After Motes and Multihop: Mobile Phones and the Global Mobile Sensor Network

Andrew T. Campbell, Dartmouth College

Wireless sensor networks have driven many great innovations over the last decade - represents a very active area of on-going research

The mote have been a superb platform for research

But, challenges remain

- Not ubiquitous
- Energy problem
- Scaling (cost and performance) problem
- Event unpredictability
- Don't have economy of scale

Importantly, sensor networks don't impact our everyday lives. Why?

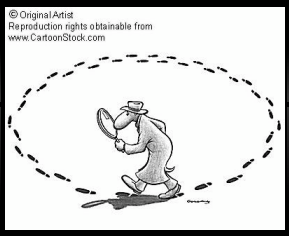

People are out of the loop

That's all just changed because of this …

Embedded sensors:
- 3-axis accelerometer
- Proximity sensor
- Microphone
- Camera
- GPS
- Bluetooth

Solves many outstanding problems:
- ubiquitous, energy problem, has economy of scale, scales, etc.

… or this

Nokia 6210 Navigator

Embedded sensors:
- 3-axis accelerometer
- Digital compass
- Microphone
- Camera
- GPS
- Bluetooth

or, in the future ..

The cool green “emotional” phone

Embedded sensors:
- 3-axis accelerometer
- Proximity sensor
- Digital compass
- Pollution/air quality sensor
- GSR “emotion sensor”
- RFID/NFC
- Microphone
- Camera
- GPS
- Bluetooth

You’ll be able to answer lots of questions and build new cool applications with these phones

What is my personal air quality like today? Or, the air quality of my neighborhood, school, town, or city?

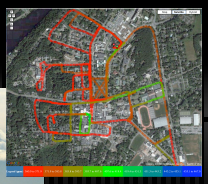

How stressed is the city this morning?

How do social “conversation” networks evolve?

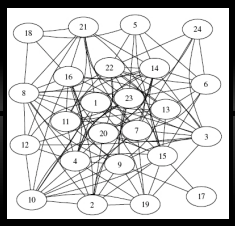

Work on audio networks by Tanzeem Choudhury (Dartmouth)

Where are my friends and what are they doing right now?

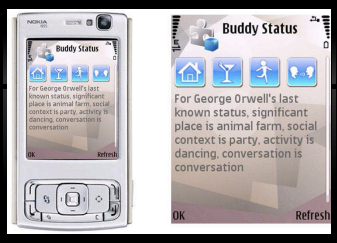

I know what you are thinking

Sounds like an Orwellian nightmare!

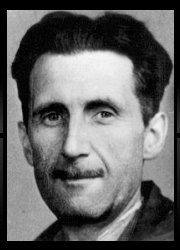

Big challenges: trust, privacy, security are critical issues (David Kotz, Dartmouth)

Now imagine 1 billion "sensor enabled mobile phones" scattered across the planet

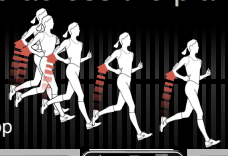

people are in the loop

This will lead to …

I know what you are thinking

On second thoughts, phones are getting cheaper and we have a ready supply of adventurous grad students ;-)

We are at a crossroads

My position for this talk

The mobile phone (and not the mote) will serve as the main platform for sensing innovation over the next 5 years.

Your mobile phone will sense your surroundings, learn your behavior (what you do, where you go and how you interact with people and your environment), and help you navigate your day.

Collectively, mobile phones will form societal scale sensor networks in support of community, urban, and global sensing applications and problem solving.

You don't have to believe me, just wait for it to evolve

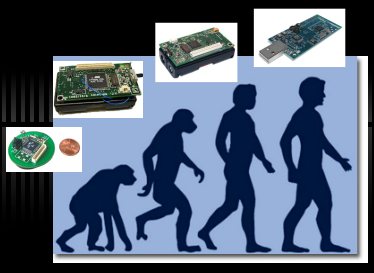

The new "mote" (minus the multihop problems)

We started in 2005 to study people-centric sensing

Characteristics of existing mote networks

- Small-scale, short-lived, mostly-static
- Application-specific
- Multi-hop wireless
- Very energy-constrained
- Mobility not an issue or driving factor
- People out of the loop

Characteristics of People-Centric Sensing

- Large-scale, long-lived, mostly-mobile
- Application-specific
- Multi-hop wireless
- Very energy-constrained
- Mobility not an issue or driving factor
- People out of the loop

Characteristics of People-Centric Sensing

- Large-scale, long-lived, mostly-mobile
- Application-agnostic
- Multi-hop wireless
- Very energy-constrained
- Mobility not an issue or driving factor
- People out of the loop

Characteristics of People-Centric Sensing

- Large-scale, long-lived, mostly-mobile
- Application-agnostic
- No multi-hop wireless
- Very energy-constrained
- Mobility not an issue or driving factor
- People out of the loop

Characteristics of People-Centric Sensing

- Large-scale, long-lived, mostly-mobile
- Application-agnostic
- No multi-hop wireless
- Periodic recharging
- Mobility not an issue or driving factor
- People out of the loop

Characteristics of People-Centric Sensing

- Large-scale, long-lived, mostly-mobile
- Application-agnostic
- No multi-hop wireless
- Periodic recharging
- Mobility is a driving factor
- People out of the loop

Characteristics of People-Centric Sensing

- Large-scale, long-lived, mostly-mobile
- Application-agnostic
- No multi-hop wireless
- Periodic recharging
- Mobility is a driving factor
- People in the loop

Characteristics of People-Centric Sensing

- Large-scale, long-lived, mostly-mobile
- Application-agnostic
- No multi-hop wireless
- Periodic recharging
- Mobility is a driving factor
- People in the loop
- Security and privacy of data critical

People-centric sensing application domains

Andrew Campbell, et al, "The Rise of People-Centric Sensing", IEEE Internet Computing, July/August 2008

I know what you are thinking

Static sensor networks don't scale to large areas, sensing coverage is costly, performance doesn't scale either, and events are unpredictable in time and space.

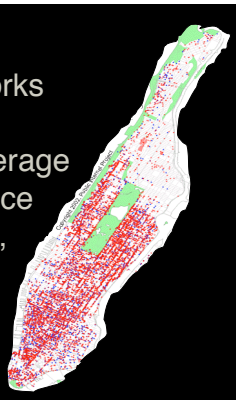

Public sensing gains scalability and sensing coverage by using people opportunistically as mobile sensors

The beauty is that the infrastructure already exists (i.e., people and the global cellular network)

People-centric sensing is based on an "opportunistic sensing paradigm" and an "interaction model" that captures interaction between people, and, between people and their surroundings

Emerging sensing paradigms

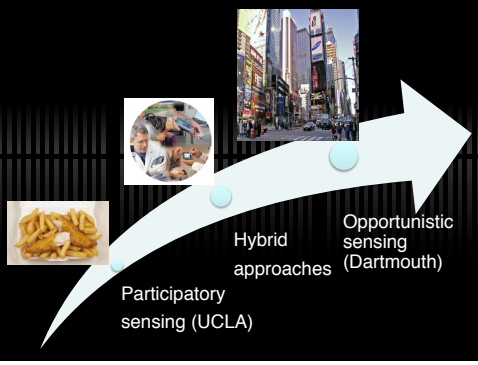

Next frontier in sensing is "people-centric" urban sensing

Ron Fricke, timescape is a day in the life of a city (edited version)

Remainder of my talk

- Three people-centric sensing applications we developed
  - BikeNet (personal/public sensing)
  - CenceMe (social sensing app)
  - SoundSense (personal sensing app)
- Need for open sensing/comms software for mobile phones
- Wrap up

The BikeNet system – sensor bikes

BikeNet - sensor bikes

Shane Eisenman, et al, "The BikeNet mobile sensing system for cyclist experience mapping", ACM SenSys '07

We can answer many questions from sensor data

- How fit are you?
- Many cars along the route?
- What was the air quality and noise like?
- Lots of trivia: slopes, coasting, braking, working hard
- Overall health and performance along the route
- How did you compare to your buddies, community?

CO2 Map of Hanover

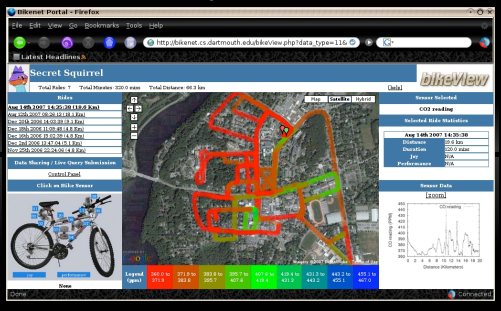

http://bikenet.cs.dartmouth.edu

Lots of cars on that route?

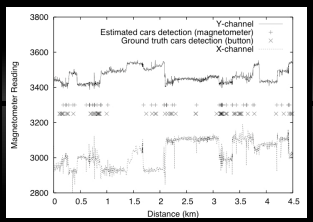

I know what you are thinking

How do you do ground truth?

Performance index

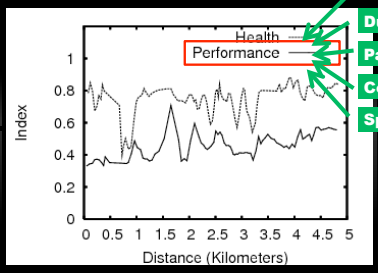

Distance
Duration
Path Slope
Coasting
Speed

Health index

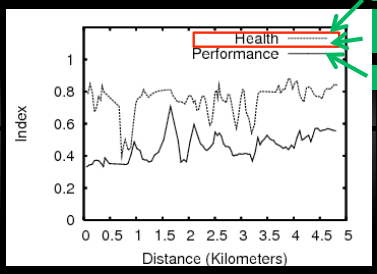

Noise
Traffic Density
CO_2 Level

Many lessons learnt

Debugging on the go is hard

CenceMe - sensing meets mobile social networks

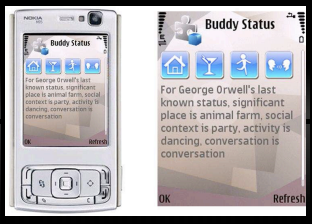

Emiliano Miluzzo et al., "Sensing Meets Mobile Social Networks: The Design, Implementation and Evaluation of the CenceMe Application", ACM SenSys 2008

sensing with cenceme

cencing with cenceme

cencing with cenceme

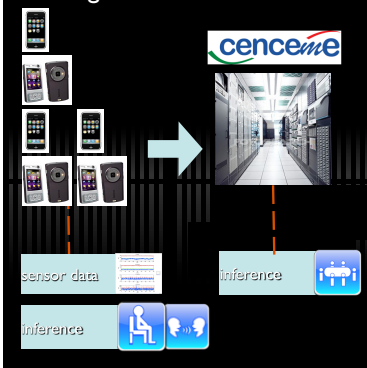

cencing with cenceme

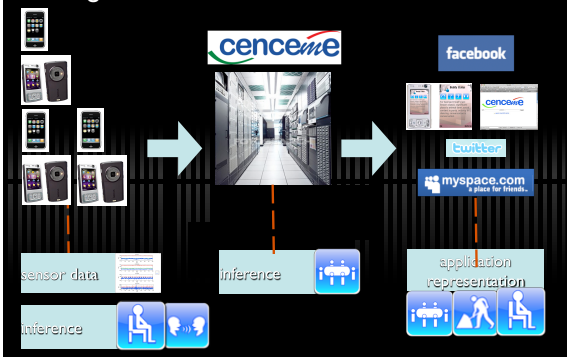

supported inferences: sensing presence

activity

supported inferences: sensing presence

activity

social context

supported inferences: sensing presence

activity

social context

significant places

supported inferences: sensing presence

activity

social context

significant places

behaviour

greeny | nerdy | party animal | cultured | healthy

CenceMe demo

Sensor presence is published on Facebook

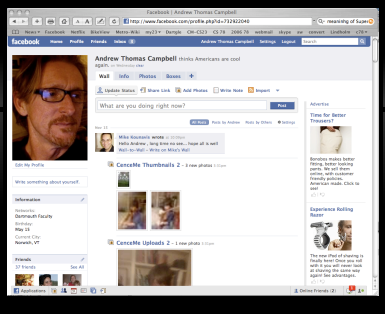

Classifying activity

sitting

standing

walking

running

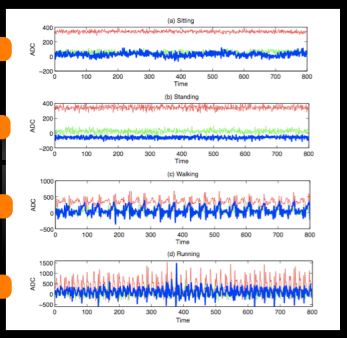

Activity classifier confusion matrix

	Sitting	Standing	Walking	Running
Sitting	0.6818	0.2818	0.0364	0.0000
Standing	0.2096	0.7844	0.0060	0.0000
Walking	0.0025	0.0455	0.9444	0.0076
Running	0.0084	0.0700	0.1765	0.7451

Supervised learning approach

Differentiated between sitting and standing is hard

Custom sensing hardware (e.g., Intel's MSP) can do better but these results are from the Nokia N95

Classifying conversation

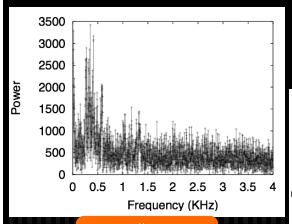

talking

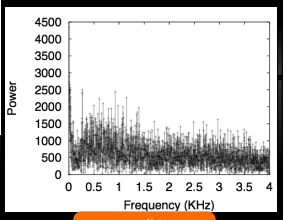

no talking

Classifying talking/ non-talking

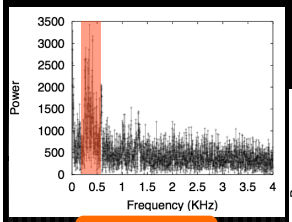

talking

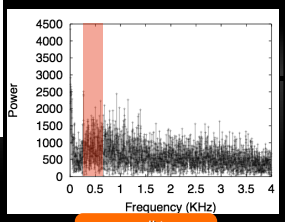

no talking

Classifying talking/ non-talking

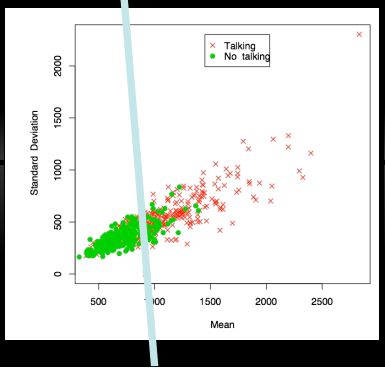

Conservation classifier confusion matrix

	Conversation	Non-Conversation
Conversation	0.83.82	0.1618
Non-Conversation	0.3678	0.6322

Design decision of 2/5 talk primitives to get into conversation and 4/5 to get out – more conservative

Poor performance for non conservation results because people aren't talking but others nearby are.

Duty-cycling on the phone: energy vs fidelity

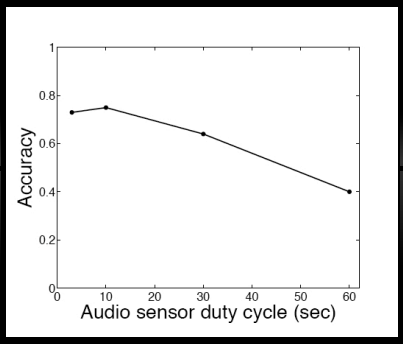

Deployment, user study – 22 people over three weeks lets look at the portal

Results: Location and activity

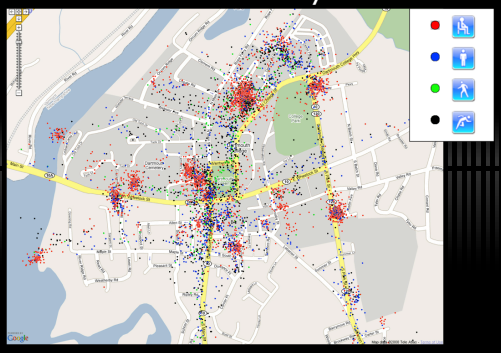

Results: Location and activity

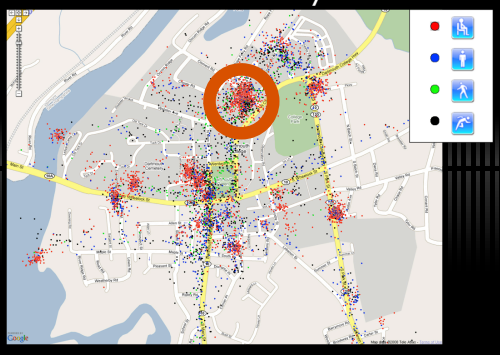

Results: Location and activity

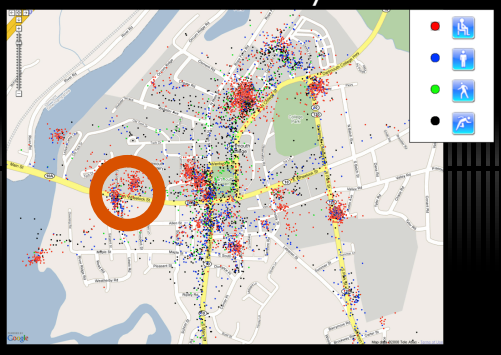

Results: Location and activity

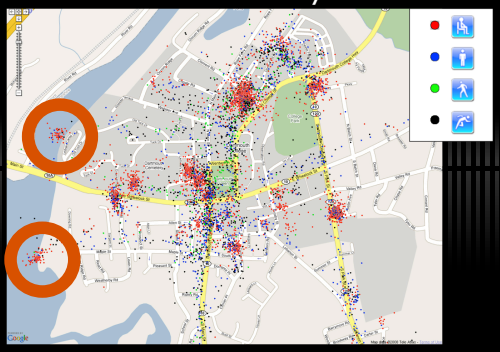

Observations from the 22 user study

Willingness to share sensitive presence information with friends when 'sensor off' button is clearly available

Valued access to presence information from the phone

Enjoyed comparing themselves with friends and the population

"CenceMe made me realize I'm lazier than I thought and encouraged me to exercise a bit more"

cenceme

iphone edition

Available on the iPhone App Store

Injecting sensor presence into second life

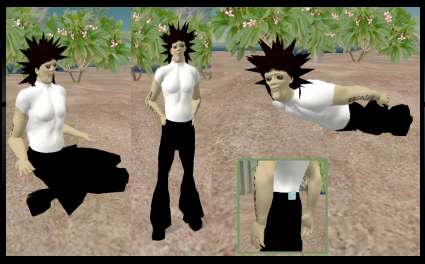

SoundSense – listen to the "significant sound events" in your life

SoundSense

Hong Lu, Tanzeem Choudury, et al., "Scalable Sound Sensing for People-Centric Sensing using Mobile Phones", ACM MobiSys 2009 (provisionally accepted)

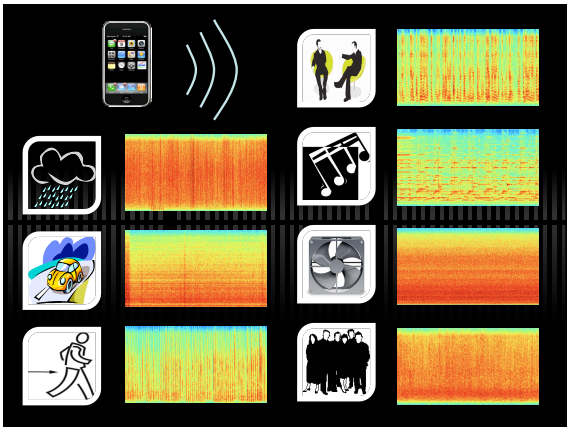

System design for phone

Real-time sound sensing system with privacy
- No backend server
- No raw audio stored on phone only features, models

Computationally efficient feature extraction and classification
- Features provide good discrimination but remain low cost
- Features must be robust to noise and energy variance
- Multi-stage classification structures

Combination of supervised learning and semi-supervised learning technique

SoundSense system

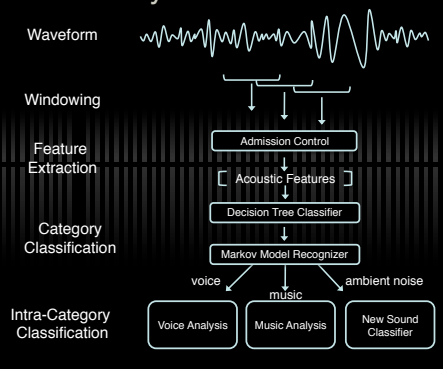

Classification confusion matrix

	Ambient Noise	Music	Speech
Ambient Noise	0.9159	0.0634	0.0207
Music	0.1359	0.8116	0.0525
Speech	0.0671	0.1444	0.7885

Accuracy of the decision tree classifier

	Ambient Noise	Music	Speech
Ambient noise	0.9494	0.0402	0.0104
Music	0.0379	0.9178	0.0444
Speech	0.0310	0.0657	0.9033

Accuracy of the markov model recognizer output

Daily diary app

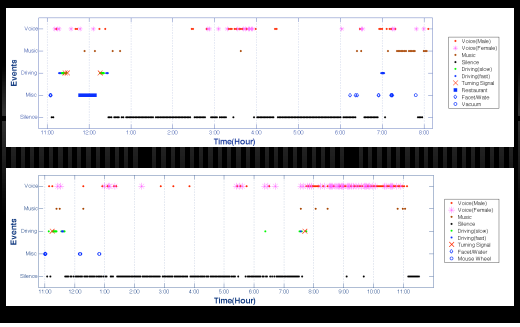

Music tagger app

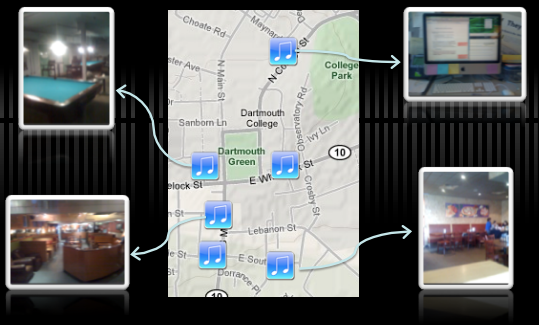

If the phone is the new mote – then, do we need a new TinyOS?

Yes – not an OS, but some libraries or dare I say, sensor phone middleware.

Toward Sensor PhoneWare

Supporting continuous sensing significant challenge
Many open challenges

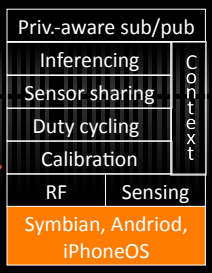

Growing interest in sensing on mobile phones

- Applications
  - WatchMe, iCAMS, PEIR, Nericell
- Sensing with mobile phones
  - UCLA, UIUC, Intel, Nokia, Microsoft, Motorola, UW, Duke, start ups: e.g., Sense Networks
- Human activity inferencing
  - MIT, Intel, UW
- Workshops
  - UrbanSense 08, MODUS 08

Finally, ……

My title is a little loaded, isn't it?

The title implies that the phone is the “new mote”, multihop is dead, and that the sensor network community should now direct its intellectual energy toward programming phones not motes.

Why do this?

Today you can ship your cool new sensor app to thousands, perhaps millions of phones.

Today you can build a global “sensor network” of thousands, perhaps millions of “nodes” if you have a really good idea.

Interesting problems will emerge. Your ideas can have significant impact.

Tomorrow? You’ll be able to reach billions of phones instantly forming societal scale sensor networks.

PS You don't need any infrastructure to do this! ☺

What's MetroSense?

It's about a societal scale sensor network built on everyday mobile phones

Metro Sense

http://metrosense.cs.dartmouth.edu/

Thanks for listening!

- Project page, papers, etc:

http://metrosense.cs.dartmouth.edu

- Thanks to many people's contributions

http://metrosense.cs.dartmouth.edu/metro-people.html

- Sponsors