

MINING SMARTPHONE MOBILITY DATA

Mobile technology overview

Spiros Papadimitriou, Tina Eliassi-Rad,
Rutgers University Northeastern
Katharina Morik, Dimitrios Gunopulos
TU Dortmund University of Athens



Mobile OSes

- **iOS**
- **Android**
- Windows Phone
- Blackberry
- Symbian
- (Ubuntu, Mozilla, OpenMoko, ...)

Why do you care?

- What is possible?
- What might be possible?
- What is not possible?

A basic understanding of the realities helps make realistic assumptions about

- Collection
- Transmission
- Processing

Overview

- **Sensors & location API**
- Network connectivity
- Power
- Accounts (“identity”)
- Mobile app basics


Smartphone sensors (Android)

Sensor	Type	Avail. since
Accelerometer	HW	1.5
Light	HW	1.5
Magnetic field	HW	1.5
Proximity	HW	1.5
Temperature	HW	1.5 (4.0)
Orientation	SW	1.5
Gyroscope	HW	2.3
Pressure	HW	2.3
Gravity	SW/HW	2.3
Linear acceleration	SW/HW	2.3
Rotation vector	SW/HW	2.3
Relative humidity	HW	4.0

- Hardware or software (“virtual”)
- iPhone: more standardized
- Android: greater variety, no minimum required, varied APIs across versions

Accelerometer

- Measures acceleration
 - Static (gravity) → orientation
 - Dynamic (linear motion)
- Example*: LIS33DLH (iPhone)
 - ~\$1.5 (DigiKey, bulk)
 - 0.7mW on / 0.03mW low power:
325 days / 20 years (1440mAh @ 3.8V)
 - $\pm 2g$ / $\pm 4g$ / $\pm 8g$ selectable range
 - 16-bit dynamic range
 - 0.5Hz ~ 1KHz sample rate
 - Simple interrupt generators (free fall, motion)



<http://www.memsjournal.com/2010/12/motion-sensing-in-the-iphone-4-mems-accelerometer.html>

RUTGERS Mining Smartphone Mobility Data 7

Accelerometer

Basics

- Measures force exerted on device (vector)
- Stationary device, lying flat:
 - Force preventing it from falling (opposite to gravity)
 - Zero force → Free fall
- Stationary device, after 45° rotation:
 - Same magnitude, but rotated

$a_x = a_y = 0$
 $a_z = g$
 $a_y = 0$
 $a_z = -a_x = g/\sqrt{2}$

RUTGERS Mining Smartphone Mobility Data 8

Accelerometer

Example code (1/3)

Android (1/2)

```

public class SensorActivity extends Activity
    implements SensorEventListener {
    private SensorManager mSensorManager;
    private Sensor mAccelerometer;

    @Override
    public final onCreate(Bundle savedInstanceState) {
        mSensorManager = \
            (SensorManager) getSystemService(Context.SENSOR_SERVICE);
        mAccelerometer = \
            mSensorManager.getDefaultSensor(Sensor.TYPE_ACCELEROMETER);
        mSensorManager.registerListener(
            this, mAccelerometer, \
            SensorManager.SENSOR_DELAY_NORMAL);
    }
    
```

[continued]

RUTGERS Mining Smartphone Mobility Data 9

Accelerometer

Example code (2/3)

Android (2/2)

```

[continued]

@Override
public final void onSensorChanged(SensorEvent event) {
    float ax = event.values[0];
    float ay = event.values[1];
    float az = event.values[2];
    ...
}

@Override
public final void onAccuracyChanged(Sensor sensor, int accuracy) {
    ...
}
} // SensorActivity
    
```

RUTGERS Mining Smartphone Mobility Data 10

Accelerometer

Shortcomings

- Cannot distinguish between gravity and acceleration
 - Impossible: "equivalence principle"
 - Solution: use low-pass filter to estimate gravity
- What if device simultaneously rotates & linearly accelerates?
 - Confused; need more data → gyroscope

RUTGERS Mining Smartphone Mobility Data 11

Gyroscope

Introduction

- Measures angular speed
 - Degrees per second (dps)
- Example*: L3G4200D (iPhone)
 - ~\$6.5 (DigiKey, bulk)
 - 18mW on / 4.5mW sleep (0.02mW off): 12.5 days / 50 days
 - ±250dps / ±500dps / ±2000dps range
 - 16-bit dynamic range
 - 100 / 200 / 400 / 800Hz sample rate
 - Temperature sensor (8-bit range)
 - Simple interrupt generator & FIFO

STMicro L3G4200D

<http://www.micropower.com/2011/01/10/1000-reading-in-the-iphone-4-micro-gyroscope.html>

RUTGERS Mining Smartphone Mobility Data 12

Gyroscope

Basics

- Measures angular speed of rotation
 - Represented by numbers for each axis (but: rotation axis is different)
 - Right-hand rule
- Integrate to obtain orientation
 - ...with care, since non-collinear rotations are not commutative

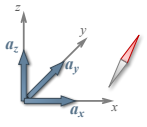
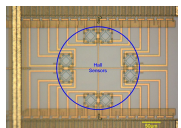
$\omega_x = \omega_y = \omega_z = 0$
 $\omega_x = \omega_z = 0$
 $\omega_y = 45^\circ/\text{sec}$

RUTGERS Mining Smartphone Mobility Data 13

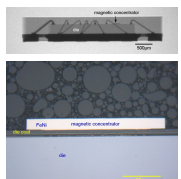
Magnetometer (compass)

Introduction

- Measures direction and magnitude of (Earth's) magnetic field
- Example*: AK8973/5 (iPhone)
 - <\$1 (Wikipedia), ~\$2 (AliExpress; non-bulk)
 - 20mW sensor on / 3mW @10Hz: 11 days / 76 days

AKM AK8975



<http://www.micromajournal.com/2011/02/motion-sensing-in-the-iphone-4-electronic-compass.html>

RUTGERS Mining Smartphone Mobility Data 14

Inertial measurement & navigation

- Summarizing: 9 DoF (degrees of freedom) available

• 3-axis accelerometer	Linear acceleration (translation)	} Dead reckoning
• 3-axis gyro	Angular velocity (rotation)	
• 3-axis magnetometer	Global reference (Earth)	
- Combine in software for accurate:
 - Position, velocity, and acceleration (linear and angular)
 - Dead reckoning

RUTGERS Mining Smartphone Mobility Data 15

Inertial measurement

Gravity estimation

Low-pass filter on acceleration data;
e.g., on Android Gingerbread:

$$g_t = \lambda_0(a_t + a_{t-2}) + \lambda_1 a_{t-1} - \kappa_1 g_{t-1} - \kappa_2 g_{t-2}$$

Basic operation, also used for:

- Linear acceleration estimate $a_t - g_t$
- Rotation vector (orientation wrt. magnetic north)

RUTGERS Mining Smartphone Mobility Data 16

Inertial navigation

Dead reckoning

- Integrate gyro to obtain orientation
- Use accelerometer and gyro (orientation) data to estimate linear acceleration
- Doubly integrate acceleration to obtain position change

- Errors accumulate over time ($\sim t^3$)
- Error depends on sampling rate
- How accurate is it?
- "Pro" (air navigation) answer:
 - GPS: better than 9m
 - Inertial: ~650m after one hour

RUTGERS Mining Smartphone Mobility Data 17

Other sensors

- Proximity

Less common:

- Thermometer
 - But: gyroscope & compass often output temperature as well
- Light intensity
- Pressure (barometric) → altitude

Future trend:

- Independent motion / sensor processors
 - Already on iPhones for a couple of years

RUTGERS Mining Smartphone Mobility Data 18

Location APIs

- Low-level location providers:
 - GPS
 - WiFi
 - Cell tower
 - ...
- Mid-level:
 - Fused location providers
- Higher-level:
 - Geo-fencing
 - Activity recognition

RUTGERS Mining Smartphone Mobility Data 19

Location APIs

Location providers (low-level)

- GPS, WiFi, cell tower, ...
- Differ in:
 - Accuracy
 - Availability / freshness
 - Power consumption
- Listen for location updates
- Choose how to update location estimate

Timeline events: Application starts, Listen for GPS and Network updates, Cached network location is retrieved, New Cell-ID fix is received, A WiFi-based location is obtained, New WiFi-based location is dismissed due to larger error estimates, A GPS location replaces current best estimate, Stop listening for updates, Best estimate of the location is used in the application, Time (t)

RUTGERS Mining Smartphone Mobility Data 20

Location APIs

Fused location services

- Combine different location providers
- User specifies:
 - Min and max update period
 - Accuracy preferences
- Location service takes care of managing different low-level providers, to obtain best accuracy at low(est) power

RUTGERS Mining Smartphone Mobility Data 21

Location APIs

Geo-fences and activities

Geo-fencing:

- User specifies a POI or map area
- Requests to receive alerts when user is near / inside fence

Activity recognition (Android APIs):

- Use sensor and location data to detect what user is doing
- Walking vs cycling vs driving
- Provides probability for each activity

RUTGERS Mining Smartphone Mobility Data 22

Indoor localization

- No O/S level APIs (?)
- Google Maps offers indoors navigation
 - Mix of WiFi-based localization and (very rough) dead reckoning
- May be possible to obtain WiFi RSS data
 - Android offers APIs, iOS is restricted
 - Also: iBeacon (BLE-based, proximity)
- Other applications have used other signals (like audio)
- Custom solutions also exist (e.g., ultrasound-based)
- More on this later ...

RUTGERS Mining Smartphone Mobility Data 23

Overview

- Sensors
- **Network connectivity**
- Power
- Accounts ("identity")
- Mobile app basics

RUTGERS Mining Smartphone Mobility Data 24

Cellular

- Various standards have evolved over the years
- Hard to track (too much marketing hype...)
- But: rapidly increasing bandwidth and decreasing cost

Year	Bandwidth (Mbps)
1996	~1
2000	~2
2004	~5
2008	~20
2012	~110

RUTGERS Mining Smartphone Mobility Data 25

Cellular

- Historical technologies:
 - 2G: GPRS (GSM)
 - 2.5G: EDGE
- Current technologies:
 - "3G": UMTS (HSDPA, HSUPA, HS[DP]A+)
 - "4G": LTE
- Other technologies (failed adoption):
 - WiBro, WiMax, ...

Standard	Year	D/L (max)*	U/L (max)*
GPRS	1997	60-80kb/s	20-40kb/s
EDGE	2003	177-237kb/s	60-118kb/s
HSPA	2006	14-42mb/s	1(?) - 6mb/s
HSPA+	2008	28-168mb/s	11-22mb/s
LTE	2010	12-300mb/s	5-75mb/s

*numbers based on marketing claims, please take with grain(s) of salt
**approx. year (first major deployment)

RUTGERS Mining Smartphone Mobility Data 26

802.11b/g/n (WiFi)

- Introduced around the same time as GPRS
- Evolved over time: bandwidth and ubiquity

Protocol	Year	B/W (Mbit/s max)
—	1997	2
b	1999	11
g	2003	54
n	2009	72.2 (2.4GHz) 150 (5GHz)
ac (draft)	2012	88-867

RUTGERS Mining Smartphone Mobility Data 27

Bluetooth

- Developed by Ericsson in 1994
- Standardized in 1998
- Developed over years
- Designed almost concurrently with WiFi; designed for short-range communications with peripherals (not Ethernet/IP packets only)
 - Fairly complex
 - Fairly ubiquitous

Class	mW (max)	Range (m)
1	100	100
2	2.5	10
3	1	1

Version	Data rate (Mbit/s)
1.2	1
2.0 + EDR	3
3.0 + HS	24

RUTGERS Mining Smartphone Mobility Data 28

Bluetooth Low Energy (BLE)

Spec	Bluetooth "Classic"	BLE
Range (max)	100m	50m
Data rate	1-3Mbit/s	1Mbit/s
App. throughput	0.7-2.1Mbit/s	0.27Mbit/s
Latency	100ms (typ.)	6ms
Time to send data	100ms	~3ms
Peak current	<30 mA	<15 mA
Power consumption	100% (reference)	10-50% (use case dep.)

- Entirely separate stack (Zigbee derivative)
- Goals: low power, low latency, low(er) cost
- Initially developed by Nokia, became standard in 2010
- Standard on both iPhone and Android

RUTGERS Mining Smartphone Mobility Data 29

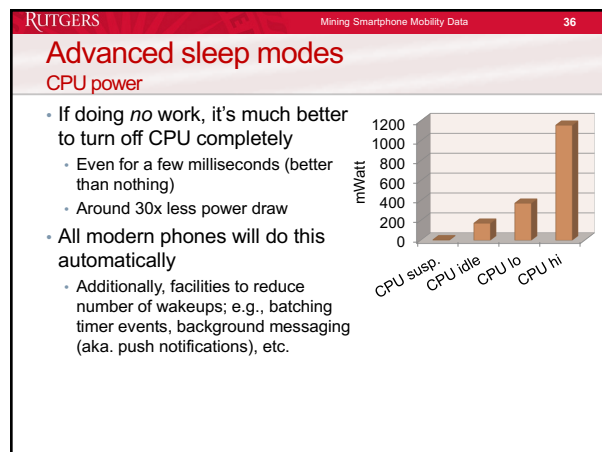

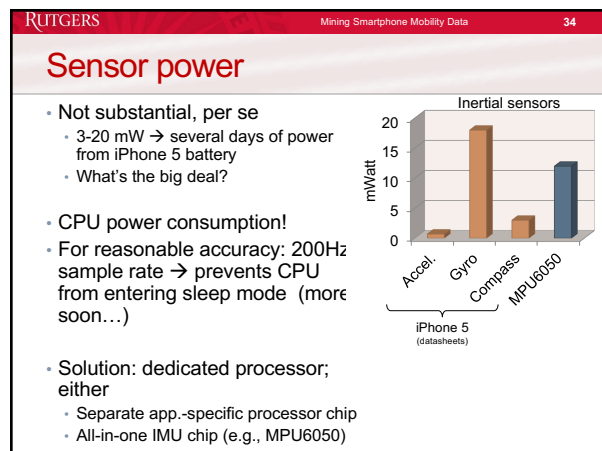
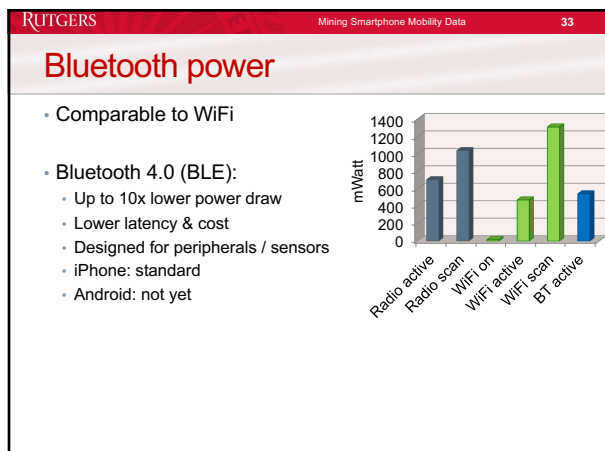
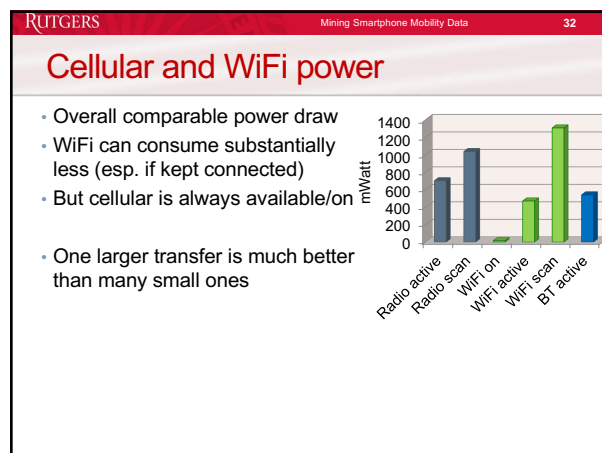

Other

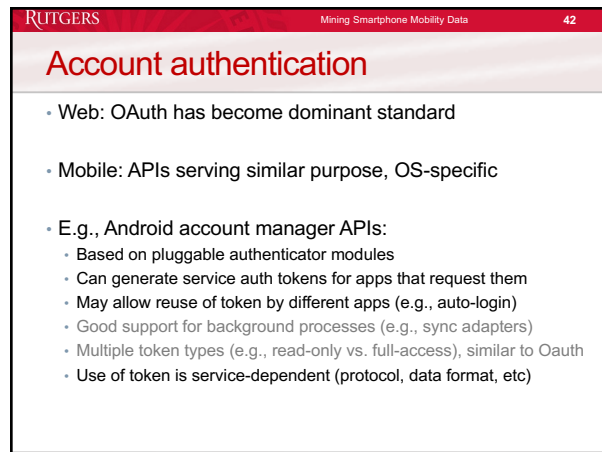
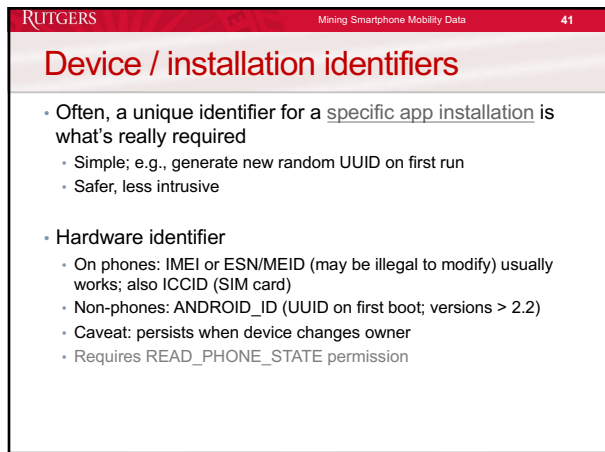
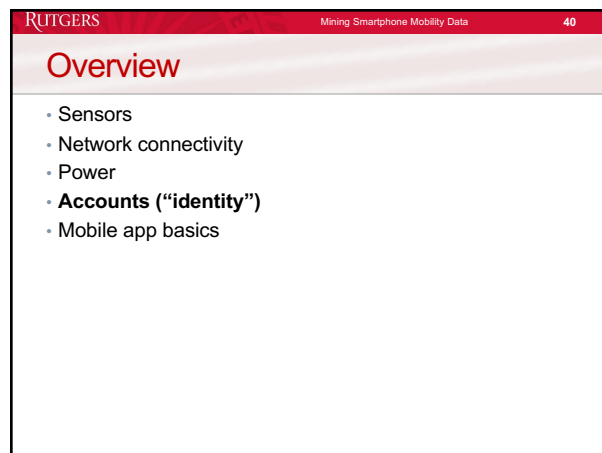


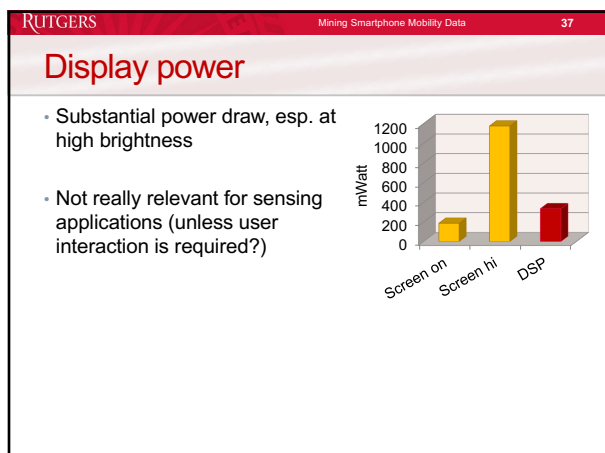
- Zigbee / XBee
- Cheap transceivers (e.g., Nordic chipsets)
- Non-standard (on phones), require ugly dongles, etc.
- But, might be worth it for prototyping

RUTGERS Mining Smartphone Mobility Data 30

Overview

- Sensors
- Network connectivity
- Power**
- Accounts ("identity")
- Mobile app basics





Sync adapters & content providers

- Sync adapter (Android):
 - Centralized facility for background syncing of on-device and remote data (e.g., contacts, posts, etc)
 - Allows optimizations (e.g., to conserve power)
- Content provider (Android):
 - APIs to expose (synced) data to other applications
 - Well-defined query endpoints and schemas
 - Cursor-like API

Overview

- Sensors
- Network connectivity
- Power
- **Mobile app basics**

Programming paradigm

Heavily event-oriented !

Application must respond to its environment; e.g.

- Network connectivity changes
- Incoming calls / messages / events
- Sensor / location data
- ...

Application must use resources efficiently; e.g.

- May be pre-empted and/or killed at any time
- May choose to respond to status information (e.g, battery level)
- ...

Cannot:


- Assume a single main() thread with sequential flow
- Control lifetime of thread(s)

Programming paradigm

Android: *activity* paradigm

- You can think of an activity as a screen
- Can be in different *states* during its *lifecycle*
- Need to respond to state-change events
- System determines state based on:
 - User interactions (e.g., start a different activity)
 - External events (e.g., screen rotation, incoming call, ...)
 - Available resources (memory, CPU, etc)
- Execute in the main app thread
- Responsible for persisting any app-specific state, as necessary

Activity lifecycle (Android)



A note on "deep linking"

- Android: an app can declare which activities handle a particular URL pattern
 - E.g., YouTube app could declare that URLs with `http://` protocol and `youtube.com` domain can be handled by its video player activity
 - If multiple apps can handle the same URL type, the user will be prompted to choose
- iOS: similar, but not as general/broad (?)

RUTGERS Mining Smartphone Mobility Data 49

Background tasks

Short lived (e.g., fetch a URL):

- Can be started in separate threads
- But: need to be prepared for activity termination/restart

Long lived:

- Need to use system APIs to register themselves and allow system to manage them
- Timers, background services, RPC interfaces

Avoid whenever possible!!

- Use system services instead, e.g., geo-fences, push notifications, etc

RUTGERS Mining Smartphone Mobility Data 50

Tutorial plan

11:55 – 13:00 Katharina

- Resource-constrained graphical models
- App usage mining and traffic prediction

14:30 – 15:15 Tina

- Local-based social networks
- Mobile advertising and search

15:15 – 16:10 Dimitrios

- Learning from trajectory data
- Crowdsourcing and applications

Lunch break (13:30-14:30)

RUTGERS THE STATE UNIVERSITY OF NEW JERSEY Mining Smartphone Mobility Data 51

MINING SMARTPHONE MOBILITY DATA

Mobile technology overview


Spiros Papadimitriou, Tina Eliassi-Rad,
Rutgers University Northeastern

Katharina Morik, Dimitrios Gunopulos
TU Dortmund University of Athens



RUTGERS Mining Smartphone Mobility Data 52

License



These slides are made available under a Creative Commons Attribution-ShareAlike license (CC BY-SA 3.0):
<http://creativecommons.org/licenses/by-sa/3.0/>

You can share and remix this work, provided that you keep the attribution to the original authors intact, and that, if you alter, transform, or build upon this work, you may distribute the resulting work only under the same or similar license to this one.

© 2013–6 Spiros Papadimitriou, Tina Eliassi-Rad