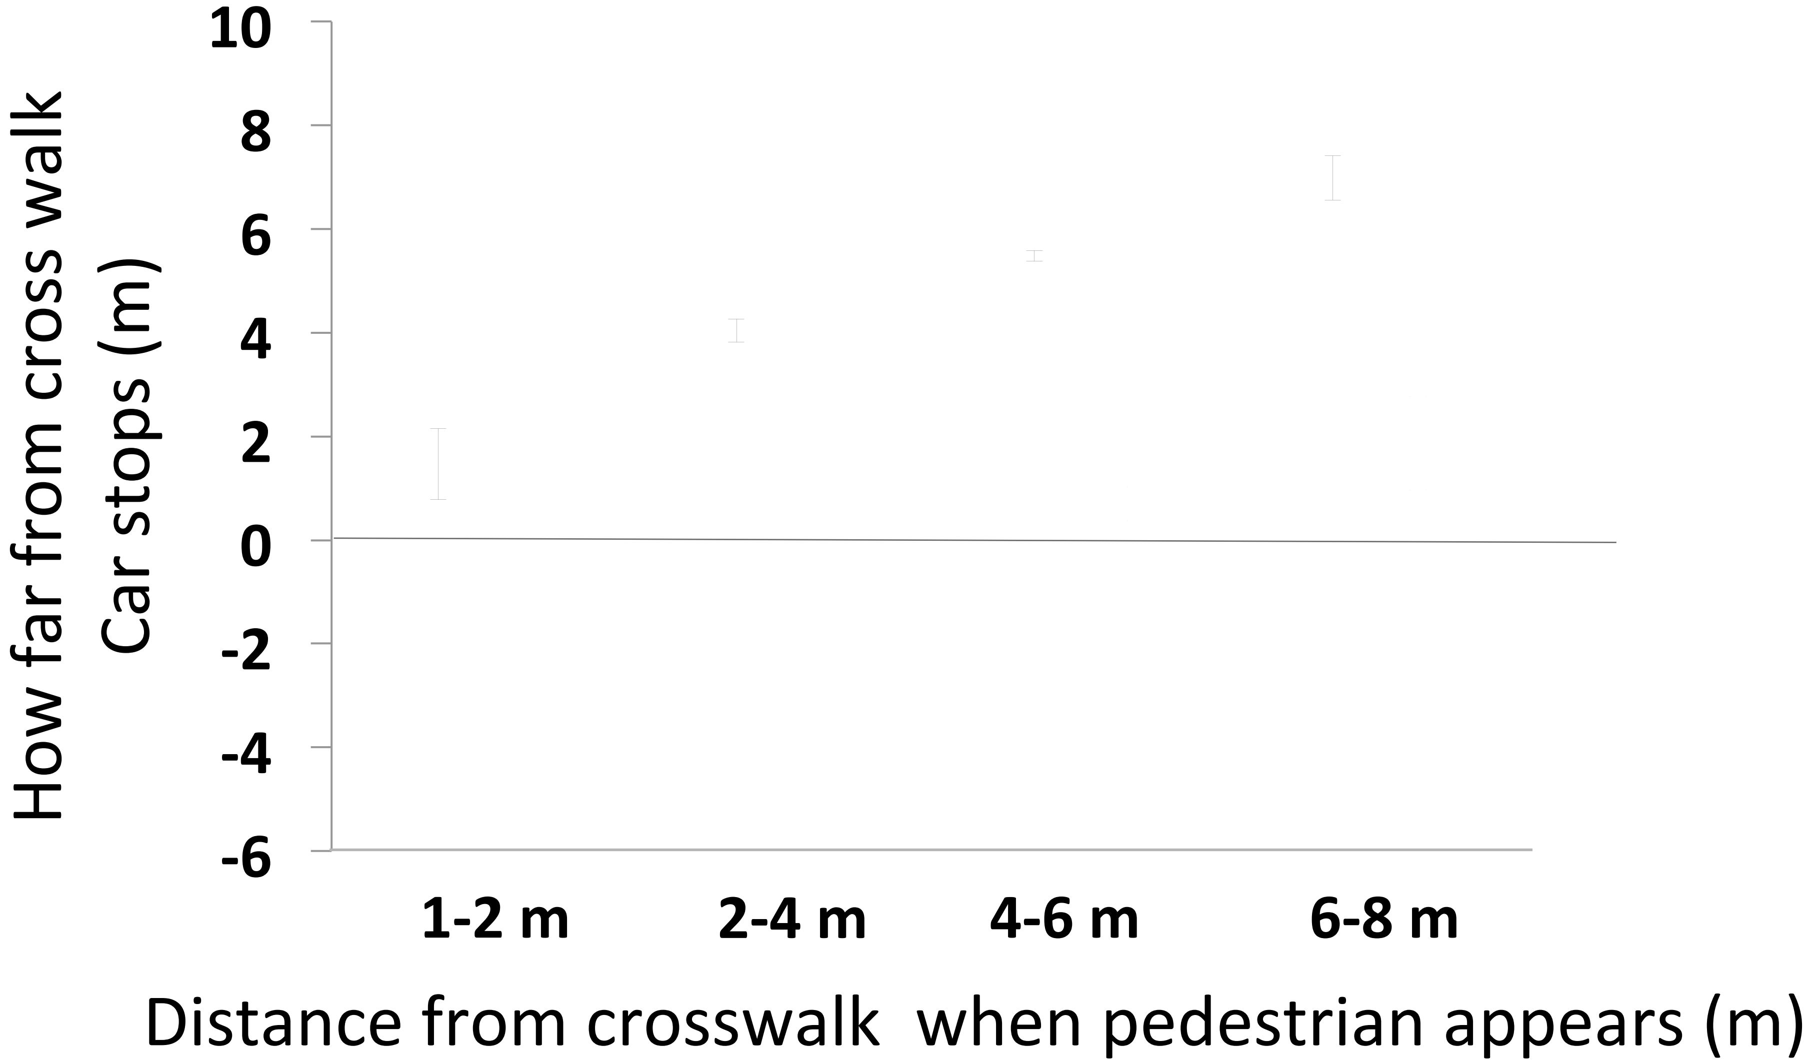

Outdoor Results

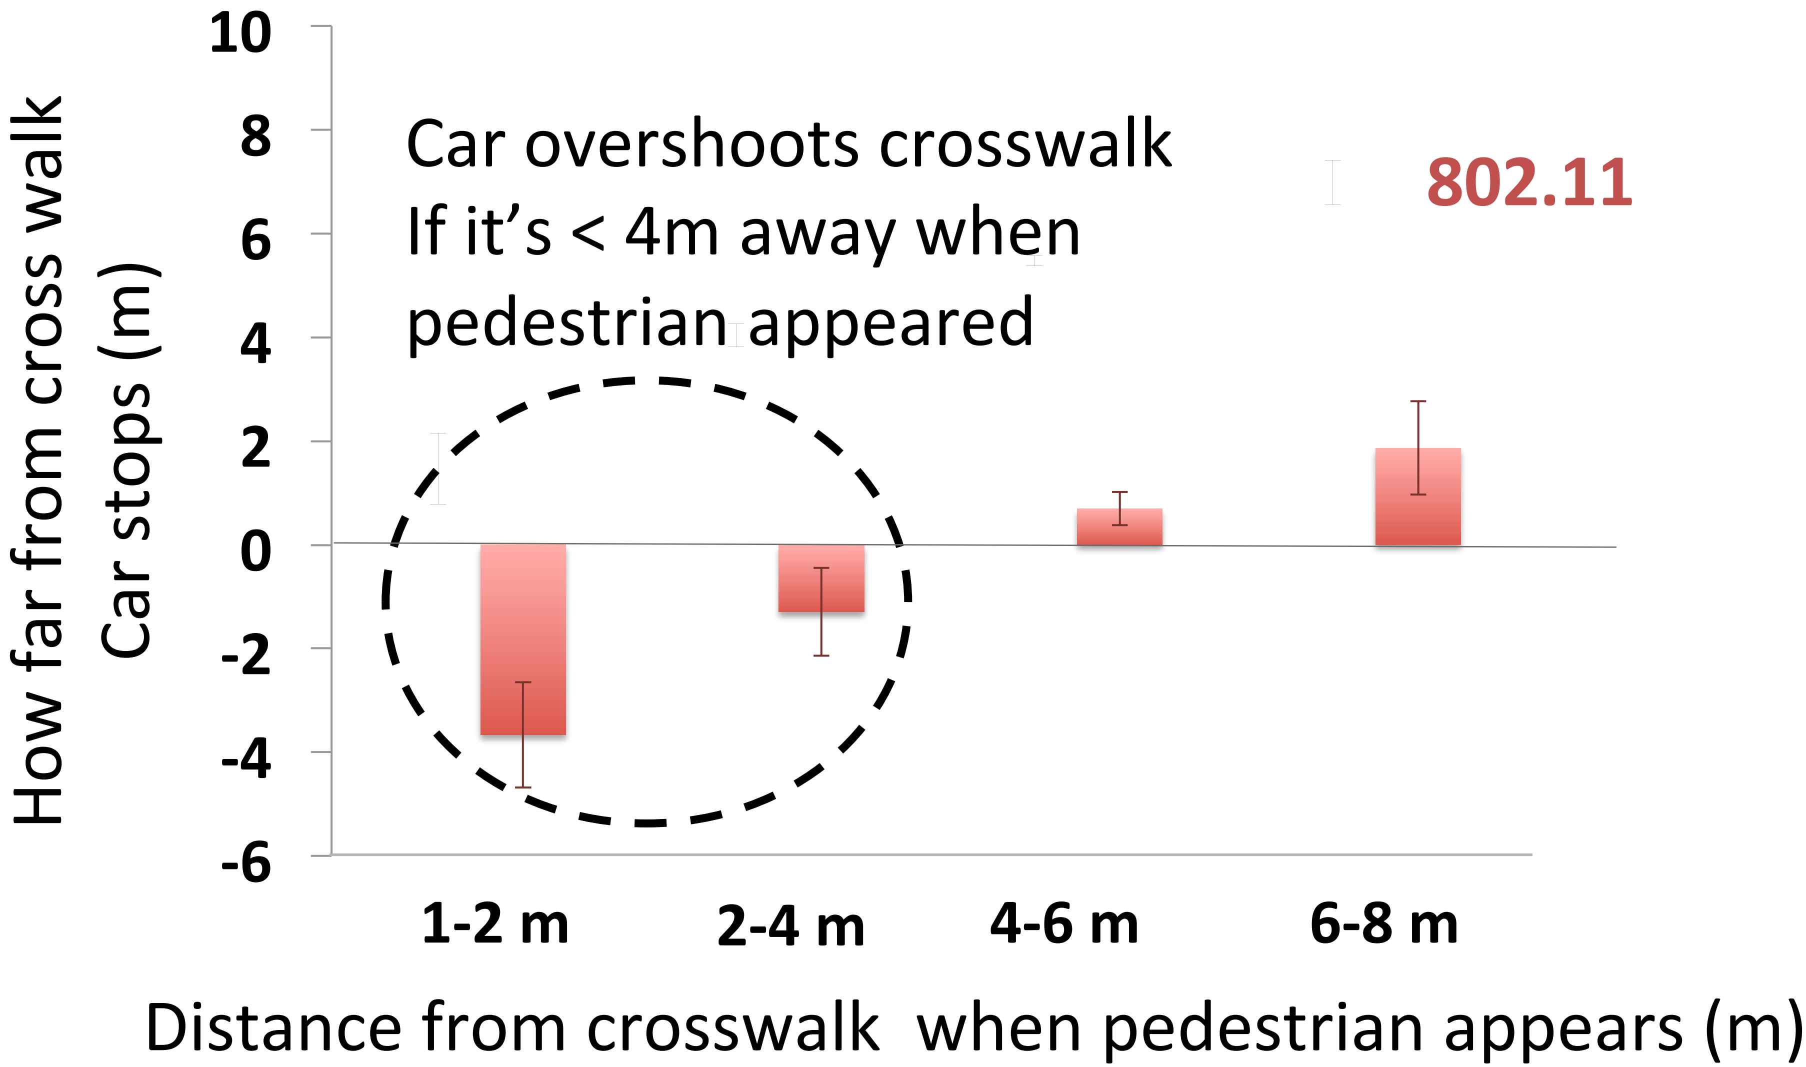

Outdoor Results

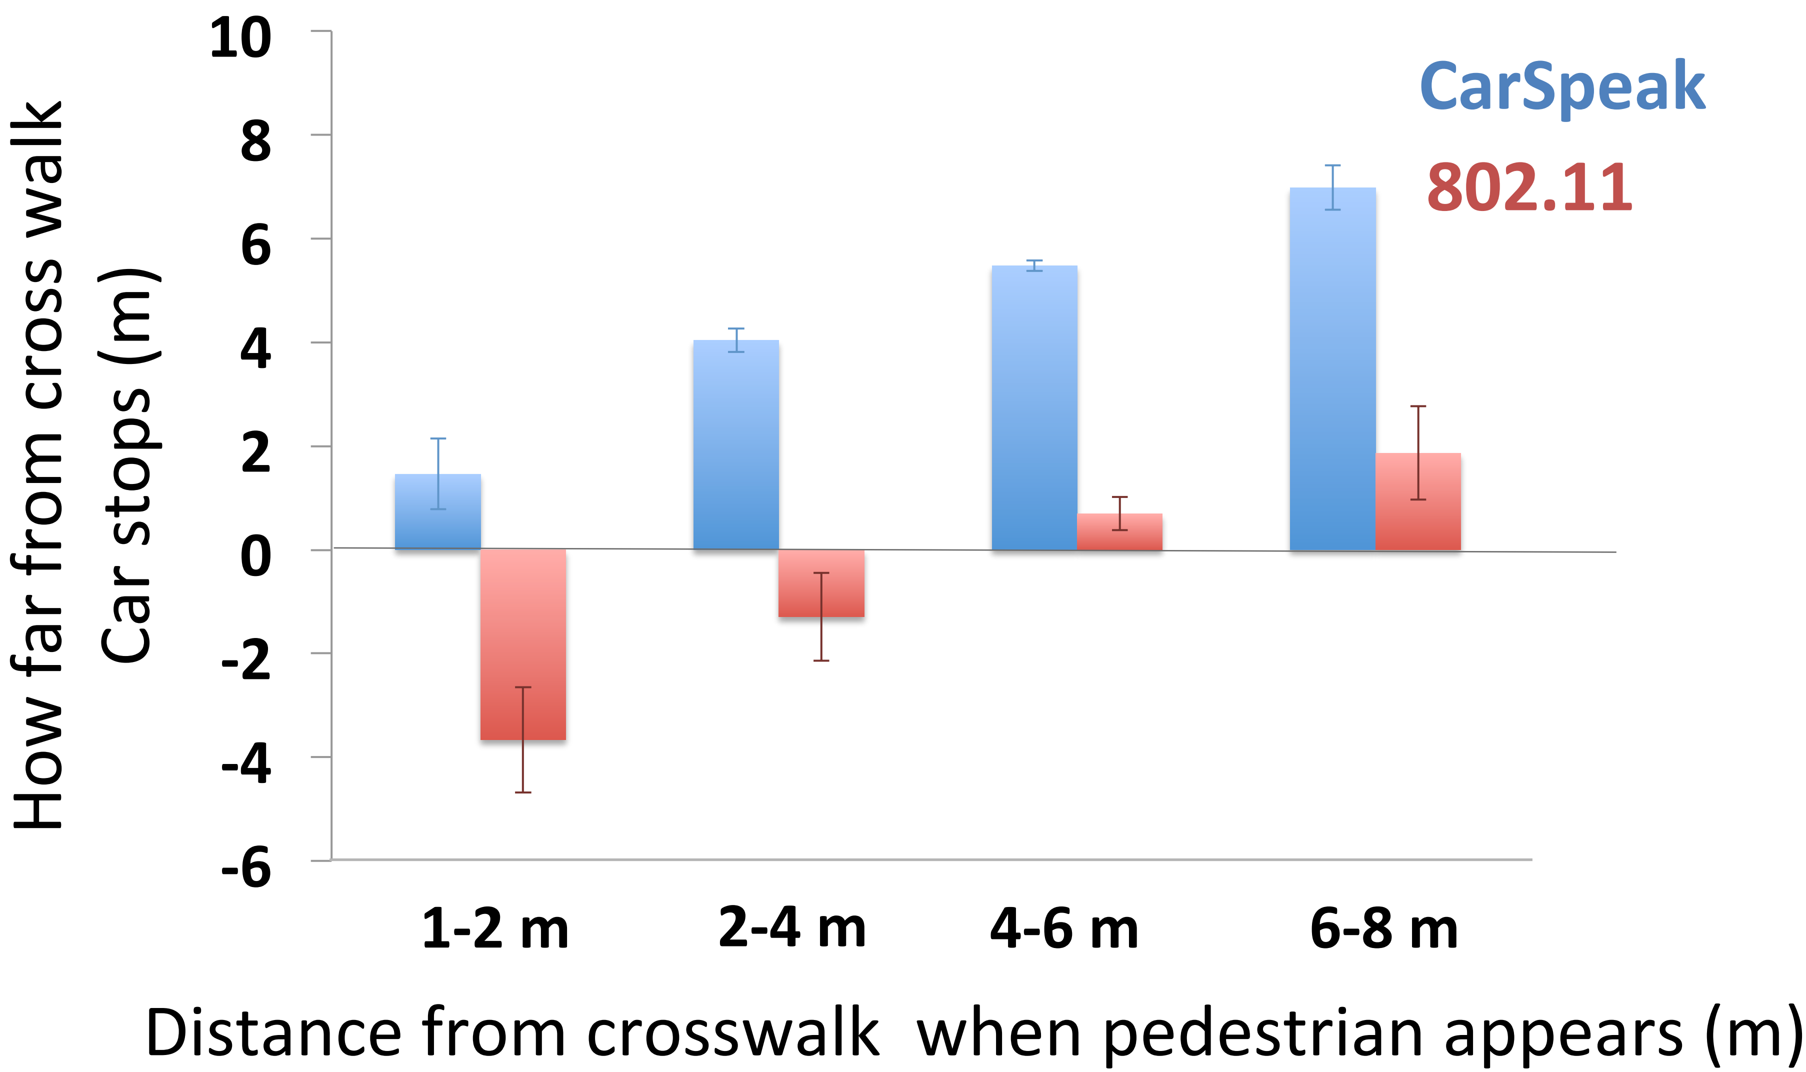

Conclusion

- A communication system fully integrated with autonomous driving
- Content-Centric approach to data access & MAC
- Generally applies to collaborative robotics, virtual reality and virtual games

Papers

Paper 4: Wagh, Aditya, et al. "Human centric data fusion in vehicular cyber-physical systems." Computer Communications Workshops (INFOCOM WKSHPS), 2011 IEEE Conference on. IEEE, 2011.

Introduction

- Vehicular Cyber-Physical Systems (VCPS), also known as Vehicular Ad Hoc Networks (VANET)
- Human Factors (HF) are critical
  - Final control unit of all fused information
  - Too many messages can overwhelm the driver
  - Need efficient and effective data fusion scheme

Overall Architecture

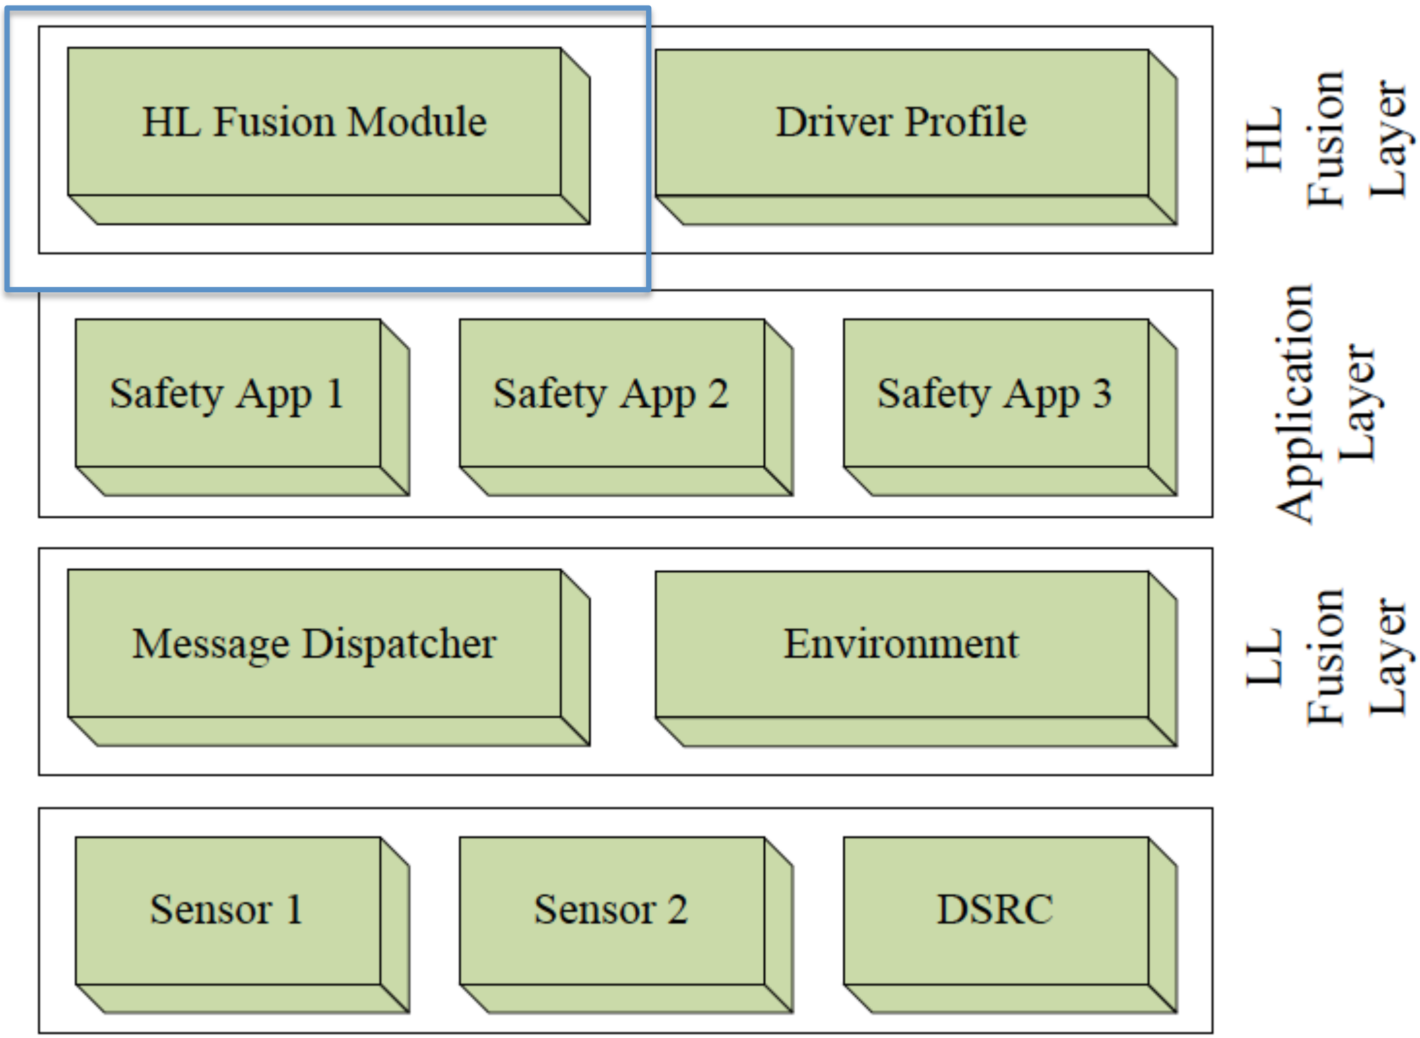

High-level Data Fusion

- Implication of warning message repetition
  - If a driver has responded to a particular warning (e.g., specific hazard) multiple times in the past, it is more strongly linked in the memory
  - If the required evasive actions are the same for two warning messages, notify the driver with the warning he/she is more familiar with
  - The response to repeated events is better when different kinds of events are alternated

High-level Data Fusion

- Evasive Actions:
  - Three key parameters:

  (Lance Choice, Direction, Speed)

- Goal: pick a small subset of evasive actions to avoid all hazards

High-level Data Fusion

- High-level Fusion Algorithm

Hazard	Location	Action Set
H_1	L_1	$(L_1, D_1, S_1), (L_2, D_2, S_2)$
H_2	L_1	(L_1, D_1, S_2)
H_3	L_3	$(L_1, D_1, S_1), (L_2, D_2, S_2)$
H_4	L_4	(L_4, D_4, S_4)
H_5	L_5	$(L_4, D_4, S_4), (L_5, D_5, S_5)$

Original: H1-H5

Q: How to remove redundant actions while still avoiding all hazards?

High-level Data Fusion

- High-level Fusion Algorithm
  - STEP 1: Location-based

Hazard	Location	Action Set
H_1	L_1	$(L_1, D_1, S_1), (L_2, D_2, S_2)$
H_2	L_1	(L_1, D_1, S_2)
H_3	L_3	$(L_1, D_1, S_1), (L_2, D_2, S_2)$
H_4	L_4	(L_4, D_4, S_4)
H_5	L_5	$(L_4, D_4, S_4), (L_5, D_5, S_5)$

Original: H1-H5 → Reduced: H1, H3-H5

High-level Data Fusion

- High-level Fusion Algorithm
  - STEP 2: Matching Actions

Hazard	Location	Action Set
H_1	L_1	$(L_1, D_1, S_1), (L_2, D_2, S_2)$
H_2	L_1	(L_1, D_1, S_2)
H_3	L_3	$(L_1, D_1, S_1), (L_2, D_2, S_2)$
H_4	L_4	(L_4, D_4, S_4)
H_5	L_5	$(L_4, D_4, S_4), (L_5, D_5, S_5)$

Original: H1-H5 → Reduced: H3-H5

High-level Data Fusion

- High-level Fusion Algorithm
  - STEP 3: Action Subsets

Hazard	Location	Action Set
H_1	L_1	$(L_1, D_1, S_1), (L_2, D_2, S_2)$
H_2	L_1	(L_1, D_1, S_2)
H_3	L_3	$(L_1, D_1, S_1), (L_2, D_2, S_2)$
H_4	L_4	(L_4, D_4, S_4)
H_5	L_5	$(L_4, D_4, S_4), (L_5, D_5, S_5)$

Original: H1-H5 → Reduced: H3, H4

Experiment Setup

- Environment: STISIM driving simulator with real drivers
- Three compared schemes:
  - No warning system
  - A warning system without data fusion
  - A warning system with data fusion

Experiment Setup

TABLE II: List of Warnings used in Simulation

Warning	Abbr.	Source	Trigger
Forward	FCW	Local	Obstacle detected directly in
			posted speed limit
Overtaking Warning	–	Radio	Vehicle overtaking from the subject vehicle's rear

No Fusion: the notification system attempts to deliver as many warnings as possible

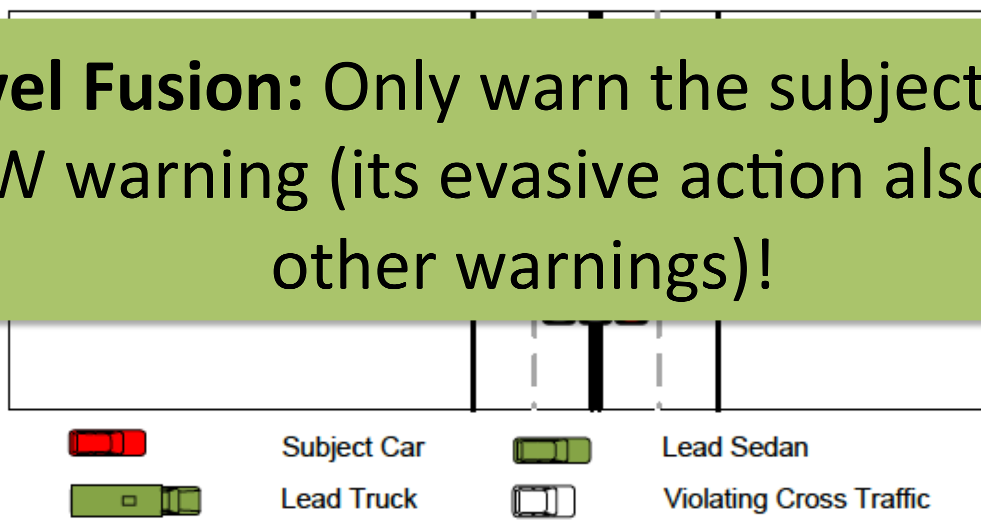

High-level Fusion: Only warn the subject car about the FCW warning (its evasive action also handles other warnings)!

Performance Analysis

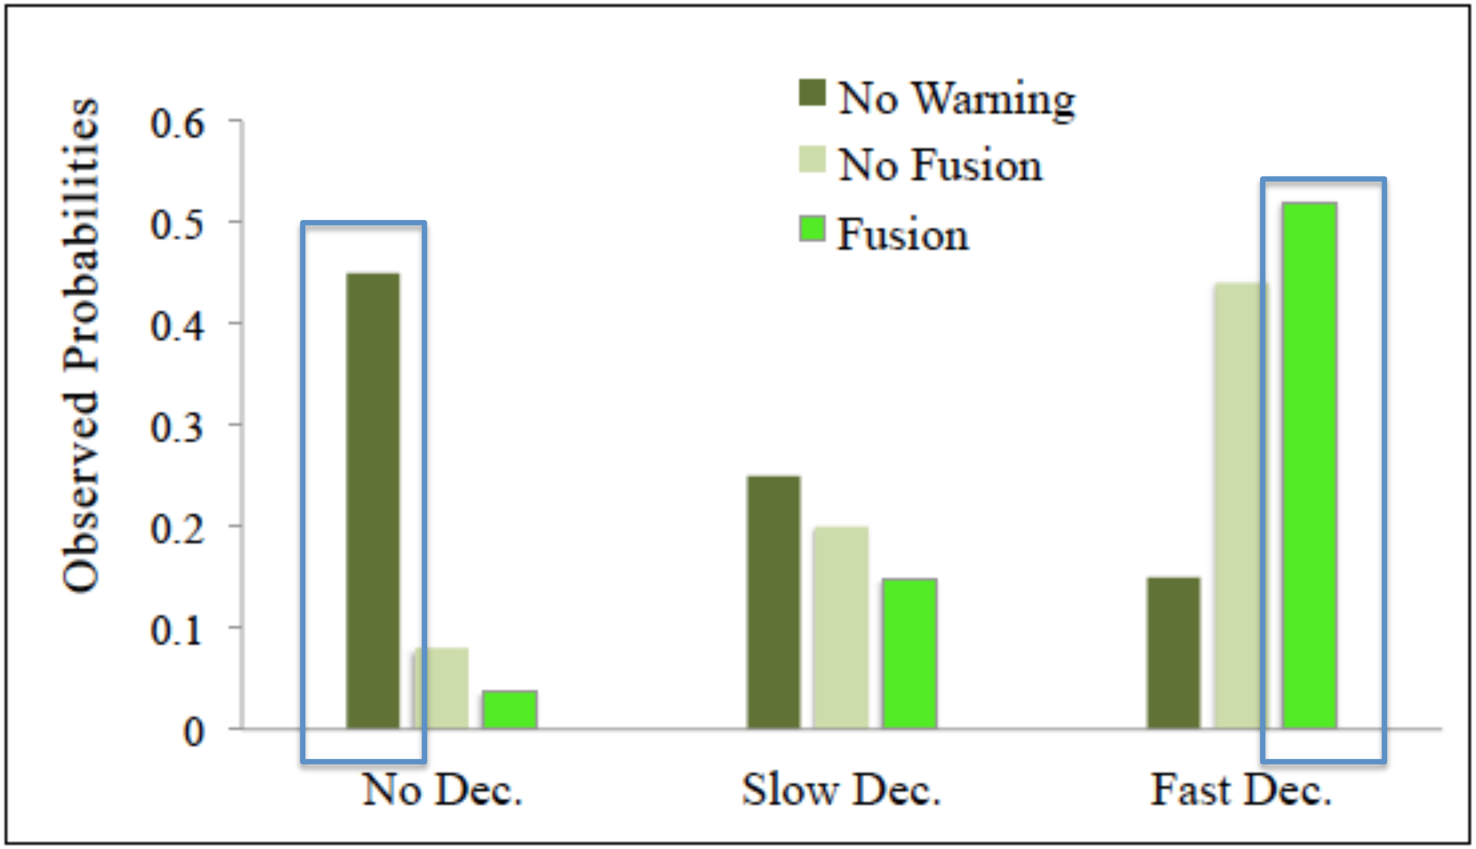

Performance Analysis

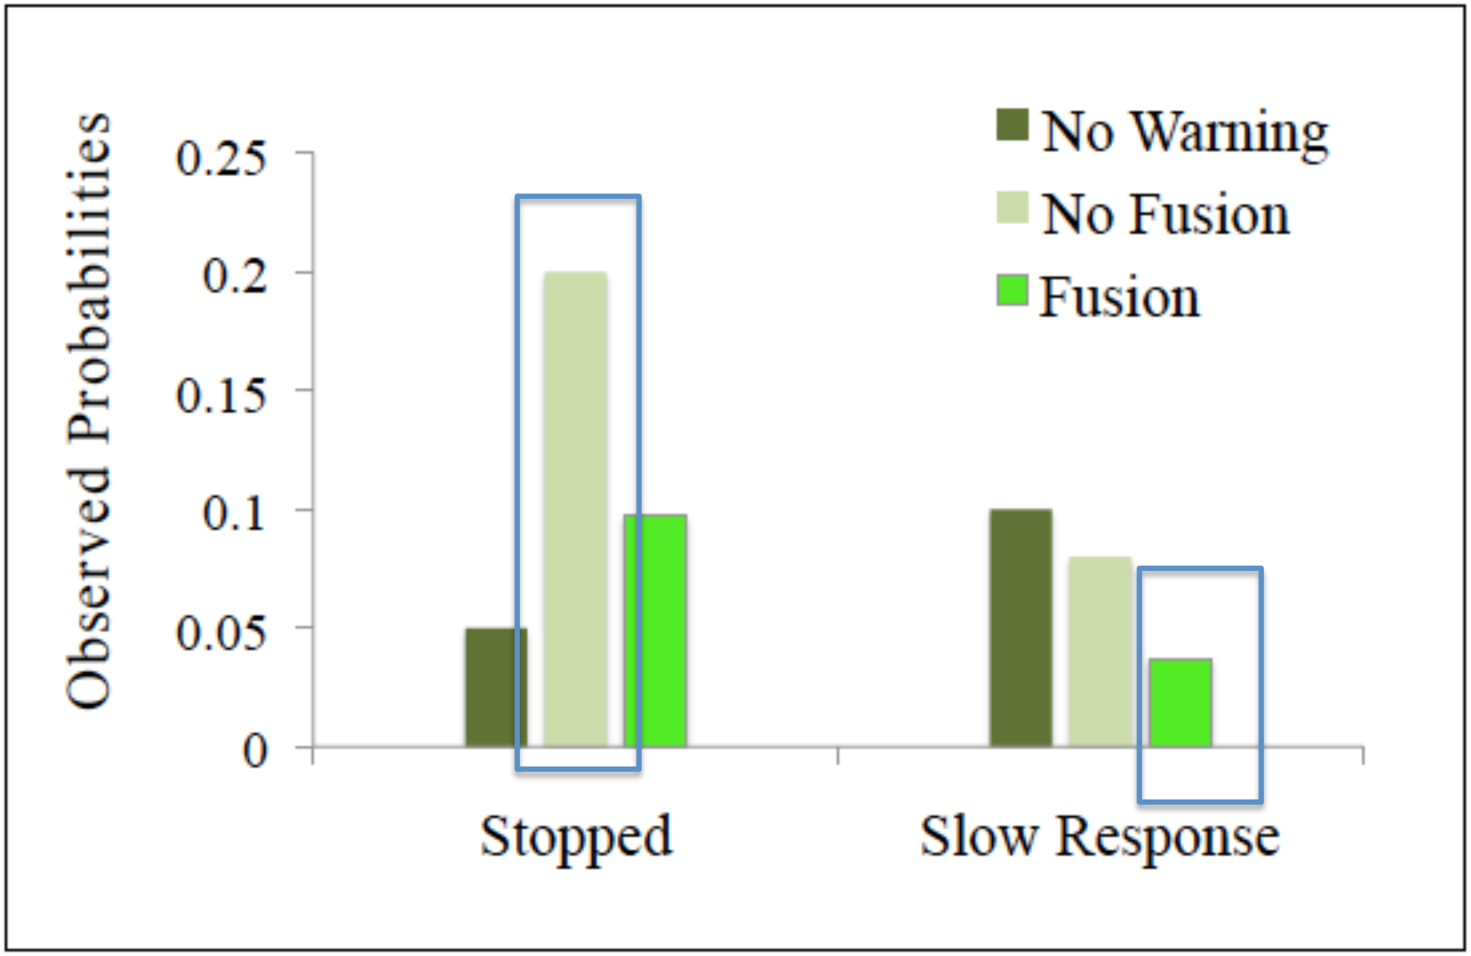

Conclusion

- A fusion system to **reduce the amount of non-beneficial information** delivered to drivers
- People **react better to fewer and more relevant warnings** rather than a deluge of correlated warnings
- This idea can also be applied in other related areas, e.g., **notification system for soldiers, aviation and naval system**