Nericell: Rich Monitoring of Roads and Traffic Using Mobile Smartphones

Ram Ramjee

Joint work with Prashanth Mohan & Venkat Padmanabhan

(in ACM SenSys 2008)

Mobility, Networks, and Systems Group

Microsoft Research India

Road and Traffic Monitoring

Courtesy: FreeDigitalPhotos.net

Bangalore ≠ <your favorite developed city>

What’s Different?

- Varied vehicle types
- Liberal honking
- Chaotic intersections
- Potholes
- Road bumps
- …

Need to go beyond GPS-based vehicle tracking!

Widespread distribution of mobile phones

Road and Traffic Monitoring

- Without deployed infrastructure
- Using existing mass of mobile phones

Mobile Phones

- ~4 billion phones worldwide
- ~400 million phones in India
- ~10 million new connections every month in India
- 115 million of 1 billion phones sold worldwide in 2007 were smartphones
- Smartphone market share expected to reach nearly 50% by 2012 in North America

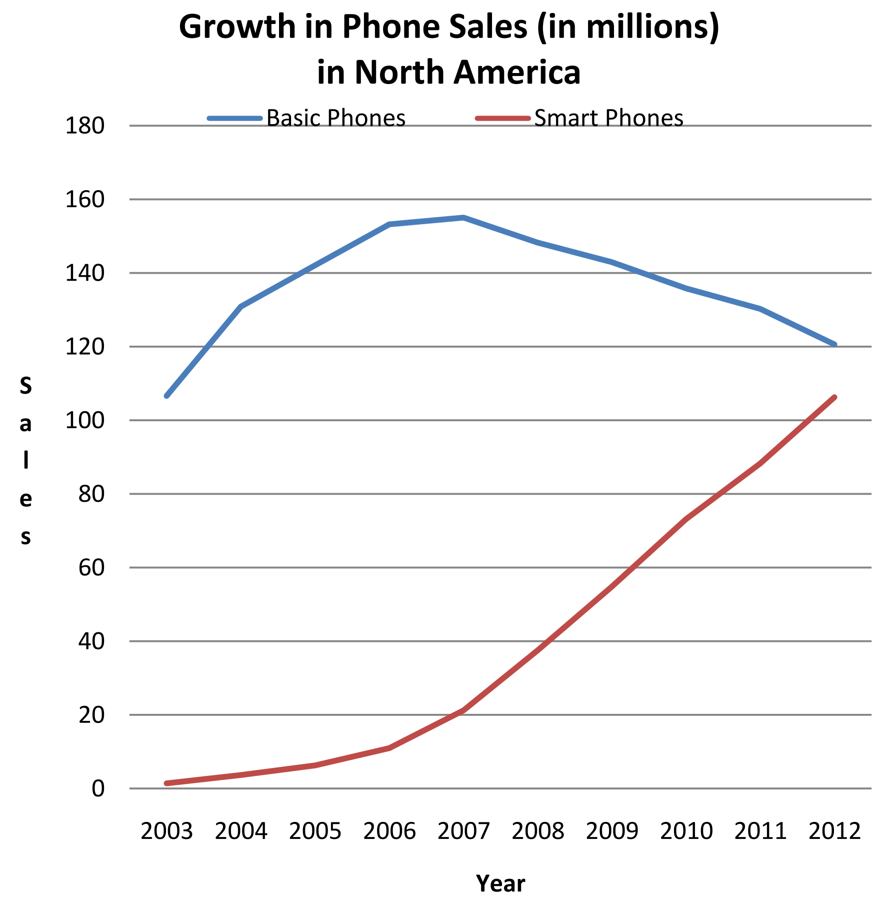

Mobile Smartphones

- Mobility
- Computing + communication + sensing
- Far more capable & ubiquitous than special-purpose sensors

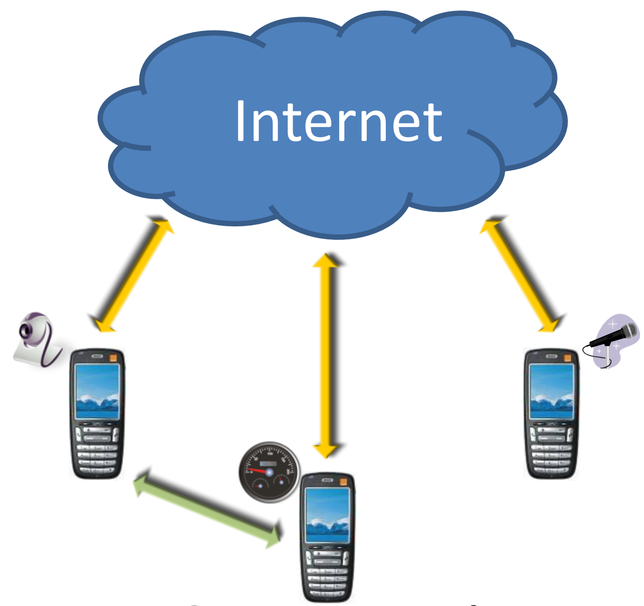

Opportunity

Distributed Sensing using Mobiles

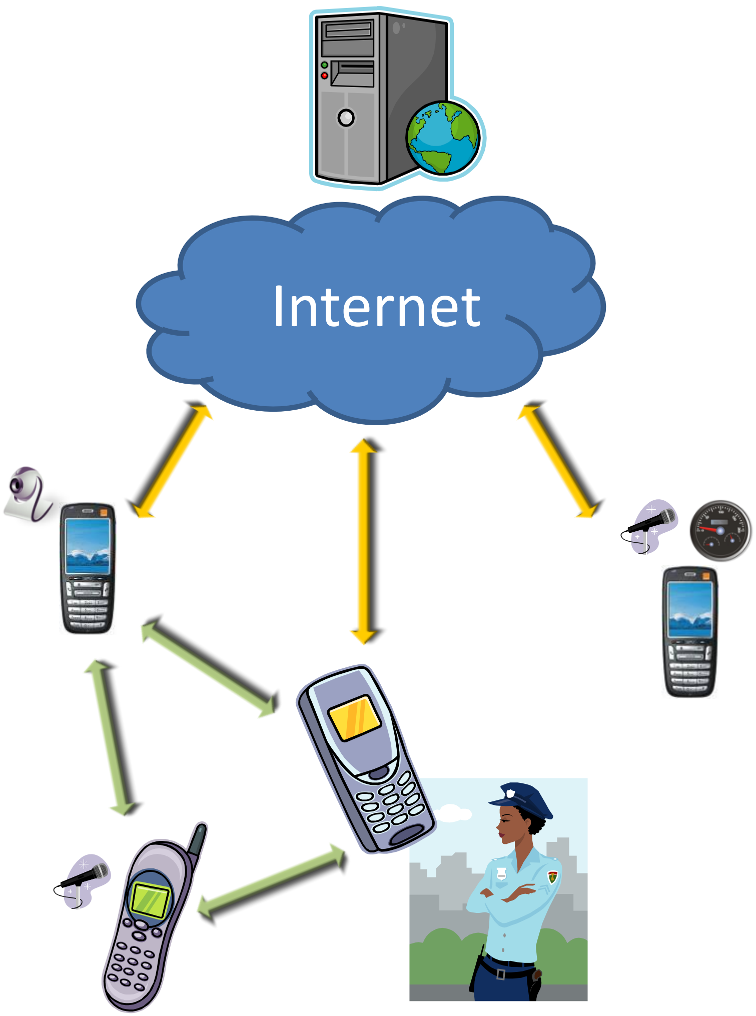

- Applications
  - Road & traffic monitoring
  - On-demand webcam
  - Human-powered search
  - …

Outline

- Mobile smartphone-based distributed sensing
- Nericell design and evaluation
- PRISM platform
- Related work
- Microsoft Research India

Nericell

- Idea: mobile smartphone-based sensing
  - enables rich monitoring
  - avoids dependence on infrastructure
- Challenges: energy, automated operation, privacy
- Key components
  - accelerometer $\Rightarrow$ road and drive quality
  - microphone $\Rightarrow$ honk detection, vehicle type
  - GSM radio $\Rightarrow$ lightweight localization
- Sensors also used in combination

Energy is a key challenge

Resource	Power (mW)
Bluetooth	20
Wi-Fi	770
GPS	620
Microphone	225
Accelerometer	2

Energy consumption on iPaq hw6965

Accelerometer-based Sensing

- Advantage: low energy cost
- Challenge: “disorientation”
- Analyses:
  - braking detection
  - bump/pothole detection
  - pedestrian versus stop-and-go traffic

Braking Detection

- Braking impacts drive quality
- Two approaches:
  - GPS: high energy cost (600 mW on iPAQ hw6965)
  - Accelerometer: much cheaper (2 mW + 30mW)
- Accelerometer-based braking detection:

Threshold T(g)	False Negative	False Positive
0.12 (4 sec)	11%	16%

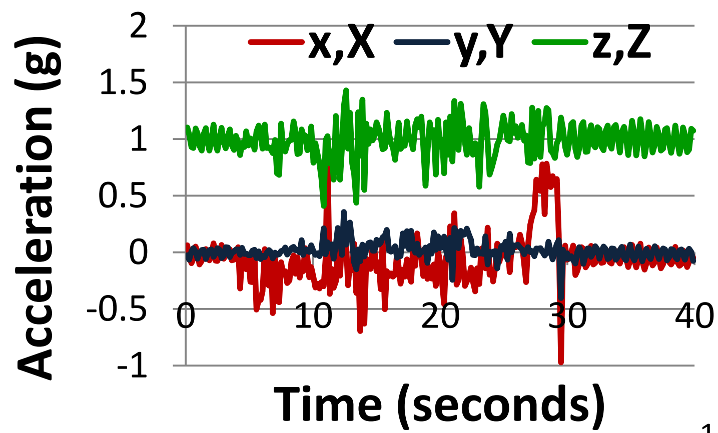

Virtual Reorientation

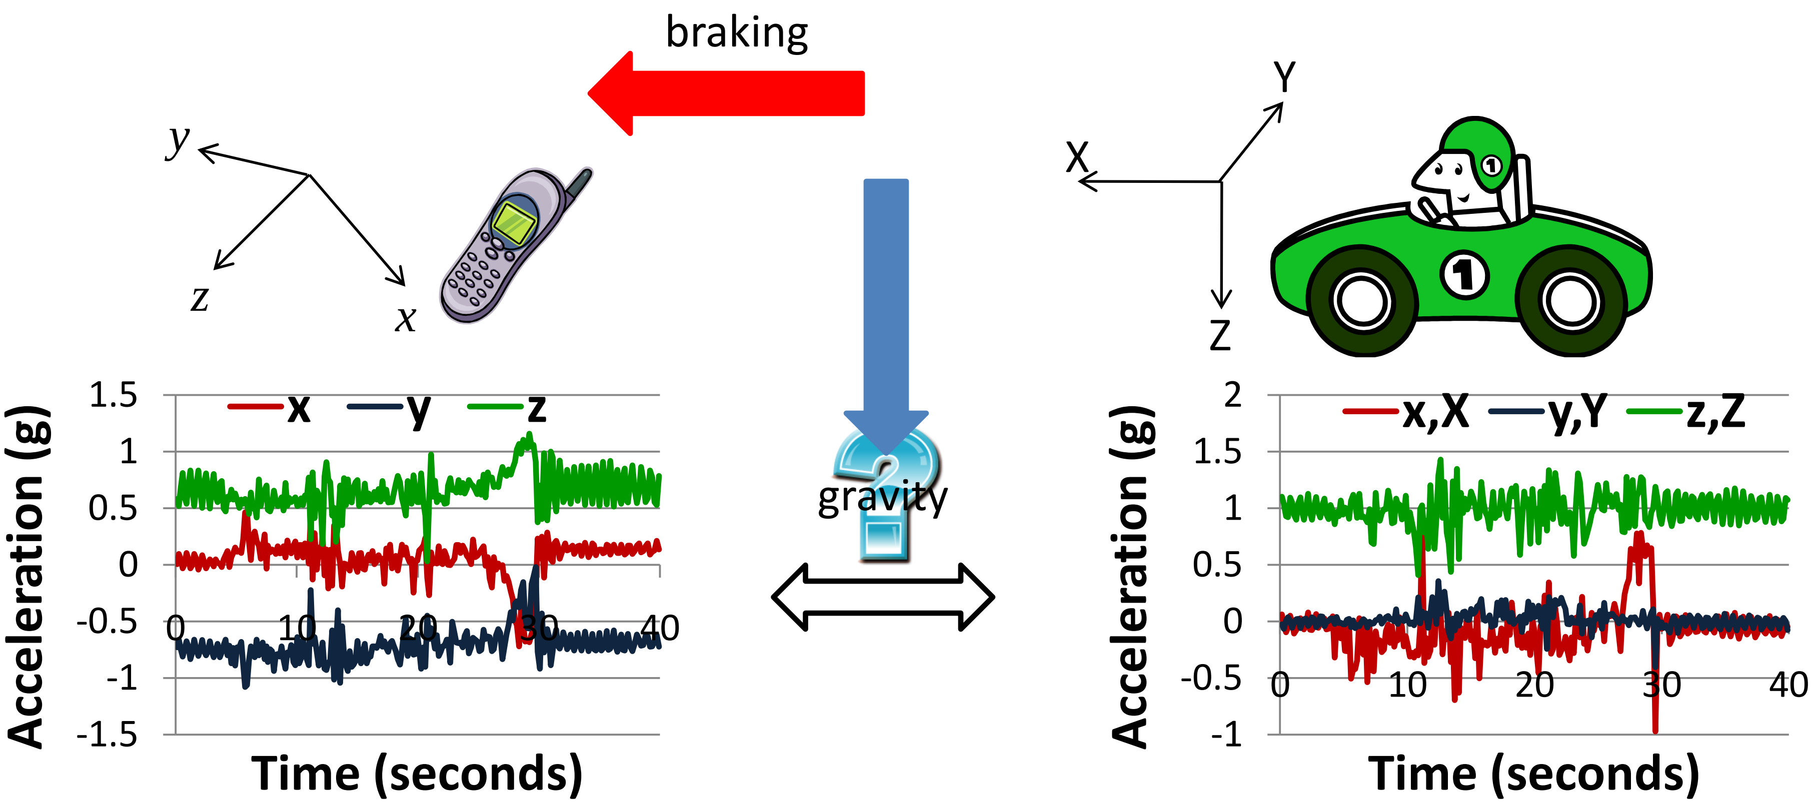

Virtual Reorientation

- Euler Angles:
  - Any orientation of the accelerometer can be represented by Z-Y-Z (and other equivalent) rotations
- Three Unknowns (angles):
  - pre-rotation (ϕ_{pre}),
  - tilt (θ_{tilt}),
  - post-rotation (ψ_{post})
- Knowns:
  - Gravity along Z
  - Braking along X

Virtual Reorientation Using Gravity

- Ideal orientation (X,Y,Z): $a_X = 0$; $a_Y=0$; $a_Z=1(g)$;
- Current orientation (x,y,z) with force a_x , a_y , a_z
- $a_z = a_Z \cos(\theta_{tilt})$

$$\boxed{\theta_{tilt} = \cos^{-1}(a_z)}$$

- $a_x = a_Z \cos(\phi_{pre}) \sin(\theta_{tilt})$
- $a_y = a_Z \sin(\phi_{pre}) \sin(\theta_{tilt})$

$$\boxed{\phi_{pre} = \tan^{-1}(a_y / a_x)}$$

Virtual Reorientation Using Braking

- Use GPS to identify braking
- Ideal orientation (X,Y,Z): a_X = large; a_Y=0; a_Z=1(g);
- Current orientation (x,y,z) with force a_x , a_y , a_z and angles θ_{tilt} and ϕ_{pre}
- Find ψ_{post} such that force along X is maximized

$$\psi_{post} = \tan^{-1}\left(\frac{-a_x \sin(\Phi_{pre}) + a_y \cos(\Phi_{pre})}{(a_x \cos(\Phi_{pre}) + a_y \sin(\Phi_{pre}))\cos(\theta_{tilt}) - a_z \sin(\theta_{tilt})}\right)$$

Automatic Virtual Reorientation

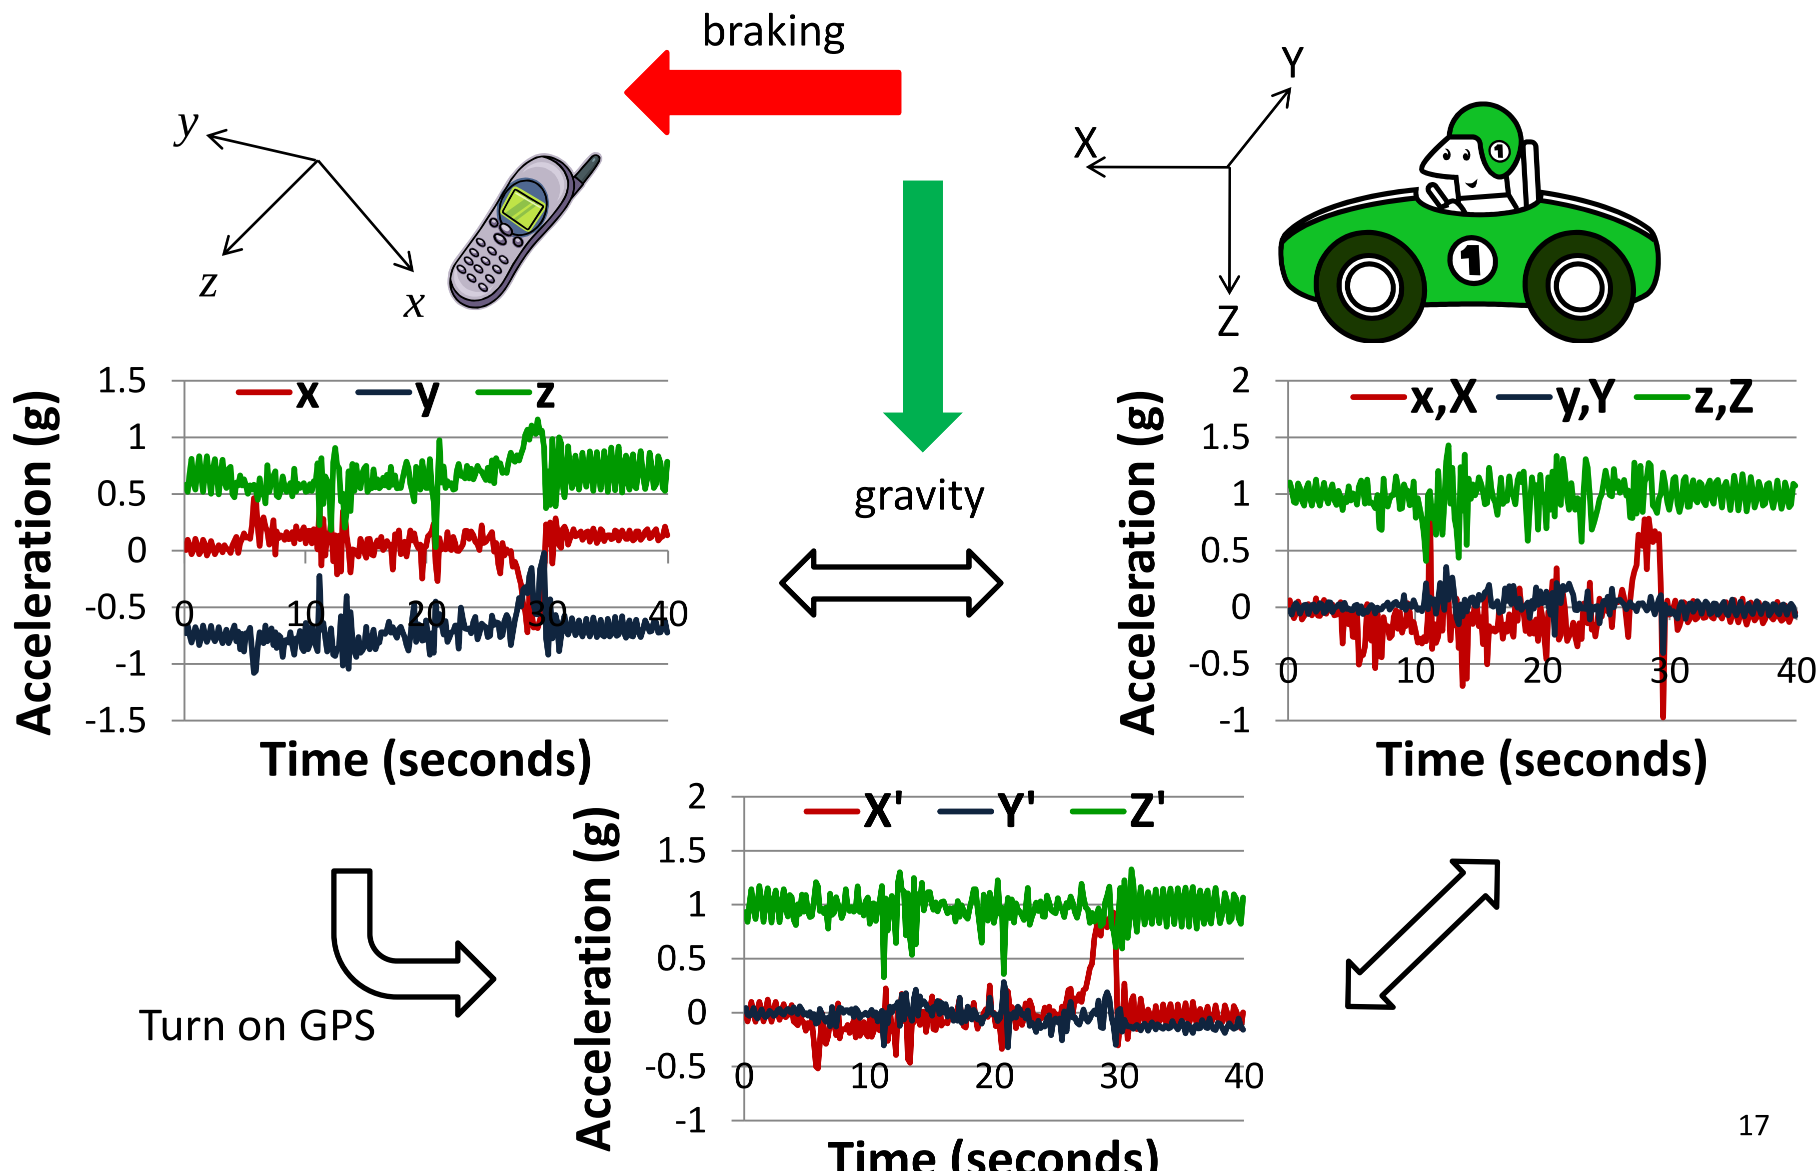

Results: Virtual Reorientation

Sr No	$\Phi_{pre}/\theta_{tilt}/\psi_{post}$	Cross correlation	
		Well oriented – Well oriented	**Reoriented – Well oriented**
1	7° /38° /106°	0.90	0.91
2	174° /34° /-107°	0.75	0.87
3	174° /34° /-107°	0.94	0.90
4	4° /42° /12°	0.74	0.68
5	3° /44° /-1°	0.76	0.79
6	-80° /42° /121°	0.78	0.73

Braking detection with Virtual Reorientation

	False negatives	False positives
Well-oriented	11%	16%
Virtually reoriented	11%	18%

Pothole Detection

Pothole Detection

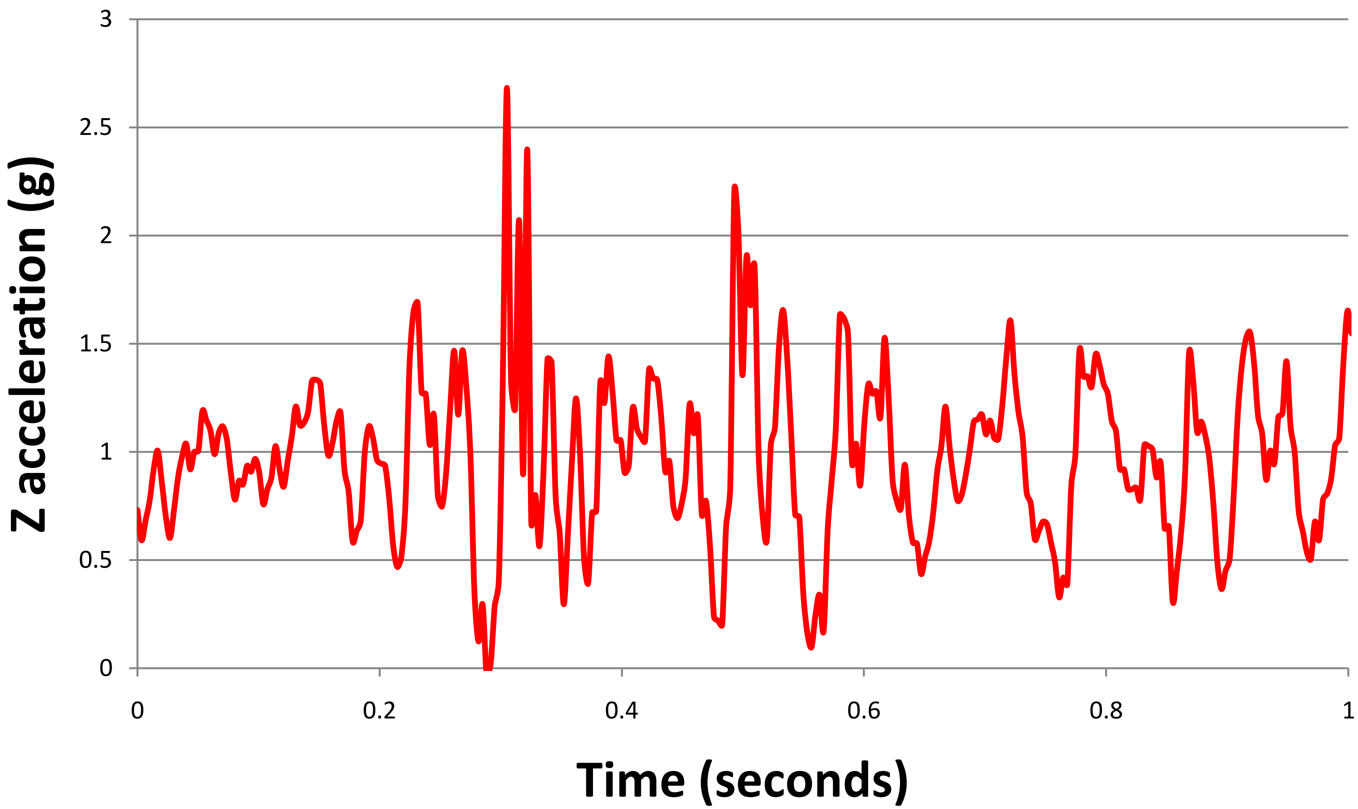

High speed (≥ 25 kmph)

z-peak: look for significant spike

Pothole Detection

Pothole Detection

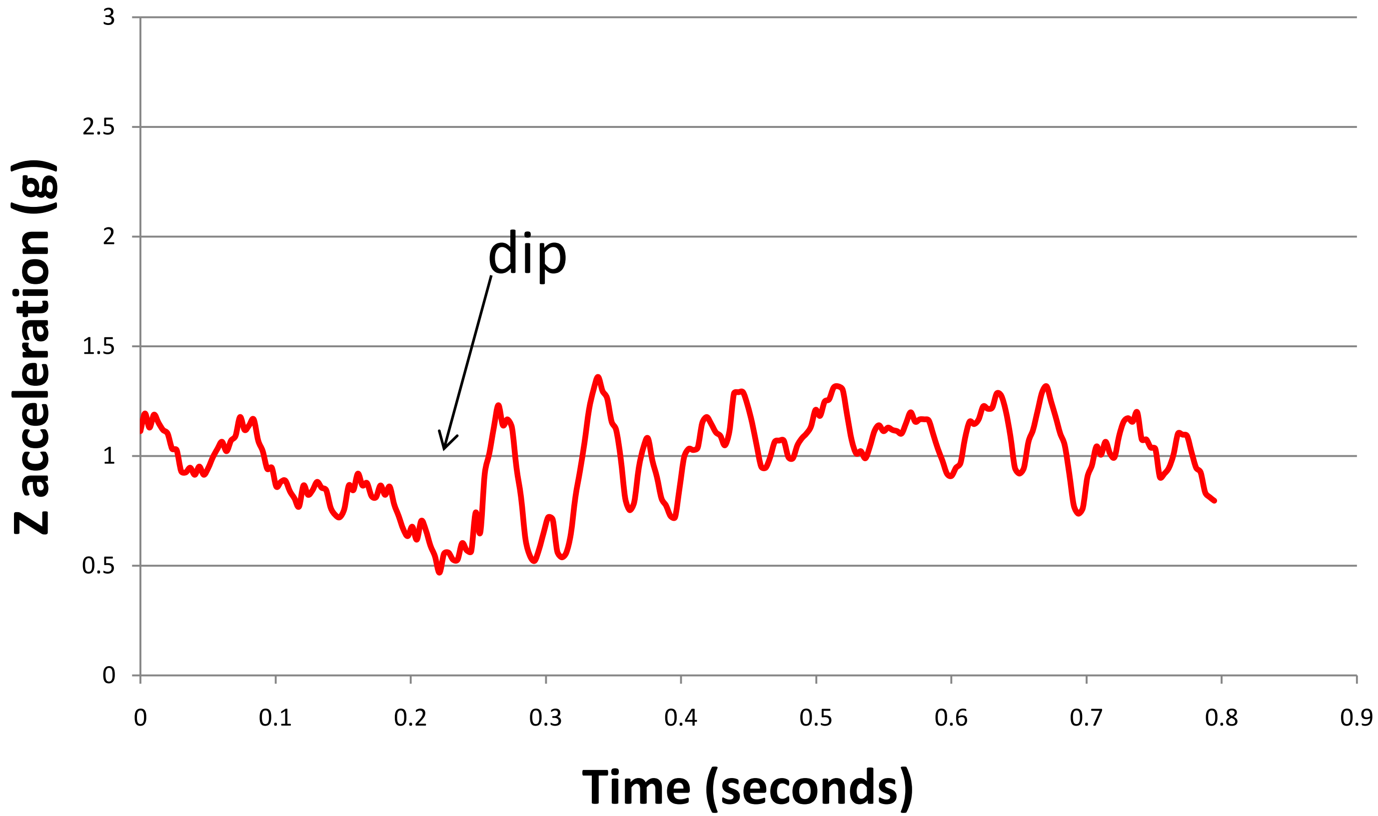

Low speed (< 25 kmph)

z-sus: look for sustained dip

Results: Pothole Detection

Training data: 5km long drive with 44 bumps
Test data: 35km long drive with 101 bumps
False Negative: missed pothole (not so bad)
False Positive: incorrectly identified as pothole (not so good)

Threshold	Speed < 25 kmph		Speed > 25 kmph	
	False Neg	False Pos	False Neg	False Pos
Z-sus (0.8g, 20ms)	37%	14%	0%	136%
Z-peak (1.45g)	65%	21%	3%	49%
Z-Peak (1.75g)	83%	0%	41%	8%

Determining Location

- Why not just GPS?
  - coverage (indoors, urban canyons, …)
  - time to lock (~26 secs even with warm start)
  - energy (~600 mW on iPAQ 6965)
  - not all phones have it
- Alternative: GSM tower matching
  - match towers seen against those in training set
  - widely accessible, fast, “zero” energy
  - location: median error: 130m, 90^{th} %tile: 610m
  - speed: median error: 3.4 kmph, 90^{th} %tile: 11.2 kmph

Locating a Pothole

- Accelerometer is cheap, so keep on continuously
- When bump is detected, use GSM tower matching to estimate location
  - median error: 130 m, 90th percentile: 610 m
- Send bump report to server
- If several reports in same vicinity, server triggers GPS on other phones for location fix
- Sample result: GPS turned on only 3.2% of the time on a 20 km drive with one point of interest

Pedestrians vs. Stop-and-Go Traffic

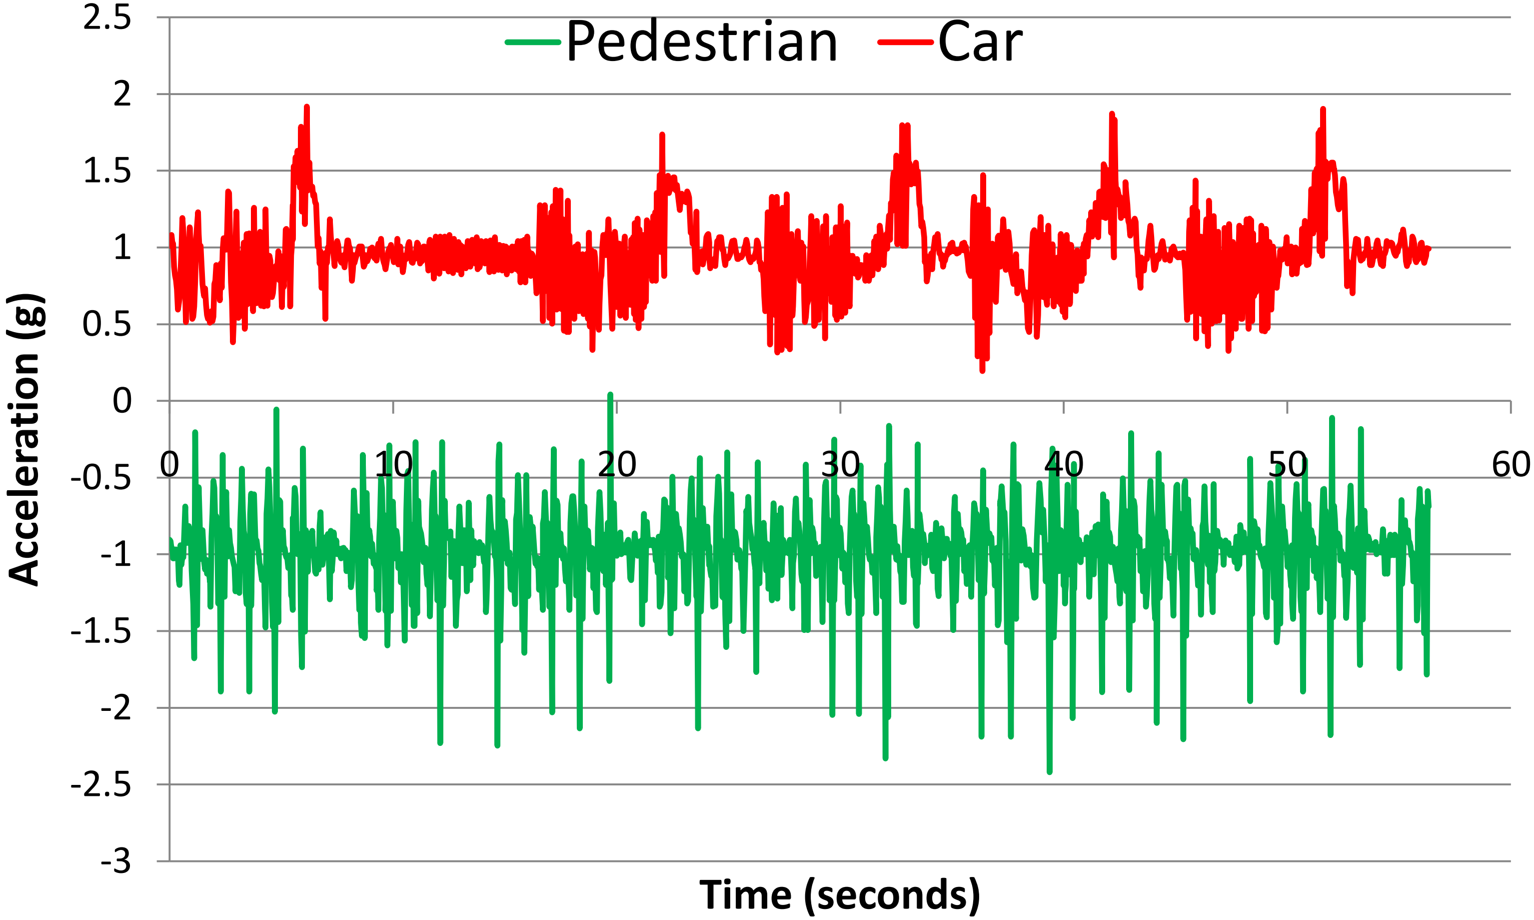

Accelerometer helps disambiguate between the two

Microphone-based Sensing

- Advantage: ubiquity
- Challenge: energy, privacy
- Analyses:
  - honk detection: triggered when accelerometer indicates a lot of braking
  - vehicle type: exposed versus enclosed vehicle

Honk Detection

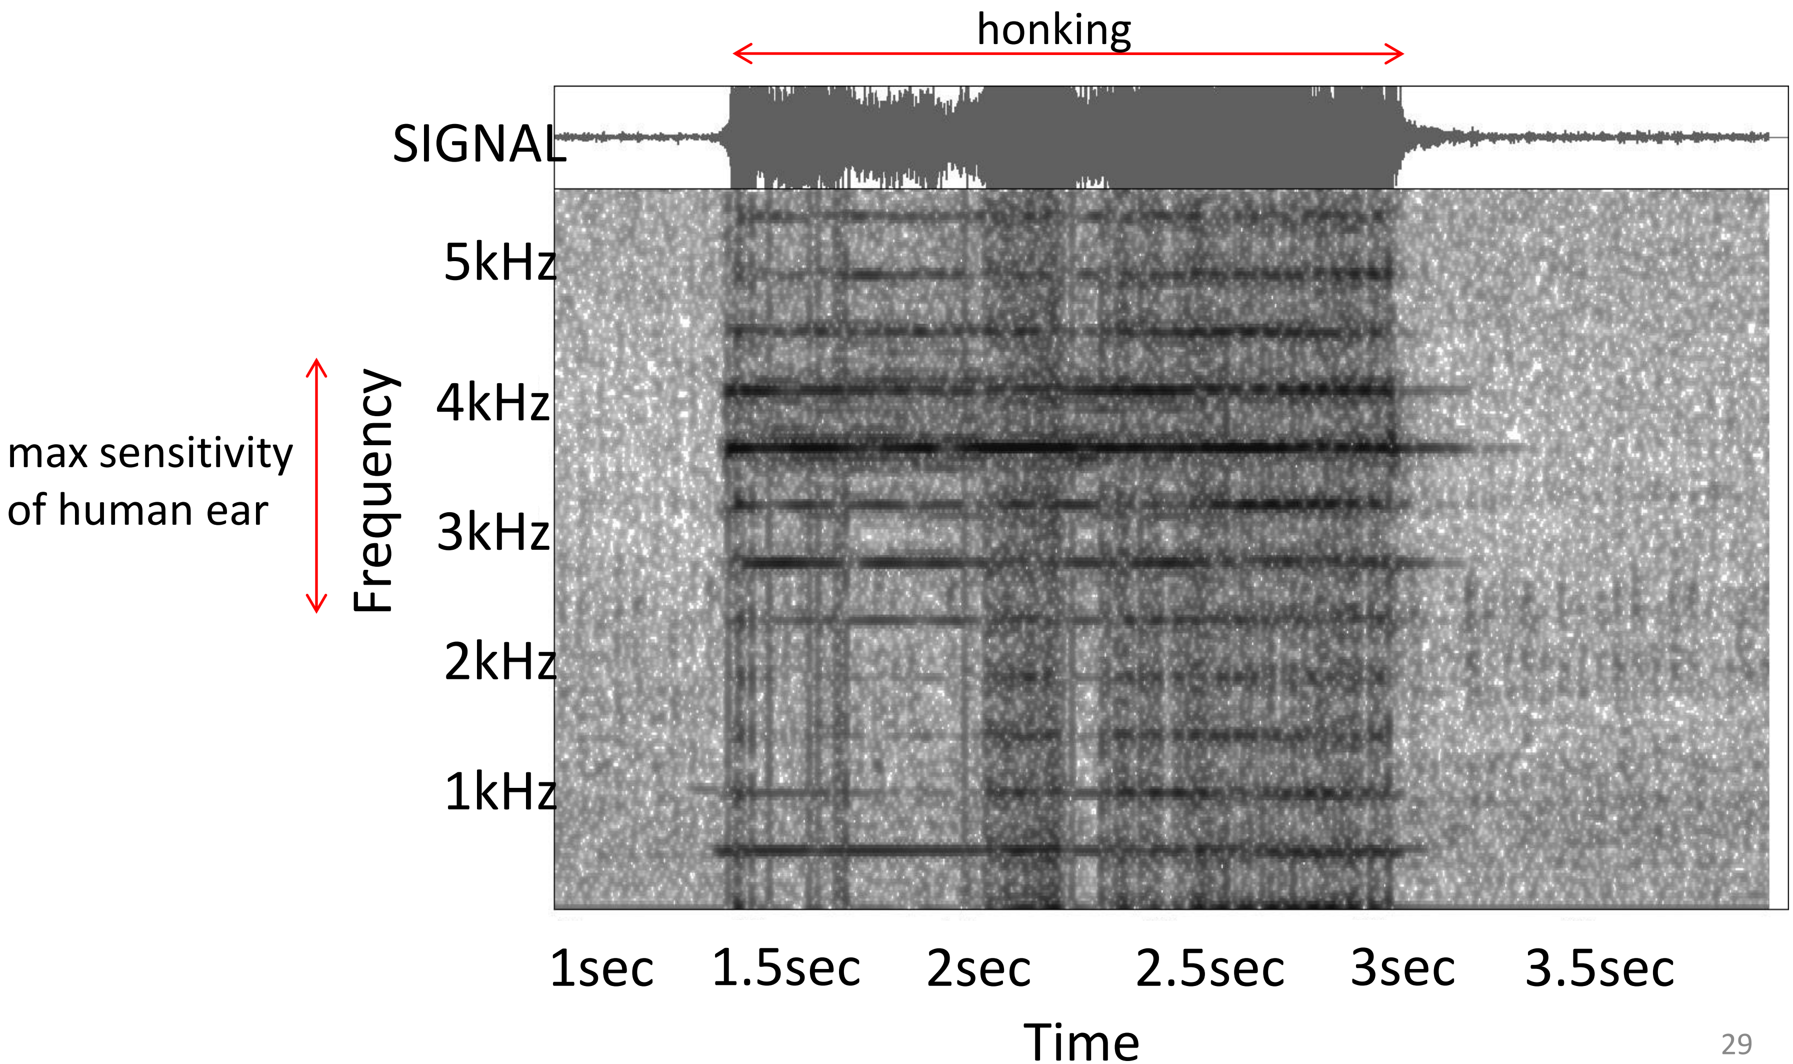

Honk Detection

- Efficient detector suitable for mobiles
  - discrete Fourier transform on 100 ms of audio
  - look for spikes in the 2.5-4 kHz range
  - spike: instantaneous > 10x mean
- Performance: 5.8% of CPU on the HP iPAQ
- Accuracy:
  - false negative rate:

Phone	Exposed	Enclosed
HP iPAQ	19%	51%
iMate KJAM	0%	23%

  - false positive rate: negligible in typical traffic conditions
    - sirens, alarms, …
    - chirping of bird!

Exposed vs. Enclosed Vehicles

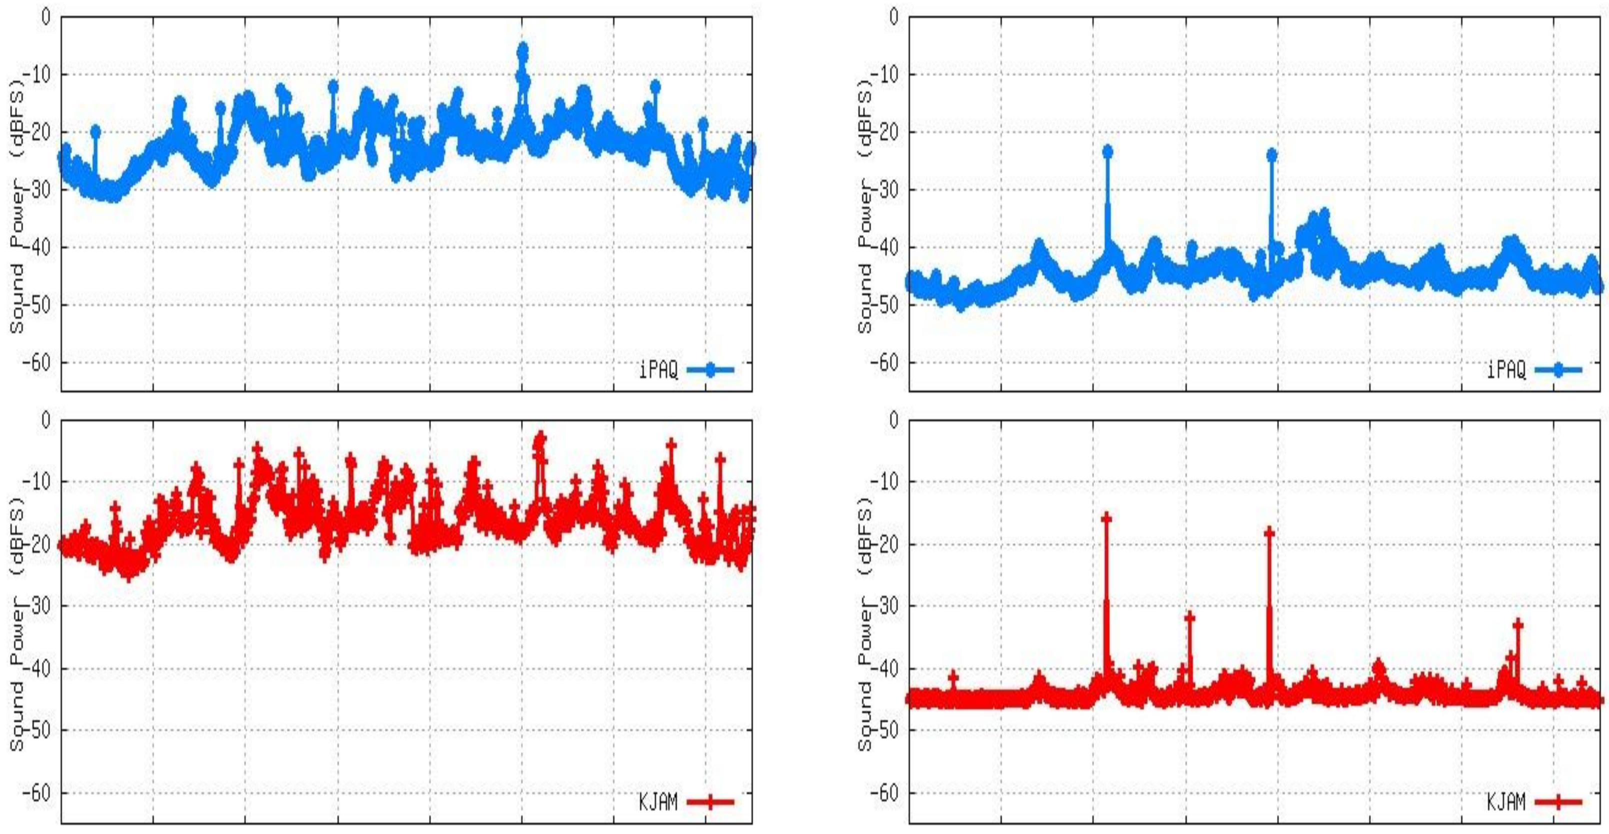

Outside car

Inside car

Neighbourhood comparison to avoid absolute thresholds

Triggered Sensing

- Use cheap sensors to trigger the activation of expensive sensors when needed

- Examples:
  - Traffic chaos: accelerometer info to trigger microphone
  - Localization: GSM tower info to trigger GPS
  - …

Integration with Maps

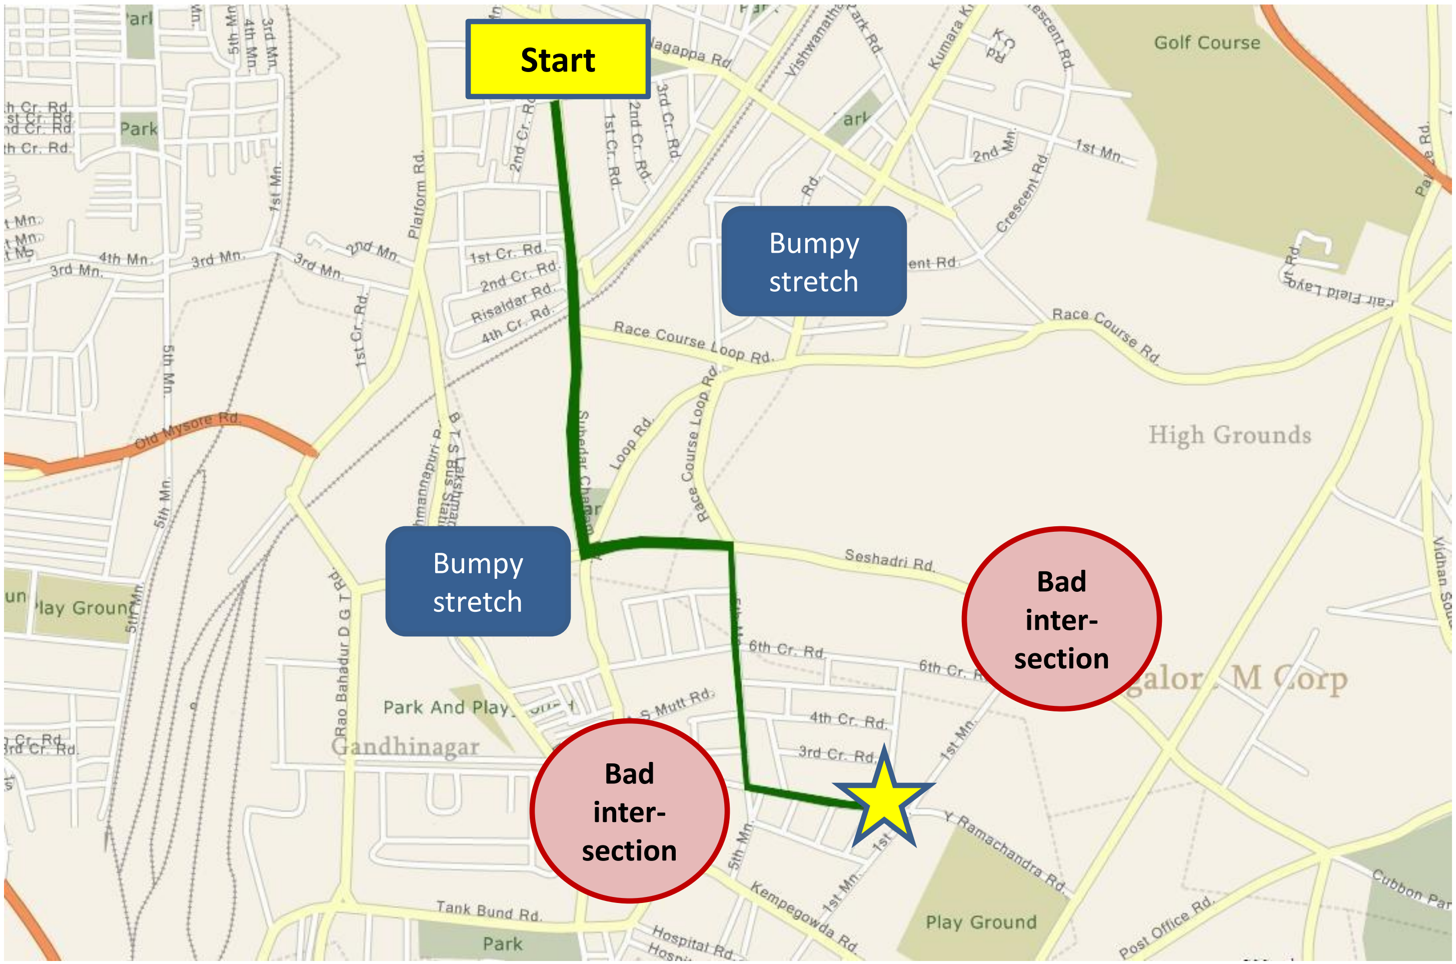

Find least stressful route

PRISM: Platform for Remote Sensing with Mobiles

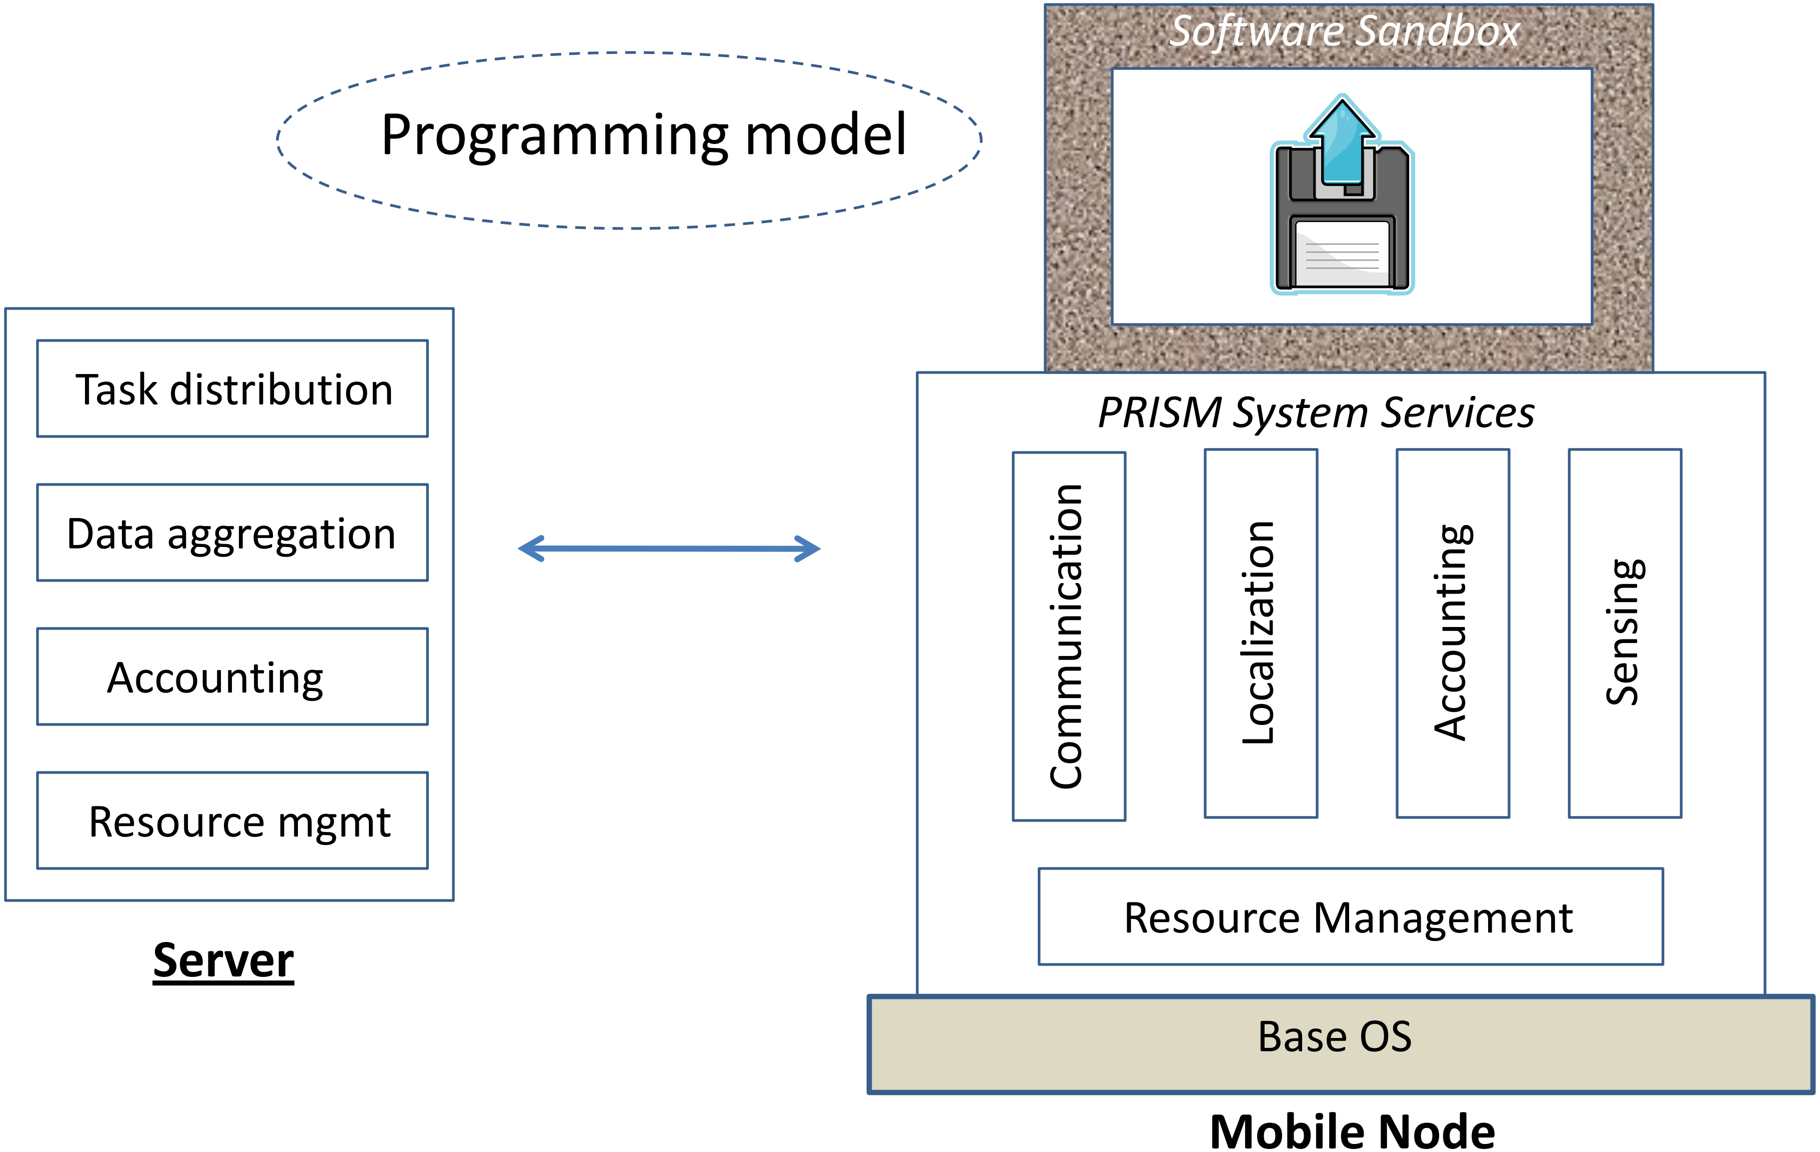

Related Work

- GPS-based vehicle tracking
  - OnStar, Surface street traffic estimation (MobiSys'07)
- Infrastructure for traffic monitoring (e.g., cameras)
  - SmartTrek, Busview
- Traffic estimation using tower-based tracking of mobiles
  - BTIS
- Dedicated vehicle based sensors
  - CarTel, Pothole Patrol

Summary

- Diversity of road and traffic conditions => need to go beyond GPS-based monitoring
- Nericell: rich monitoring of road and traffic conditions using smartphones
- http://research.microsoft.com/research/mns

Microsoft Research India

- Established in 2005
- Goals:
  - high-quality research
  - internal tech transfer
  - external collaboration/service
- ~50 full-time staff
  - researchers, post-docs, assistant researchers, software engineers
  - visiting researchers, interns
- Seven areas of research
  - algorithms
  - crypto, security, and applied math
  - digital geographics
  - mobility, networks, and systems
  - multilingual systems
  - rigourous software engineering
  - technologies for emerging markets

Mobility, Networks, And Systems Group

<u>Research Staff</u>

Bhavish Aggarwal
Assistant Researcher (IIT Mumbai)
Network diagnostics, wireless networking

Ranjita Bhagwan
Researcher (IIT Kharagpur →UCSD → IBM Research)
Network management, distributed systems

Tathagata Das
Assistant Researcher (IIT Kharagpur)
Network management, P2P systems

Vishnu Navda
Researcher (Bangalore U→MS IDC → Stonybrook)
Wireless networking, mobile systems

Venkat Padmanabhan (Research Manager)
Sr Researcher (IIT Delhi →Berkeley → MSR-Redmond)
Mobile systems, network management

Ram Ramjee
Sr Researcher (IIT Chennai →UMass → Bell Labs)
WAN acceleration, wireless networking

Group formed in Spring 2007

<u>Visiting Researchers from academia</u>

Kameswari Chebrolu (IIT Mumbai)
Bhaskaran Raman (IIT Mumbai)
Geoff Voelker (UC San Diego)

<u>Interns</u>

9 interns in 2007 (5 from India, 4 from the U.S.)

9 interns in 2008 (4 from India, 5 from the U.S.)

<u>"Graduated" Assistant Researchers</u>

Ganesh Ananthanarayanan (→Berkeley)
Lenin Ravindranath (→ MIT)
Prashanth Mohan(→ ??)

Ongoing and Completed Research

Network Management and Performance

NetPrints: Home Network Configuration Management (**USENIX NSDI 2009**)

CoCoNet: Content Compression in Networks (**ACM SIGMETRICS 2009**)

Mobile and Sensor Systems

SPACE: Lightweight Peering (**ACM HotNets 2006**)

Nericell: Rich monitoring of Roads and Traffic Using Smartphones (**ACM SenSys 2008**)

PRISM: Platform for Remote Sensing with Mobiles

SixthSense: RFID-based Enterprise Intelligence (**ACM MobiSys 2008**)

Wireless Networking

COMBINE: Collaborative Downloading Using WLAN and WWAN (**ACM MobiSys 2007**)

Neighbourcast: Enabling Communication Among Nearby Clients (**ACM HotMobile 2008**)

Multicast in Wireless LANs

Smartphone energy modeling and optimization

Other

Insight: Distributed Systems Profiling (**ACM HotMetrics 2008**)

Defending Against Code Geometry Attacks

More info: http://research.microsoft.com/research/mns/