Nericell: Rich Monitoring of Road and Traffic Conditions using Mobile Smartphones

Prashanth Mohan, Venkata Padmanabhan, Ramachandran Ramjee.
Microsoft Research India, Bangalore

Presented by: Philip Shibly

Topics

- What is Nericell?
- Framework
- Acceleration
- Orientation and Reorientation
- Validation
- Road and Traffic Conditions
    - Bump Detection
    - Stop-and-Go
    - Audio
- Localization
- Energy

What is Nericell?

- Smartphone system that monitors road and conditions in developing regions.
- Tailored to complex traffic flow and degraded roadways.
- Avoids usage of infrastructure or area networks.
- Tested in Bangalore, India.

Framework

- How to sense congestion?
- How to sense road conditions?
- How to determine location?
- How to be efficient?

Congestion: accelerometer

- Monitors acceleration to determine braking and congestion.
- Uses the 3-axis accelerometer but must account for orientation.
  - Does not make assumption that everyone has their phone in the same orientation in the car.
  - Developed an algorithm for determining orientation and then virtually reorienting the accelerometer.
  - Monitors for user-interaction with the phone and neglects those readings.
- Validates the reorientation and develops heuristics to detect bumps, potholes and braking under certain speed conditions.

Acceleration

- Defines orthogonal axes with respect to the phone as (x, y, z).
- Defines orthogonal axes with respect to the vehicle as (X, Y, Z).
- If they are equal, we say that the phone is *well-oriented*. If not, it is *disoriented.*
- Accelerometer readings relate to the frame of references as (a_x, a_y, a_z) and (a_X, a_Y, a_Z)
  -
  - If the accelerometer is well-oriented, these are equal.
- A DC accelerometer is capable of measuring static acceleration as 1G in the downward direction.

Determining Orientation

- Applying rotations about the X, Y, and Z axes are computationally intensive and require the CPU to perform costly trig functions.
- Instead, Euler angles are used and any orientation of the accelerometer can be represented by a *pre-rotation* of Φ_{pre} about the Z axis, a *tilt* of θ_{tilt} about the Y axis, and then a *post-rotation* of Ψ_{post} again about the Z axis.

Estimating Pre-Rotation and Tilt

- When the accelerometer is stationary, the only effect on it is gravity (1G) along Z. By their math, this is the only sampling that is needed to calculate pre-rotation and tilt.
- Instead of waiting for the vehicle to come to a stop in order to estimate, a rolling 10-second averaging window of the accelerations are taken to determine the median values. Since any momentary bumps would average out, then the pre-rotation and tilt would be able to be continuously calculated.

Estimating Post-Rotation

- Braking provides a significant change in acceleration that is orthogonal to Z, which is needed to estimate post-rotation.
- However, GPS is needed to sample this change and is expensive compared to the prior estimations.
- So they monitor pre-rotation and tilt for any noticeable change, then turn GPS on to estimate post-rotation.

Validating the Estimations

- How successful is reorientation?
- A well-oriented accelerometer is compared to a disoriented one and the cross-correlation is taken to determine the effectiveness of reorientation. The cross-correlation is then also compared against the cross-correlation of two well-adjusted accelerometers.

Inferring Road and Traffic Conditions

- Brake Detection is performed by monitoring a_x
- The mean is computed over a sliding window N seconds wide and if the mean exceeds a threshold T, then a braking event is detected.
- To establish the ground truth, GPS or CAN is used and a threshold of 1m/s^2 is over a duration of 4 seconds is used to quantify a braking event.
- By using T=.11G-.12G in experiments, that equates to 10-20% more conservative than the ground truth.

Brake Detection Results

	False Negatives		False Positives	
Accelerometer (threshold T (g))	Rate	Change in Speed avg(max)	Rate	Change in Speed avg(max)
ACL-1 (T=0.11) ACL-1 (T=0.12)	4.4% 11.1%	15(16) 16(18)	22.2% 15.5%	12(10) 12(9)
ACL-3 (T=0.11) ACL-3 (T=0.12)	4.4% 11.1%	15(16) 16(18)	31.1% 17.7%	12(9) 12(9)

- ACL1 and ACL3 agree quite well
- False-positives seem high, but they correlate to deceleration events at a slightly lower rate than the heuristic.
- False-negatives are low, and fall within the GPS localization error. These can be avoided by using the CAN.

Differentiating Between Stop-and-Go and Pedestrian Traffic

- Main observations are:
  - the amplitude of the surges in acceleration
  - The frequency of the surges
- No false positives or negatives were captured using the same heuristics from prior experiment.
- Different pedestrian traces did not produce any false positives.

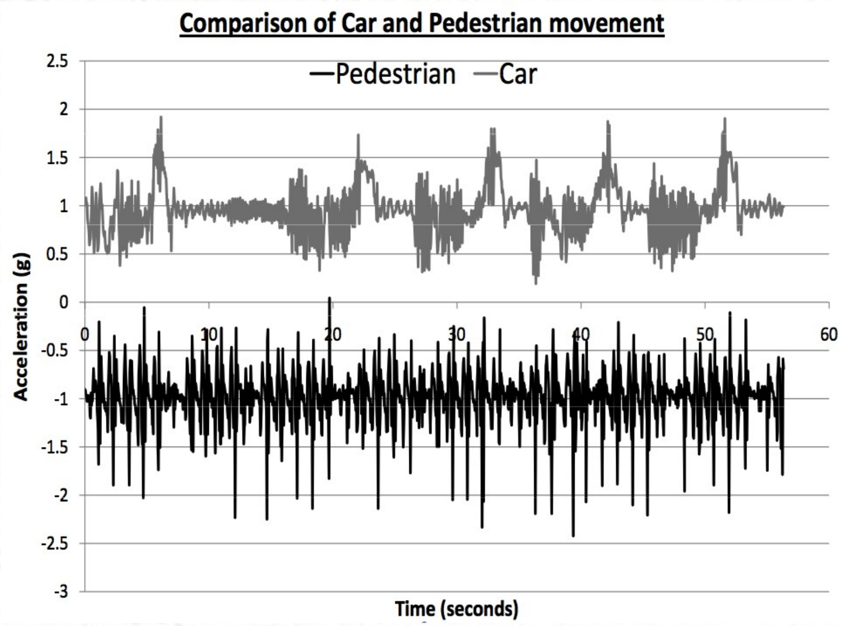

Bump Detection

- Difficult to implement for a number of reasons
  - How do you establish the ground truth?
    - Manual annotation
  - Accelerometer signal is of very short duration and different magnitude at different speeds.
    - Implement two sets of heuristics for fast and slow speed. z_{sus} and z_{peak}
      - z_sustaining looks for longer duration dips below a lower threshold.
      - z_peak looks for quicker dips below a higher threshold
  - Needs a training set to develop heuristics

Bump Detection Results

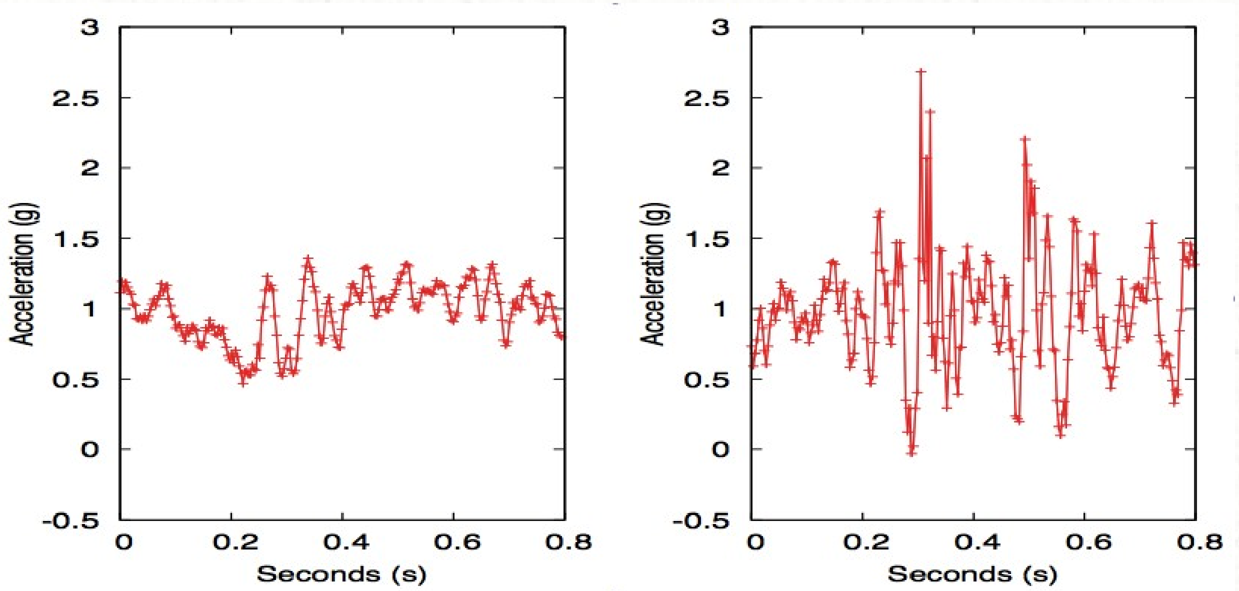

Low Speed High Speed

Bump Detection Results

Detector	Accel.	Speed < 25kmph		Speed ≥ 25kmph	
		FN	FP	FN	FP
BUMPY road		40 bumps total		4 bumps total	
z-sus	ACL-1	**25%**	**5%**	50%	0%
	ACL-2	**30%**	**0%**	25%	0%
	ACL-3	**23%**	**5%**	0%	50%
z-peak (1.45)	ACL-1	28%	15%	0%	125%
	ACL-2	20%	5%	0%	125%
	ACL-3	30%	10%	0%	200%
MIXED road		62 bumps total		39 bumps total	
z-sus	ACL-1	**29%**	**8%**	18%	80%
	ACL-3	**37%**	**14%**	0%	136%
z-peak (1.45)	ACL-1	35%	6%	5%	197%
	ACL-3	65%	21%	3%	49%
z-peak (1.75)	ACL-1	90%	0%	**51%**	**3%**
	ACL-3	83%	0%	**41%**	**8%**

- Both detectors tuned for low false-positives (<10%)
- False-negatives are high because of the difficulty of establishing the ground truth.

Honk Detection

- Implements a simple heuristic approach for both exposed and enclosed vehicles.
- Detector implements a discrete Fourier transform on 100ms audio samples and looks at spikes in the frequency domain.
- By observation, they choose to implement the heuristic by looking for two spikes in frequency with one required to be within 2.5kHz to 4.0kHz range.
- Ground truth is by manual annotation.

Honk Detection Results

Phone	Exposed vehicle		Enclosed vehicle	
	FP	FN	FP	FN
KJAM (T=5)	38%	0%	8%	15%
KJAM (T=7)	0%	0%	0%	23%
KJAM (T=10)	0%	19%	0%	54%
iPAQ (T=5)	19%	4%	0%	19%
iPAQ (T=7)	0%	8%	0%	50%
iPAQ (T=10)	0%	27%	0%	81%

- With a large enough spike to avoid false-positives, the detector performs better in exposed vehicles due to a higher received power.
- The varying sensitivity between phones produces a different number of honks detected within the same sample.
- A high spike threshold protects against false-positives but doesnt eliminate them.

Localization

- GPS and WiFi are high energy users.
- GSM is much more efficient when using signal strength localization algorithms. Although, it relies on dense network of towers that is ironically seen in dense developing cities. So naturally GSM a good option.

Item	**Median error**	90^{th} **percentile error**
Localization Distance	117m	660m
Absolute speed	3.4 kmph	11.2 kmph
Relative speed	21%	70%

- Testing shows high relative error due to the slow speeds of vehicles in dense cities. The absolute accuracy provides better information, as you can distinguish between slow moving cars and faster ones.

Energy Costs

	Power (mW)	% Time active
Audio	223.2	5
Honk Detection	63.3	5
GPS	617.3	10
Reorientation of accel values	20.9	100
Bump & Brake Detection	9.3	100
Accelerometer	1.65	100

- Benchmarks ran on a HP iPAQ for a 4-hour duration.
- Triggered sensing allowed for efficient sensors to turn on costly ones only when needed. GSM and accelerometer stays on and triggers GPS and microphone when needed.
- Showed a decrease in battery life of only 9.7%

Energy Usage for Various Activities

Mode	Life Time Includes Phone Idle	Power (mW) For given mode only
Phone Idle	24h 18m	182.7
Bluetooth (BT) Idle	22h 13m	17.1
BT Device Inquiry	10h 46m	229.5
BT Service Discovery	7h 53m	380.0
WiFi Idle	4h 39m	771.8
WiFi Beacon (Sending)	4h 36m	782.0
WiFi Scan (Receiving)	2h 59m	1298.8
GPS	5h 32m	617.3
Microphone	10h 54m	223.2
Accelerometer (per spec.)	24h 5m	1.65
Accel. with Bluetooth	19h 56m	40