

# When Smart Phones Collaborate...

Margaret Martonosi  
Princeton University

Students: Sibren Isaacman, Emmanouil Koukoumidis

Other Collaborators: Ramon Cáceres, Li-Shiuan Peh,  
Rick Becker, Jim Rowland, Alex Varshavsky



# 10 Years ago: First ZebraNet paper at ASPLOS '02



During the 5-year project:

- ~5 distinct generations of tracking hardware
- 2 deployments of our collar designs
- Additional deployments of industry-made designs

Research contributions:

- Protocols
- Compression methods to save radio energy
- Middleware designs for radio-based updates
- ...

And graduated some great PhD students:

- Chris Sadler, Pei Zhang, ...



# Then and Now

- Then: Cellular telephony not ubiquitous, especially in less-developed countries and regions
- Now: Cellphones are the world's dominant computing platform
  - ⇒ Mobile services are burgeoning
- Then: GPS was still a niche technology, with high-wattage receivers and inaccurate readings
- Now: GPS is widespread, more accurate, and dropping in power requirements
  - ⇒ Location-aware services



# Opportunities and Challenges?

- Everyone has a GPS-based, communication-enabled sensor with them nearly all the time
  - Our work asks: What can we do with this? (For humans, not zebras!)
- Our work:
  - How can we use cellphones as implicit mobile sensing devices?
    - E.g. Noise pollution maps
  - How can mobility predictions and modeling improve performance of “opportunistic” networking?
    - And in turn provide higher-performance low-cost networking approaches for the world’s rural poor.
    - C-LINK
  - How can collaboration between nearby phones lead to interesting applications with better performance than non-collaborative versions?
    - SignalGuru



# Opportunities and Challenges?

- Everyone has a GPS-based, communication-enabled sensor with them nearly all the time
  - Our work asks: What can we do with this? (For humans, not zebras!)
- Our work:
  - How can we use cellphones as implicit mobile sensing devices?
    - E.g. Noise pollution maps
  - How can mobility predictions and modeling improve performance of “opportunistic” networking?
    - And in turn provide higher-performance low-cost networking approaches for the world’s rural poor.
    - C-LINK
  - How can collaboration between nearby phones lead to interesting applications with better performance than non-collaborative versions?
    - SignalGuru



# SignalGuru: The Goal

- How fast should I drive to hit all the green lights?
- In an arbitrary region, sense the traffic signal green-red transition times
- And give drivers accurate “optimal speeds” at which to drive to hit only green lights

## Tracking Signal transitions:

- Individually: Insufficient information
- Cloud-Archived/Historic: Too static. Non-adaptive.
- **Grassroots regional collaboration: Gather info quickly. Adjust even to highly-adaptive signal schedules**



© Thanva Tivawong




**2.4L Chrysler  
PT Cruiser '01**



**Scan Tool  
OBD-LINK  
device**

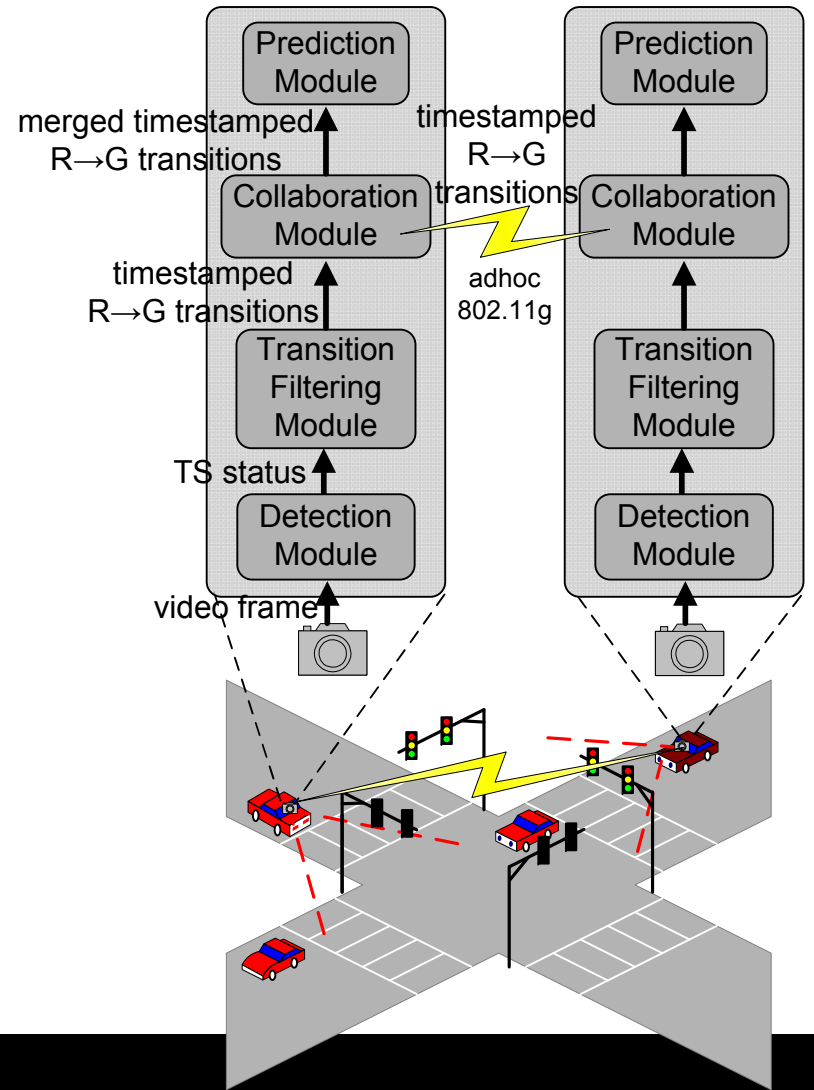


**OBDWiz software  
(IMAP)**



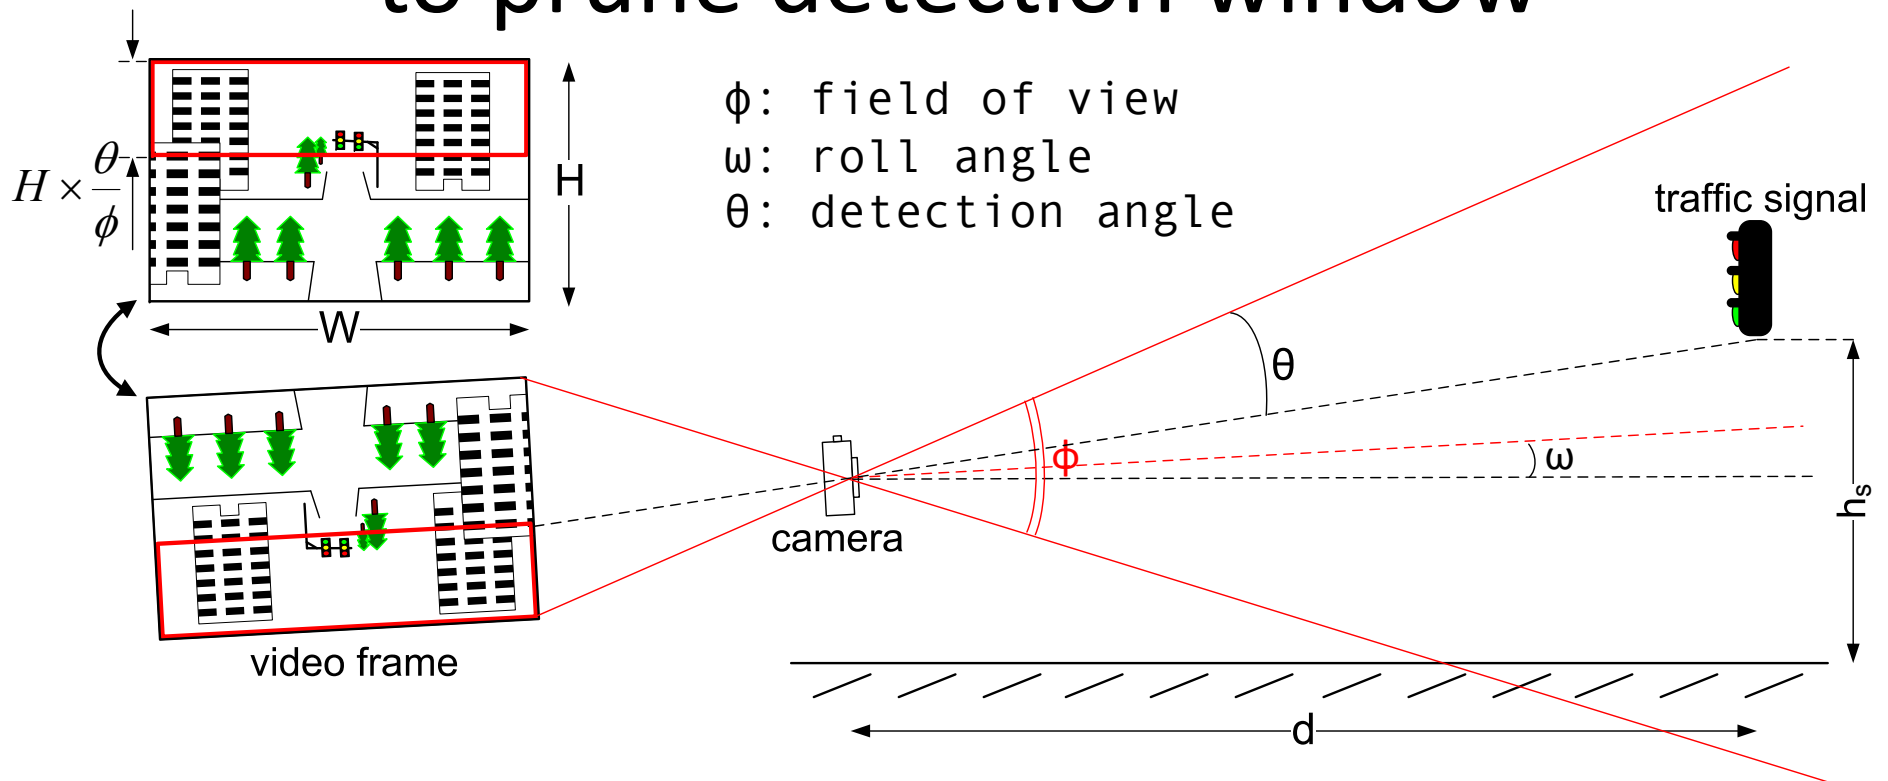
# Summary of Approach

- Sample video stream every 2 seconds
- Identify red or green stoplights
  - Local computation on iPhone
- Share with neighbors
- Predict appropriate speed





# Use phone's gyro and accelerometers to prune detection window

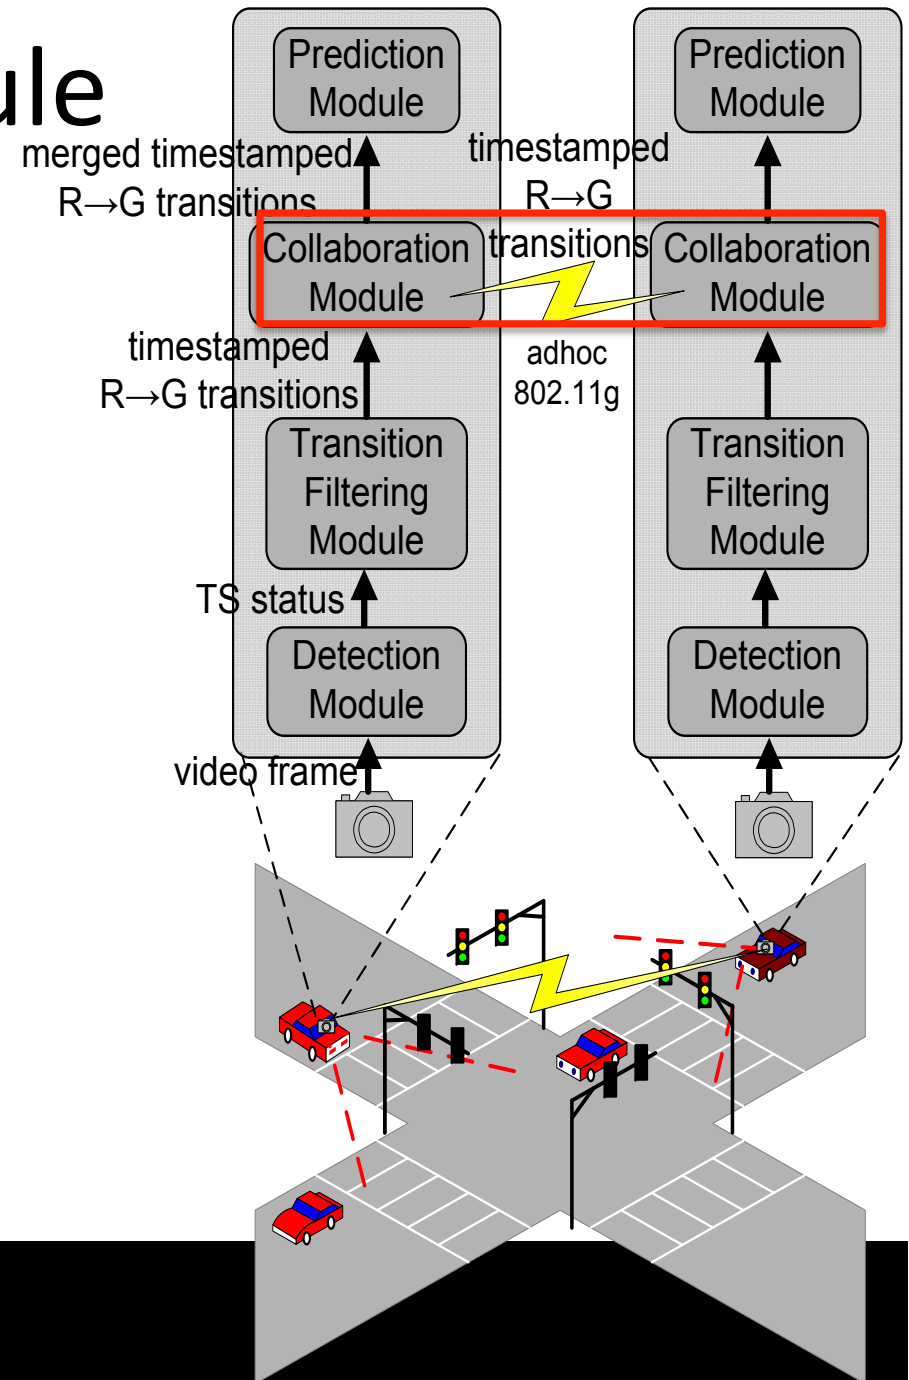


- Process only area within detection window.
  - Processing time reduced by 41%.
  - Misdetection rate reduced by 49%.



# Collaboration Module

- No cloud server.
- Real-time adhoc exchange of timestamped R→G transitions (last 5 cycles) database.
- Collaboration:
  - Improves mutual information.
  - Enables advance advisory.



# SignalGuru/GLOSA iPhone Application

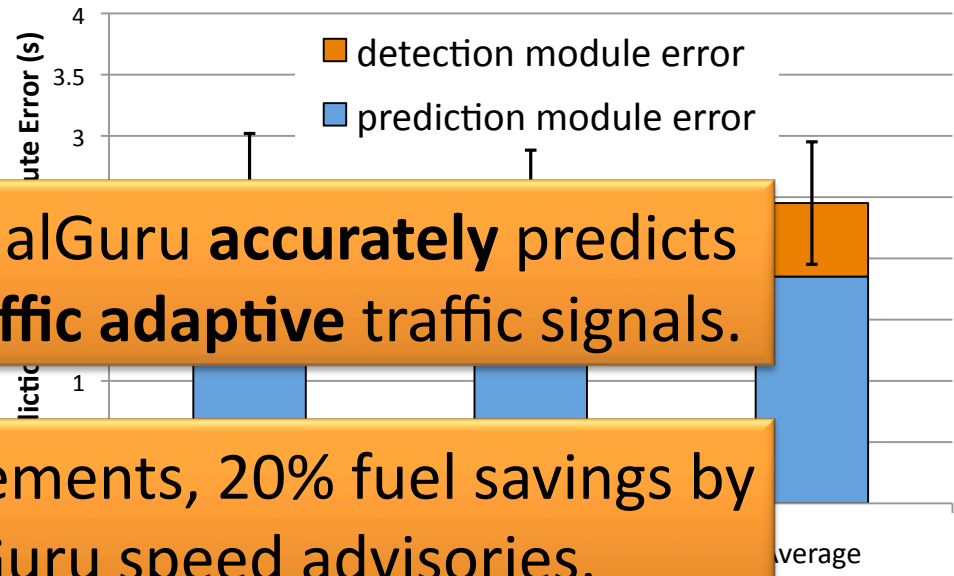
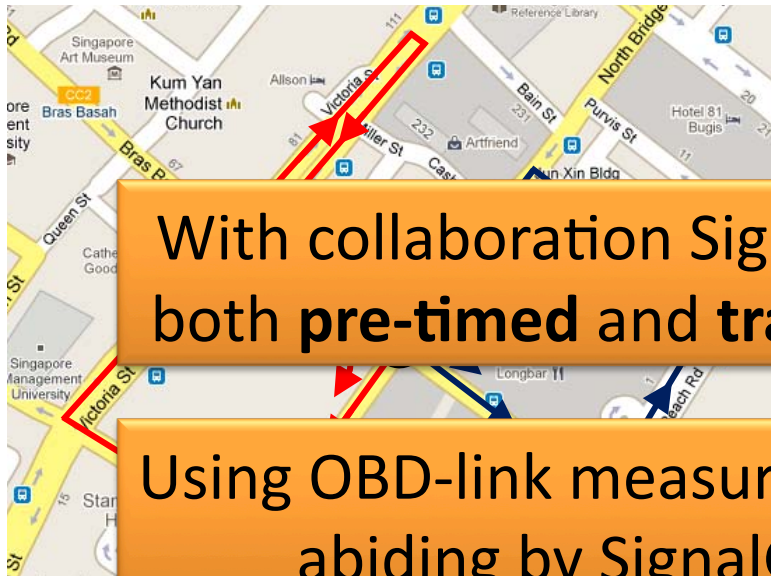
Residual amount of time in sec until the traffic signal turns green.  
Residual amount of time in sec until the traffic signal turns red again.  
Recommended speed.

The screenshot displays the SignalGuru application interface. At the top, there are control buttons: 'Connect', 'Sync', 'Camera', 'Adjust Exposure', and 'Auto Exposure'. On the right, the 'Detection' status is shown as 'ON' with a toggle switch and an information icon. The main display area is divided into two sections. The left section, titled 'Advisory', contains three data points: '3.0 sec' (highlighted with a green border), '45.0 sec' (highlighted with a red border), and '11.9 mph' (highlighted with a white border). The right section shows a live camera view of a city street with a traffic light. A red box highlights the traffic light area. In the bottom right corner of the camera view, a red-bordered box displays the value '0.69'. The 'SignalGuru' logo is visible in the bottom left corner of the interface.



# Singapore: Prediction Accuracy Evaluation

## Singapore



With collaboration SignalGuru accurately predicts both **pre-timed** and **traffic adaptive** traffic signals.

Using OBD-link measurements, 20% fuel savings by abiding by SignalGuru speed advisories.

- Traffic-adaptive traffic signals.
- Experiment in downtown:
  - 8 cars over 30 min.
  - 2 signals, 26 transitions.
- **Error<sub>Average</sub> = 2.45sec (3.8%).**
  - **Error<sub>Transition Detection</sub> = 0.60sec (0.9%).**
  - **Error<sub>Phase Length Prediction</sub> = 1.85sec (2.9%).**
- **Without collaboration:**
  - **Error<sub>Average</sub> = 11.03sec**



# SignalGuru: Leveraging Mobile Phones for Collaborative Traffic Signal Schedule Advisory

Emmanouil Koukoumidis, Li-Shiuan Peh, Margaret Martonosi

MobiSys '11 (Best Paper Award)

*<http://www.princeton.edu/~ekoukoum/SignalGuru.html>*

