Overview of the CS Department and NODES: Mobile and 5G Research

Professor Sasu Tarkoma, Head of Department

NODES Research Group

22 March 2017

University of Helsinki

- The largest and the oldest university in Finland
- Key data for 2015
  - 32 000 students
  - 7 900 employees
  - 300 subjects
  - 6 100 degrees/year
  - 530 PhDs/year
- Founded in Turku 1640
- Moved to Helsinki 1828

Faculty of Science at Kumpula Campus

Departments

- Chemistry
- **Computer science**
- Geosciences and Geography
- Mathematics and Statistics
- Physics

50 Years of Excellence

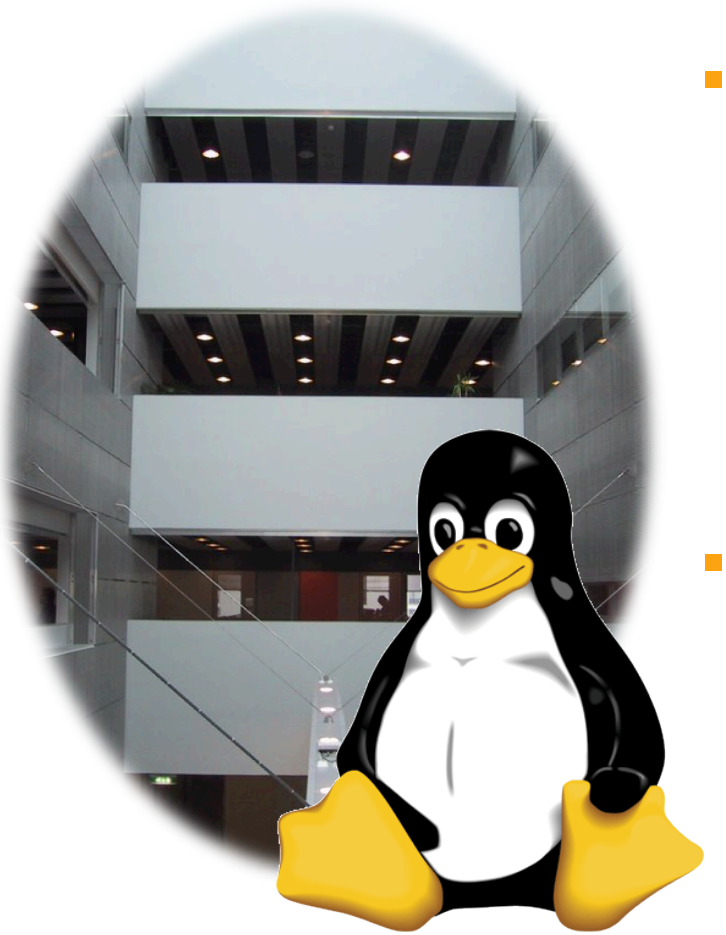

- **Department of Computer Science**
- Leading institution in Computer Science in Finland
  - #1 in Finland in QS Ranking 2017
  - #1 in Nordic Countries and overall #69 in Times Higher Education 2017

- **Core CS and Data Science**
  - Algorithms, Data Analytics and Machine Learning
  - Software Systems
  - Networking and Services (NODES)
  - Bioinformatics

carat.cs.helsinki.fi

Current research topics include:
Digital services, IoT security and privacy, software-defined networks, Data Science, ...

Mobile Edge Computing

Big Data Frameworks

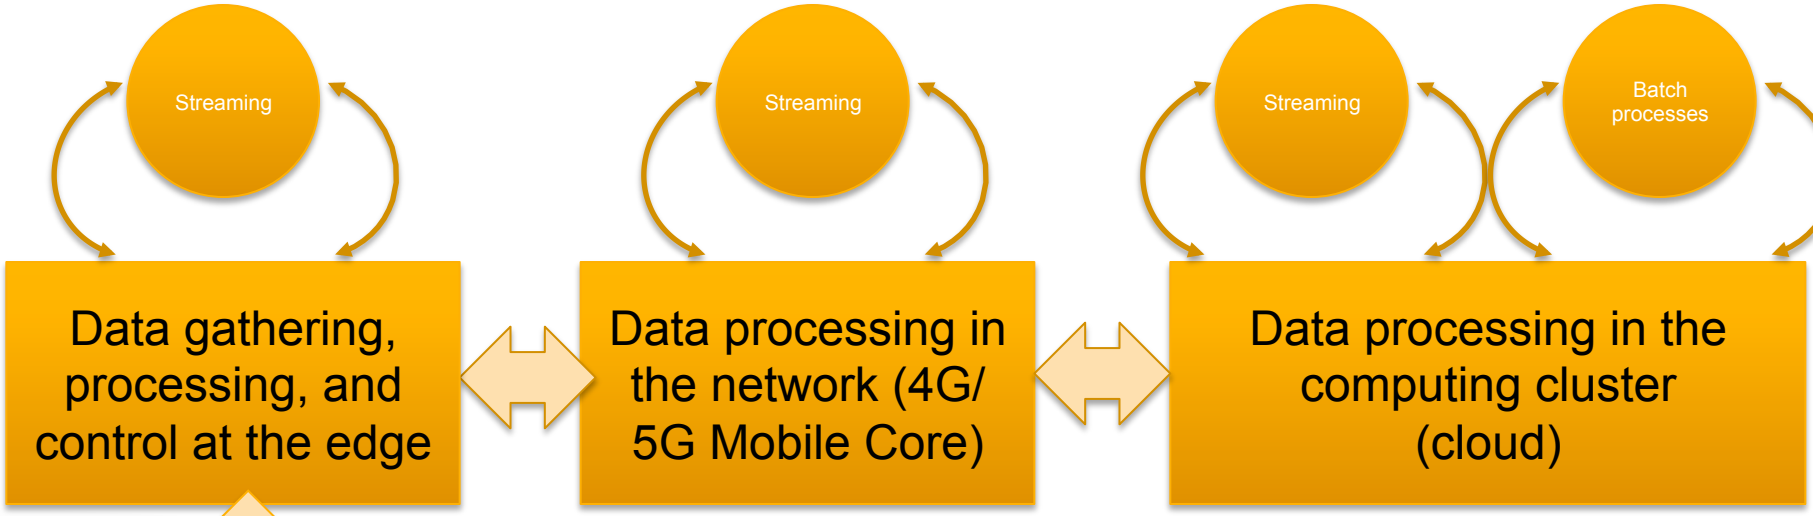

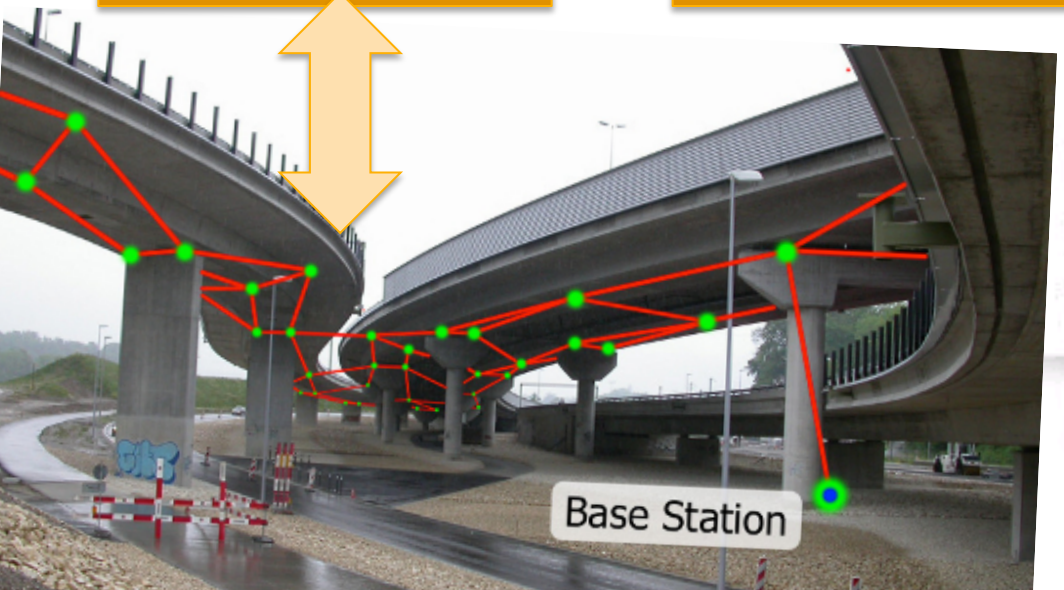

Mobile Research: Carat

Carat team (carat.cs.helsinki.fi)

Motivation

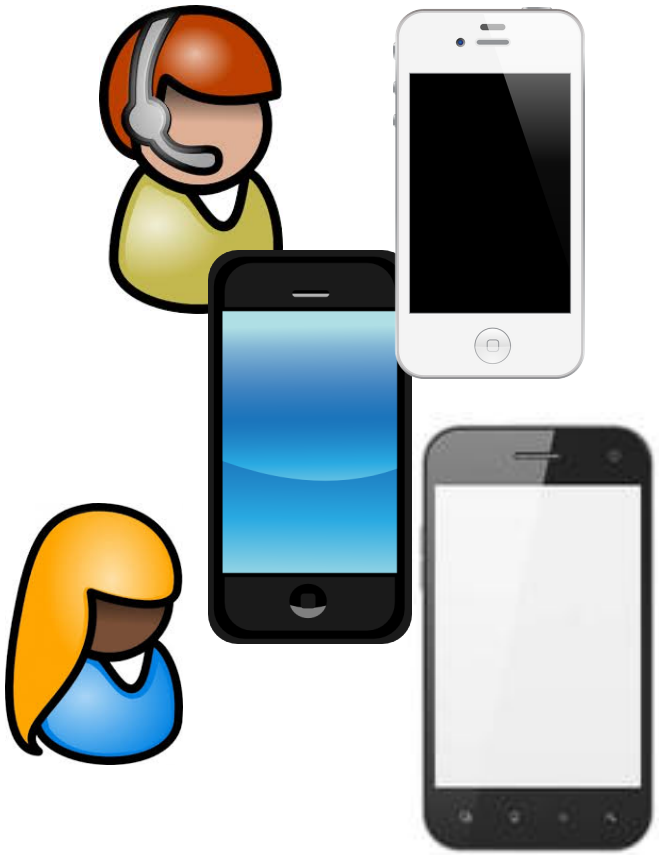

Battery lifetime?

Risk level?

Many heterogeneous, active devices and many users with different intents. – What kind of behavior is **normal** or **typical**?

Introducing Carat

Carat is the **first system** to use the mobile device community to detect and correct energy problems

Our method for **diagnosing** energy anomalies uses the community to infer a specification (expected energy use), and we call deviation from that inferred specification an anomaly

Carat

- Originated in UC Berkeley, in collaboration with University of Helsinki
- Mobile app for Android and iOS
- Currently over 850 000 users
- >2.5 TB of data, > 250 million measurements
- Research project with many directions
- http://carat.cs.helsinki.fi

UNIVERSITY OF HELSINKI

amplab UC BERKELEY

What is Carat?

- Mobile app shows users advice:
  “Kill Facebook for 16m ± 41s battery life”
- Energy **hogs** and **bugs**
- Tracks user's battery life average since installation
- Places users within community with a ranking called J-Score

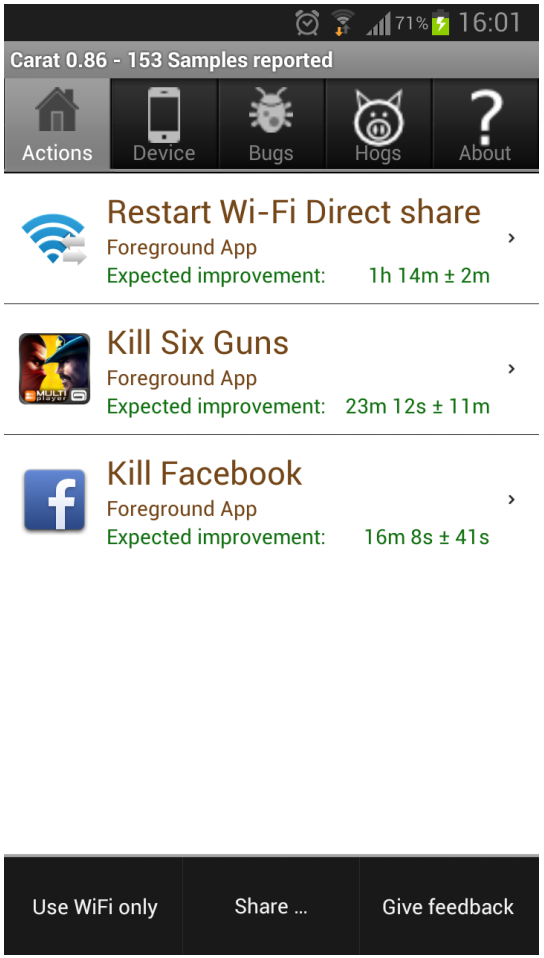

New UI

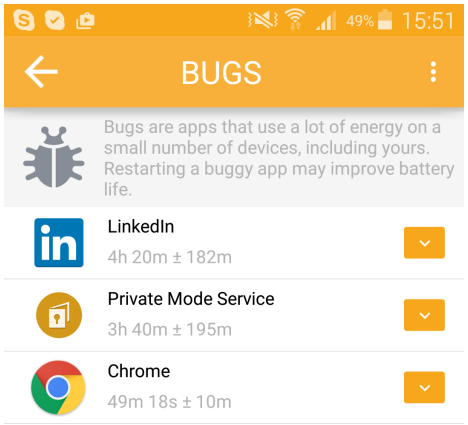
49% 15:51
BUGS
Bugs are apps that use a lot of energy on a small number of devices, including yours. Restarting a buggy app may improve battery life.
LinkedIn
4h 20m ± 182m
Private Mode Service
3h 40m ± 195m
Chrome
49m 18s ± 10m

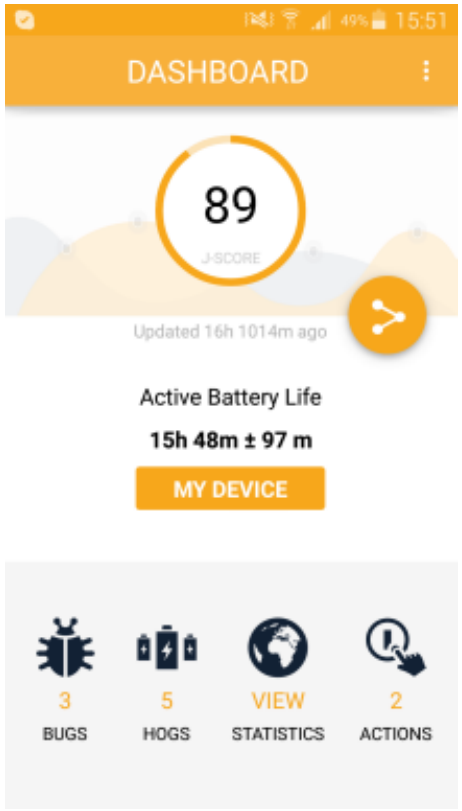
49% 15:51
DASHBOARD
89
J-SCORE
Updated 16h 1014m ago
Active Battery Life
15h 48m ± 97 m
MY DEVICE
3
BUGS
5
HOGS
VIEW
STATISTICS
2
ACTIONS

Comment 36 Like 1k Tweet Share 309 +1 171

Carat: The Brilliant App That Increases Your Battery Life By Showing What Other Apps To Kill

JOSH CONSTINE

Thursday, June 14th, 2012 **36 Comments**

"Kill Pandora – Expected Battery Life Improvement: 1 hour 50 minutes" This is what you'll learn from **Carat**, an incredibly useful free **new iOS** and **Android app** that's the first to give you personalized mobile battery life-saving recommendations.

Carat quietly takes measurements from you device, does some math, combines it with other people's anonymized data, and sends back tips on if you should update your OS, kill or restart apps, and how many more minutes of tablet or phone fiddling you'll gain.

As battery tech is expected to improve slowly, some say increasing life just 5% a year, and as we get faster processors, more powerful apps, and brighter screens, everyone could use a Carat in their pocket.

Suddenly...

lifehacker

TOP STORIES

t

By

'o Kill

36 Comments

Life Improvement: 1 hour
ırn from **Carat**, an
d **Android app** that's the
ıile battery life-saving

s from you device, does
r people's anonymized data, and sends back tips on if you should update your OS, kill or restart apps, and how many more minutes of tablet or phone fiddling you'll gain.

As battery tech is expected to improve slowly, some say increasing life just 5% a year, and as we get faster processors, more powerful apps, and brighter screens, everyone could use a Carat in their pocket.

Suddenly...

lifehacker

TOP STORIES

171

t

By

'o Kill

Free Carat app finds 'energy hogs,' 'energy bugs' on iOS or Android devices

ANDROID | JUNE 14, 2012 | BY: MICHAEL SANTO

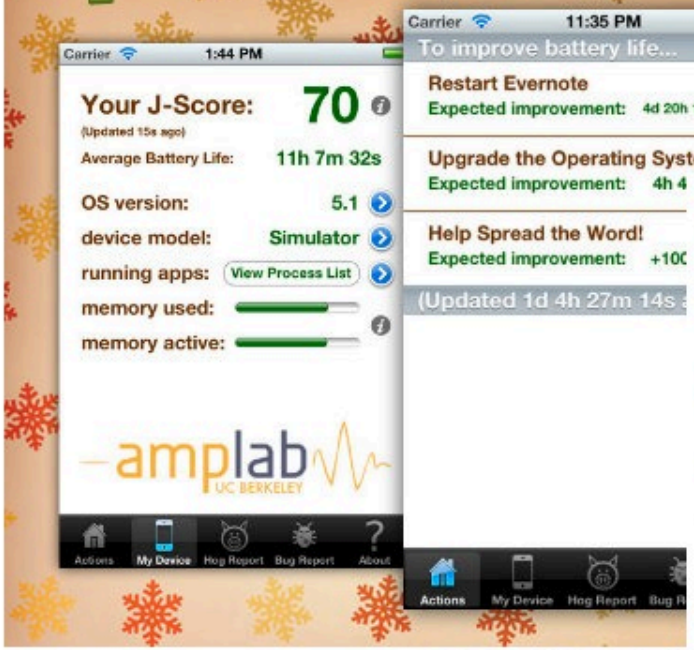

Su

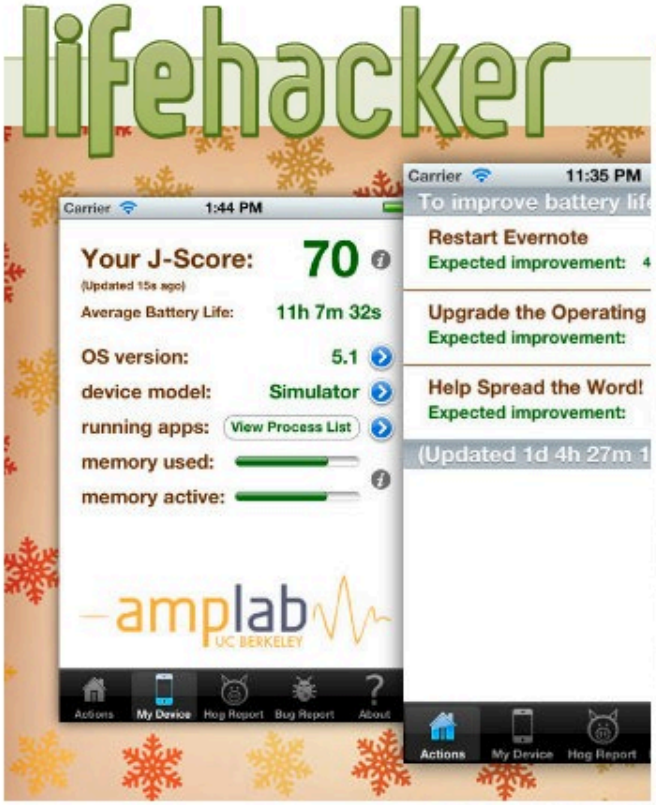

Carat: Extend Your Phone's Battery Life

LESLIE HORN JUNE 19, 2012 2:00 PM

Su

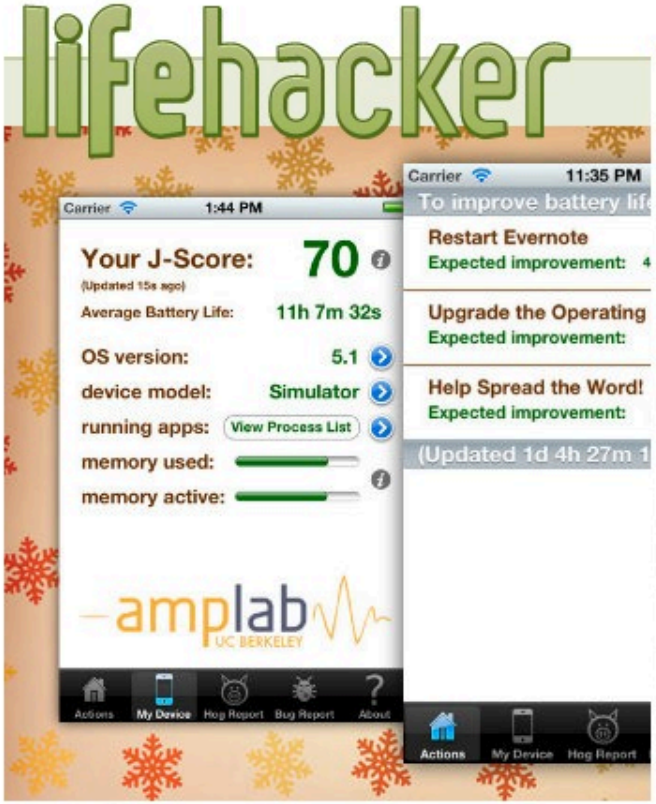

Carat: Extend Your Phone's Battery Life

LESLIE HORN JUNE 19, 2012 2:00 PM

Topic: iPhone

iOS and Android app helps you get more from your battery

Summary: *Carat has been developed by a team of scientists from the UC Berkeley electrical engineering and computer science department's Algorithms, Machines, and People Laboratory (AMP Lab).*

By Adrian Kingsley-Hughes for Hardware 2.0 | June 15, 2012 -- Updated 10:21 GMT (03:21 PDT)

Follow @the_pc_doc

lifehacker

Carat: Extend Your Phone's Battery Life

LESLIE HORN JUNE 19, 2012 2:00 PM

Carrier 1:44 PM

Your J-Score: 70

(Updated 15s ago)

Average Battery Life: 11h 7m 32s

OS version: 5.1

Carrier 11:35 PM

To improve battery lif

Restart Evernote

Expected improvement: 4

Upgrade the Operating

Expected improvement:

Carrier 1:44 PM

Your J-Score: 70

Carat went viral. In 24 hours, there were dozens of articles and we had more than 100,000 users.

Topic: iPhone

iOS and Android app helps you get more from your battery

Summary: *Carat has been developed by a team of scientists from the UC Berkeley electrical engineering and computer science department's Algorithms, Machines, and People Laboratory (AMP Lab).*

By Adrian Kingsley-Hughes for Hardware 2.0 | June 15, 2012 -- Updated 10:21 GMT (03:21 PDT)

Follow @the_pc_doc

Statistics (October 2016)

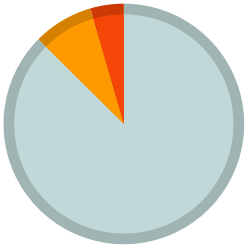

471 645 Android and iOS apps
10% energy hogs, 4% energy bugs

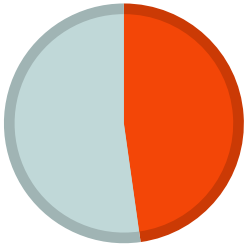

50% of devices have at least one energy bug

Android has a long tail of different device types.

carat.cs.helsinki.fi/statistics

The Carat project: System

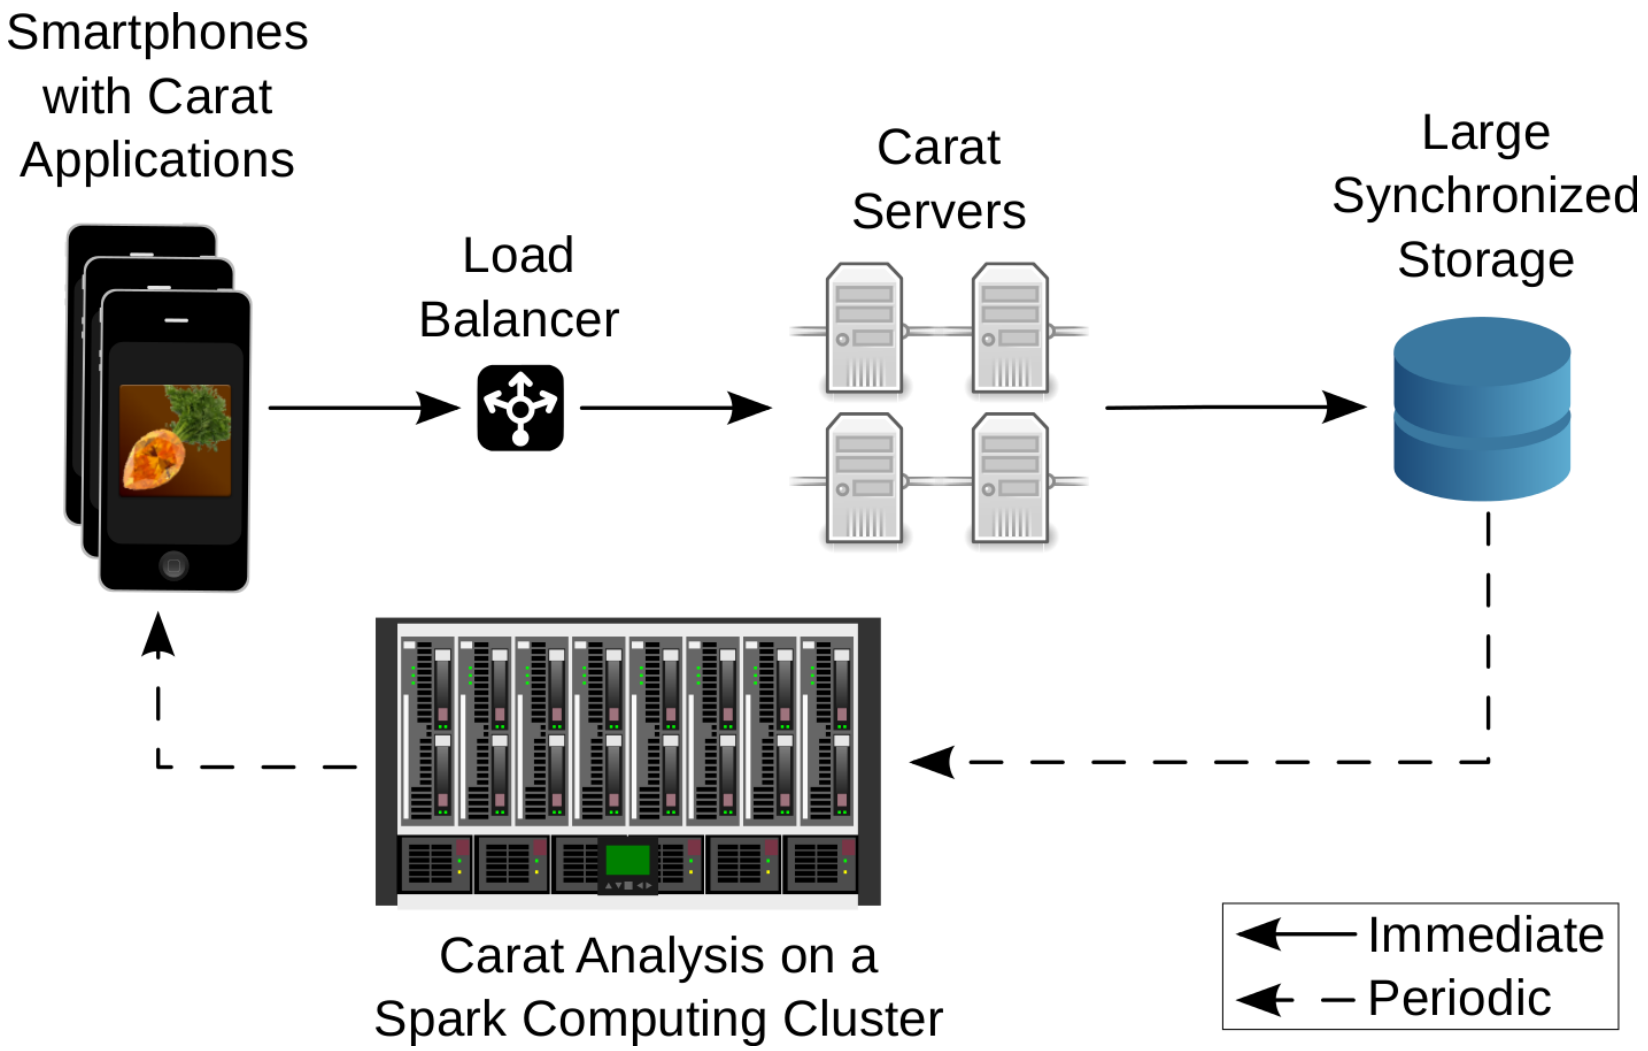

What is Carat?

- Users see Hogs, high energy use apps
- And Bugs that use energy faster on THEIR device than on others
- Users with these issues quickly see battery life benefits once they are addressed

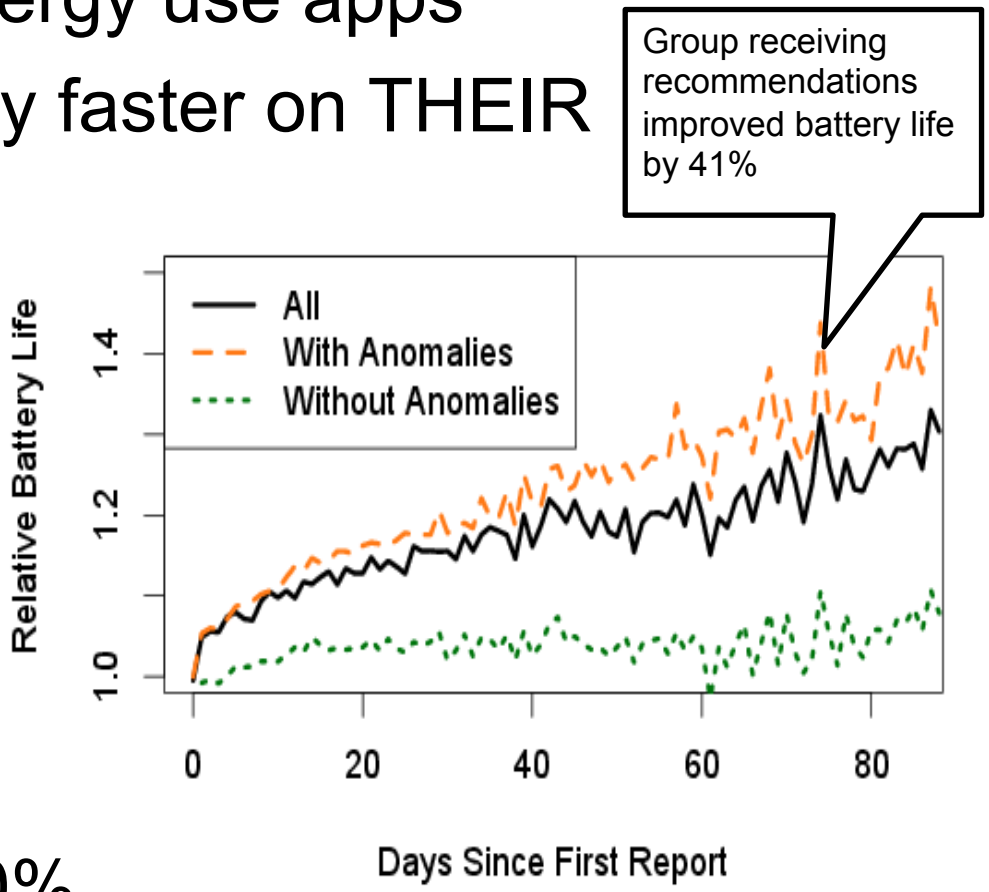

- Average improvement 20%
- Those with energy anomalies can improve 41%

The Data Analysis

- Samples are combined to obtain energy drain probability distributions (with features)
  - Users, Apps, App and User pairs, OS versions, Device models
- Distributions are compared using the distance between their 95% confidence interval error bars
  - If a distribution has a positive distance from another and a higher mean, it is a:
  - **Hog** (for an app vs the distribution for other apps)
  - **Bug** (for app & user combination vs other users of the same app)

Hogs and Bugs

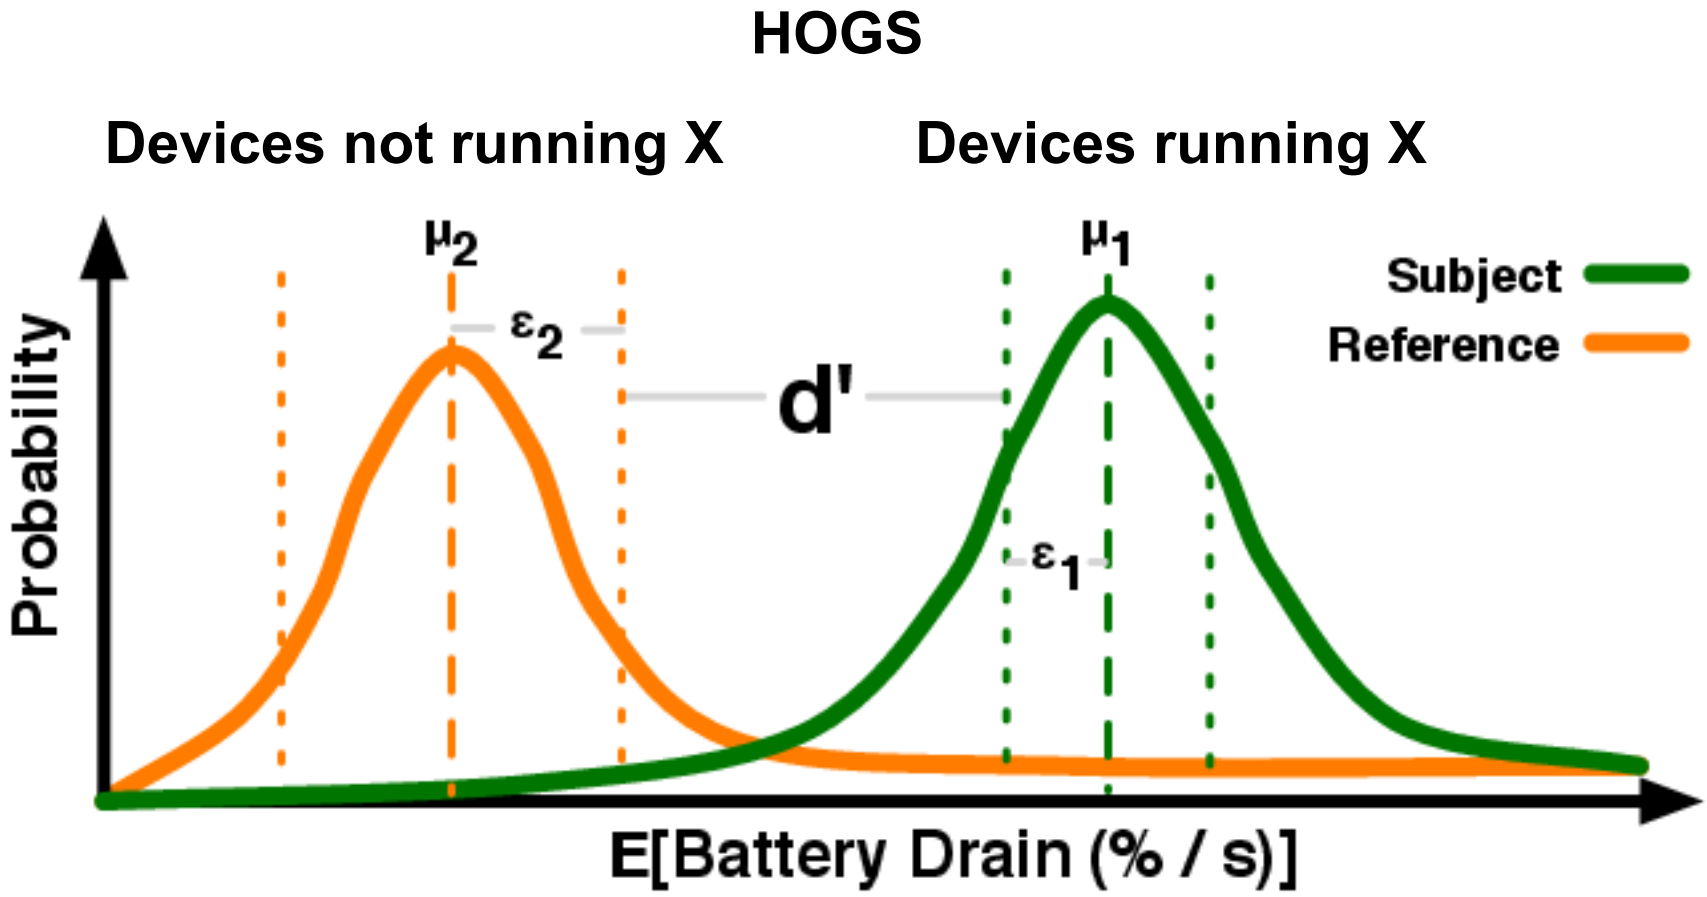

...

Hogs and Bugs

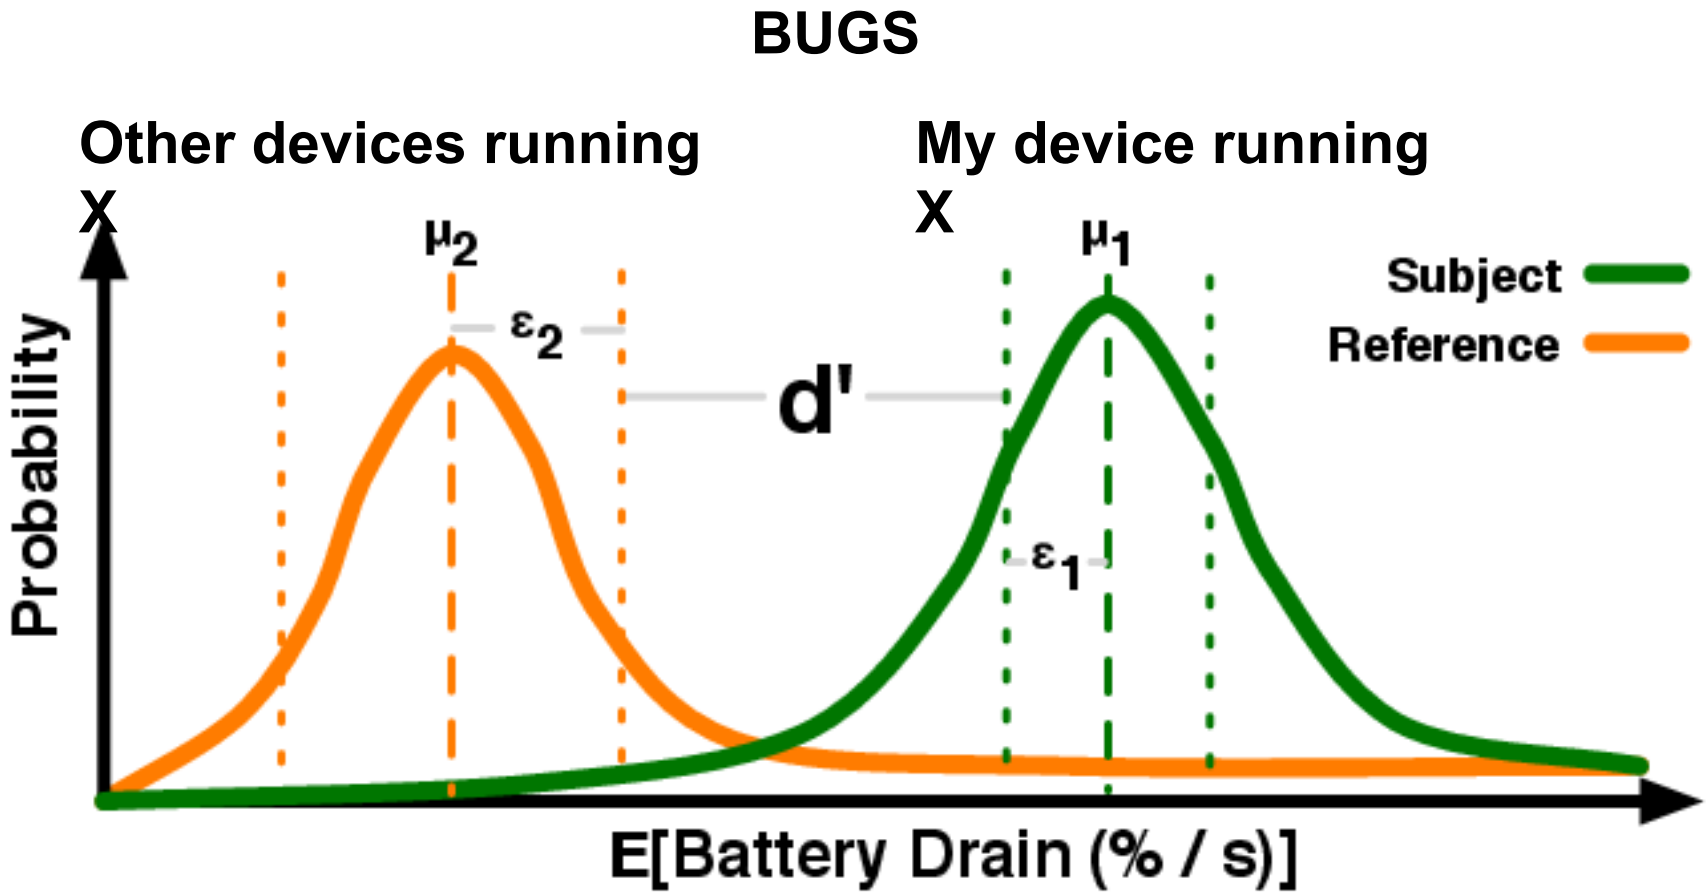

...

Collaborative Data Gathering

Each device collects: Battery life, timestamp, running apps, context/system settings

The data is combined and the results for your apps and your device are sent back to you

Context feature analysis: how various context features affect the energy consumption of the device

Collaborative aspect: We observe trends in the community, as well as how your device is different

The method can be used for phones, sensors, houses, base stations, servers, laptops, … anything that generates measurements

Example: The Kindle WhisperSync bug

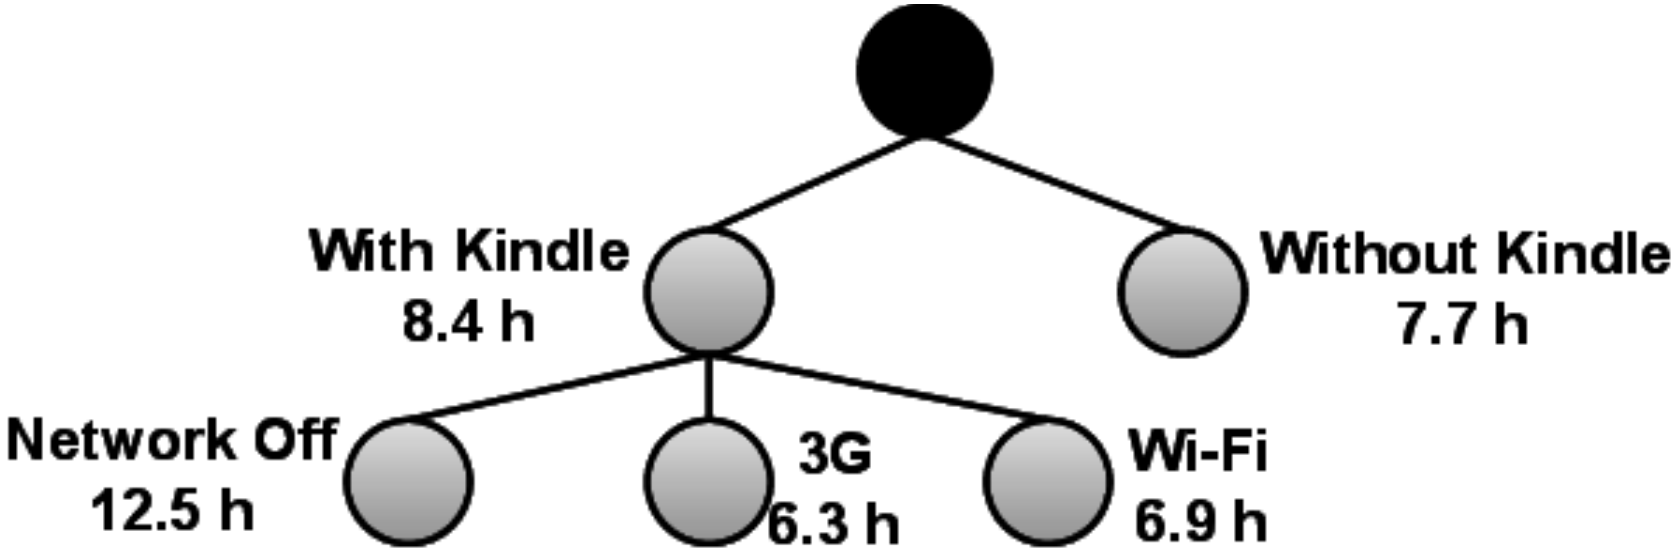

The decision tree allows "what-if" analysis and the generation of recommendations

Project Infrastructure

- Data Analysis: Amazon EC2
  - 10 x X-Large VM (4 cores, 15G memory)
- Server facing mobile devices: Amazon EC2
  - 4 x medium VM (1 core, 4G memory)
  - Load balancer, independent DNS name for easy changing of infrastructure when required
- Amazon S3
  - Storage of data (incoming 0.5-1.0 GB / week)

Lessons learned

- Research prototype != product
- It is not easy to scale
  - 100 000 users in one day when we launched
  - Scaling will cost, cloud is not free
  - Managing clusters is not easy
- Design system so that it can evolve (no hardcoding, extensible formats)
- Validation is not easy
  - Ground truth
  - Injected bugs, validated bugs

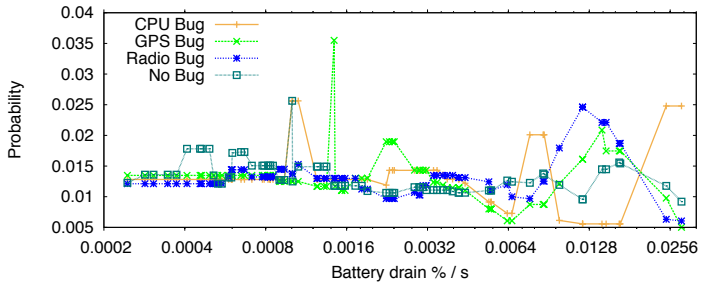

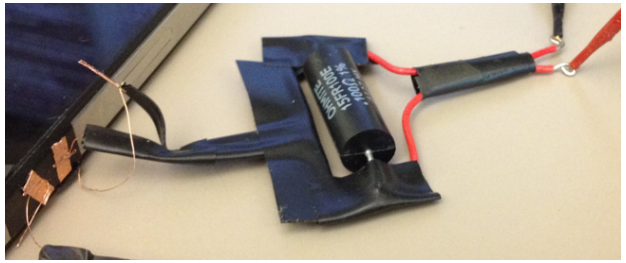

Energy efficient configuration?

Energy efficient configuration?

Network type

Energy efficient configuration?

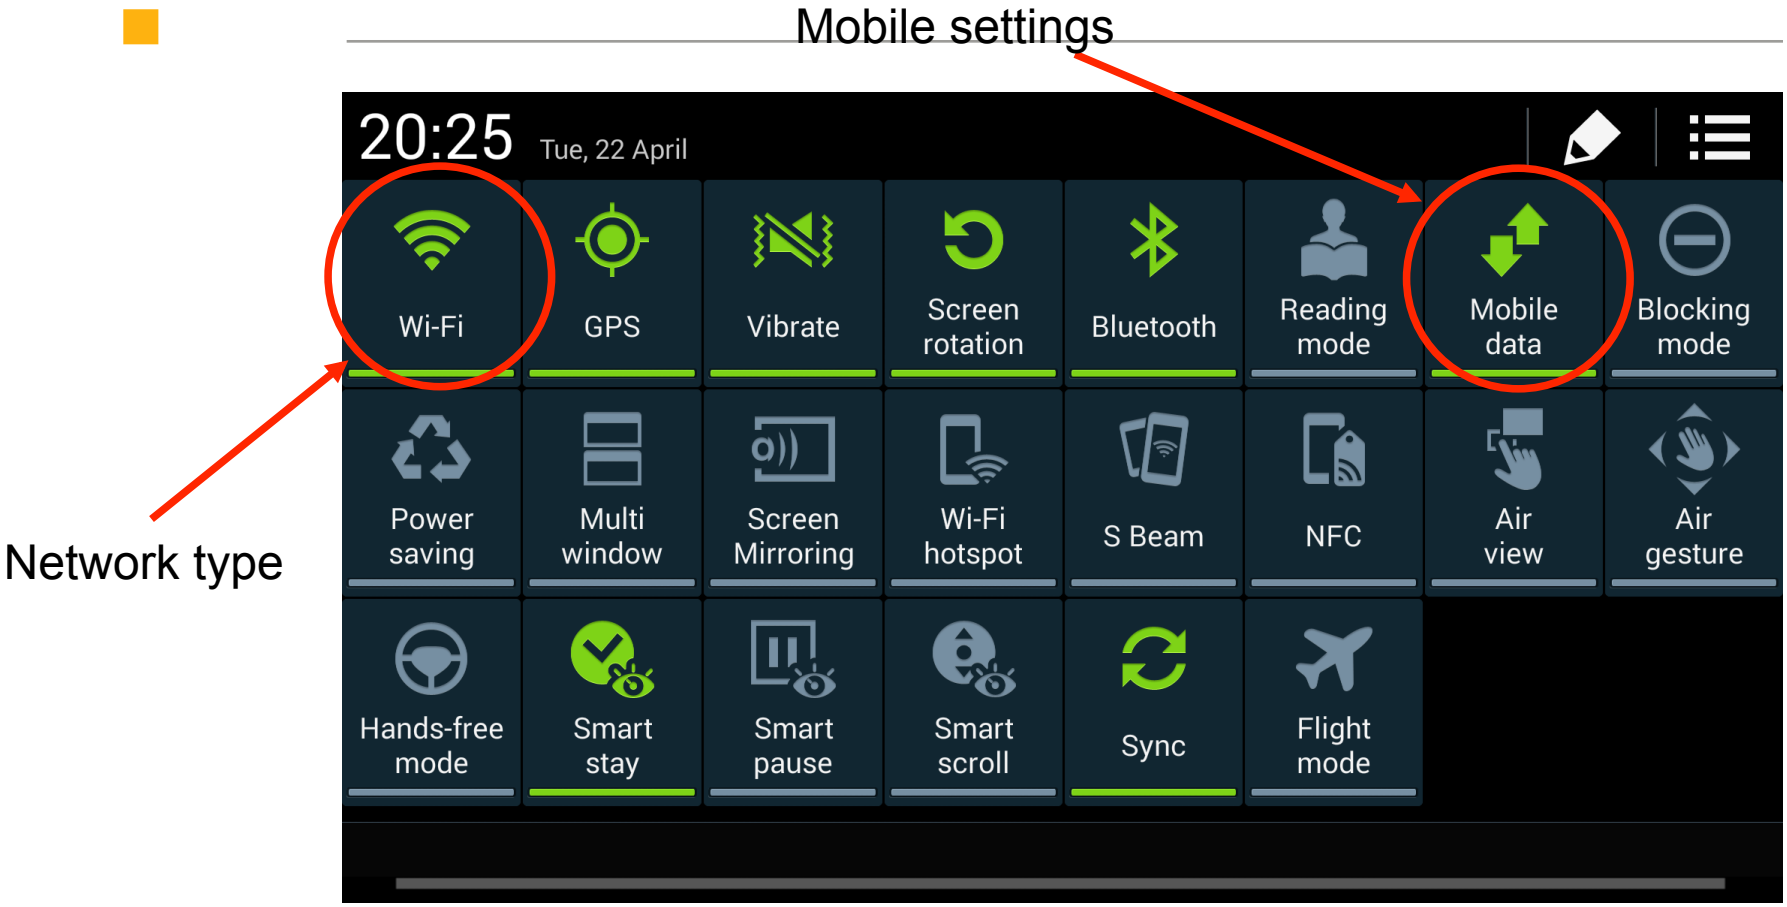

Energy efficient configuration?

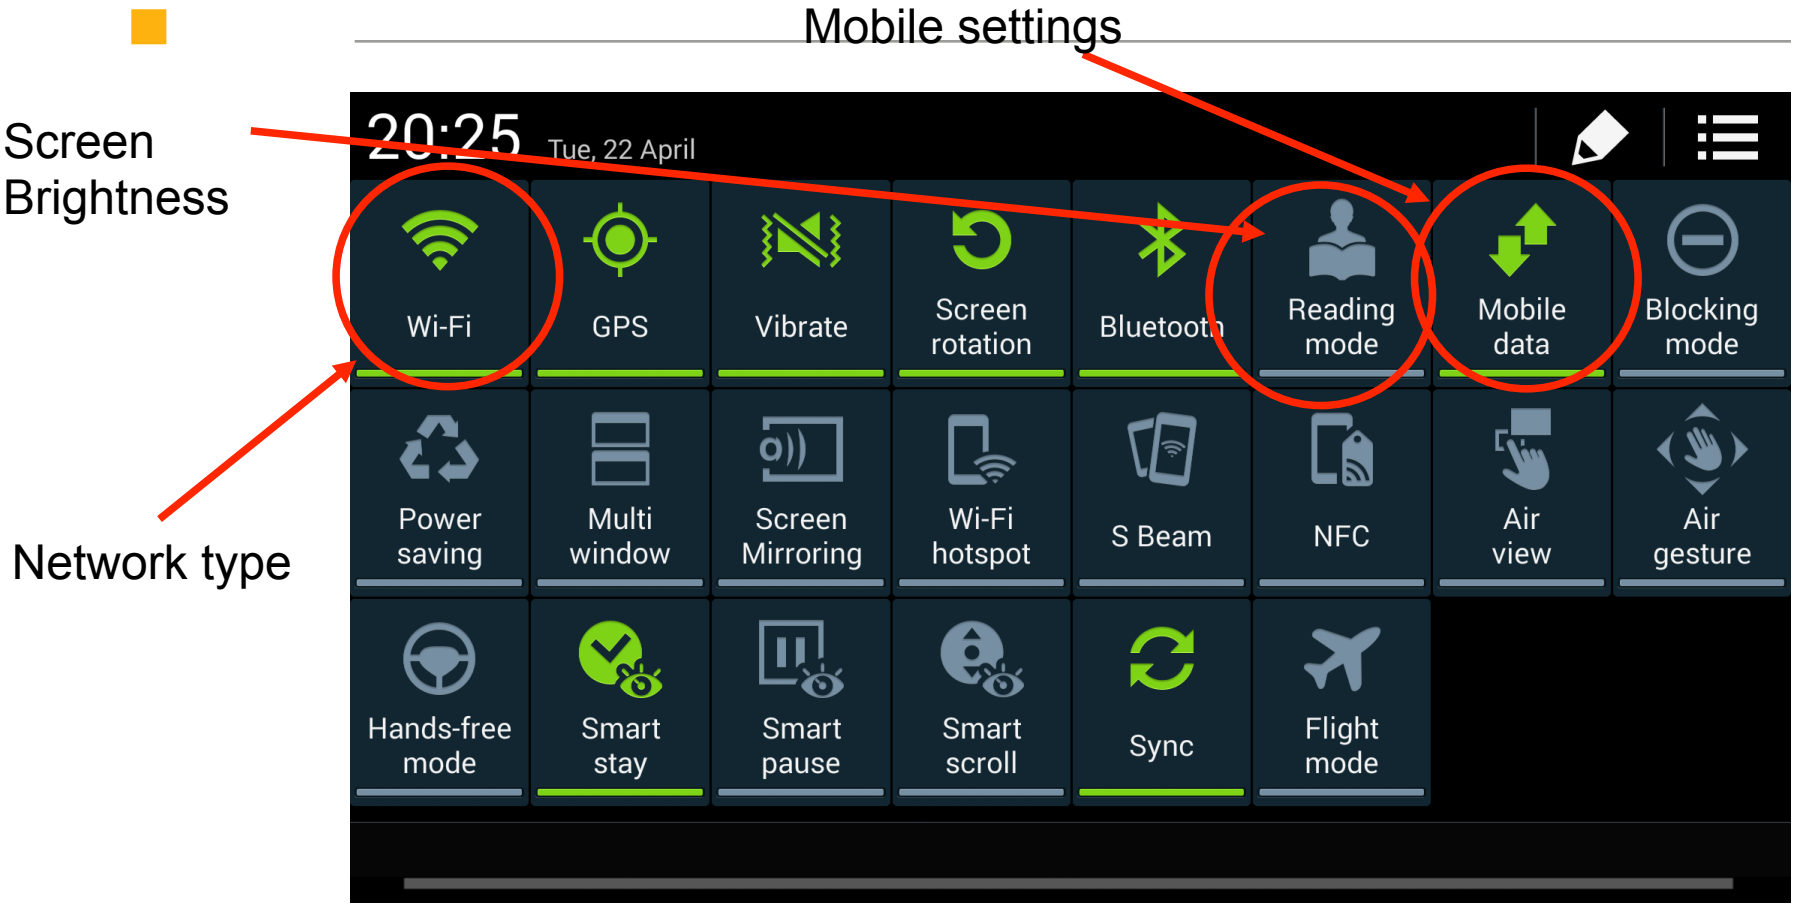

Energy efficient configuration?

Mobile settings

Screen Brightness

Network type

Moving or not?

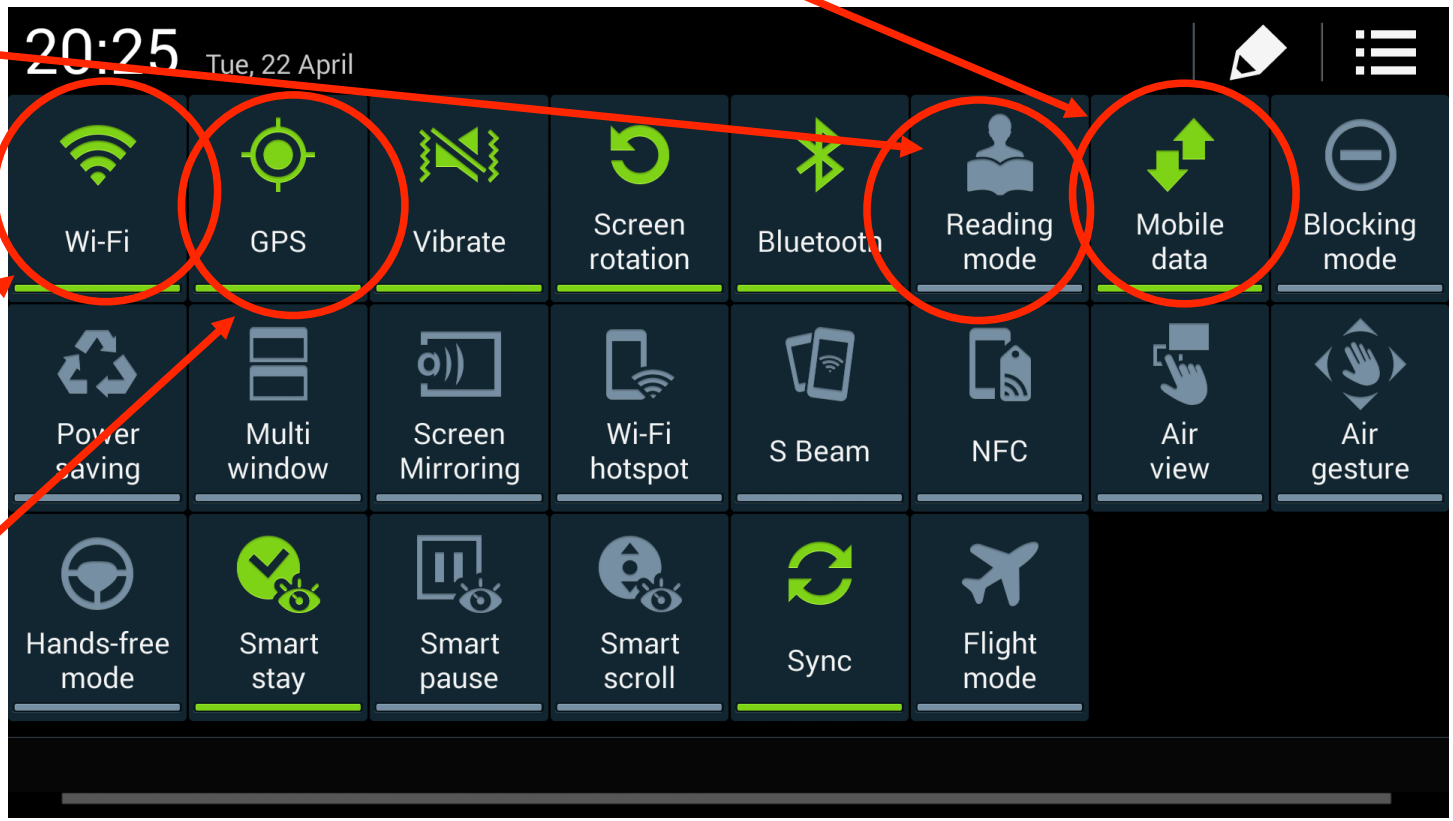

Selected findings

- Wi-Fi signal strength dropping one bar can result in over 13% battery loss
- High temperature can cause 50% battery loss, and high temperature is not always related to high CPU load
- Automatic screen brightness is, in the most cases, better than manual setting
- In addition to CPU, battery temperature and distance traveled are useful in predicting battery lifetime

Battery lifetime – an example

Battery Temperature	Distance Traveled	CPU Use Level	Screen Brightness	Estimated Battery Life (h)
Under 30°C	>0	Low	Automatic	8.83 – 9.12
Under 30°C	>0	Low	Manual	8.49 – 8.82
Under 30°C	>0	High	Automatic	8.09 – 8.24
Under 30°C	>0	Medium	Automatic	7.65 – 7.89
Under 30°C	>0	Medium	Manual	7.34 – 7.60
Under 30°C	>0	High	Manual	7.27 – 7.41
Under 30°C	None	Medium	Automatic	6.57 – 6.64
Under 30°C	None	Low	Automatic	6.28 – 6.35
Under 30°C	None	Medium	Manual	6.13 – 6.20
Under 30°C	None	Low	Manual	5.88 – 5.96
Under 30°C	None	High	Automatic	5.78 – 5.82
Over 30°C	>0	Low	Automatic	5.08 – 5.22
Under 30°C	None	High	Manual	5.00 – 5.04
Over 30°C	>0	Low	Manual	4.73 – 4.88
Over 30°C	>0	High	Automatic	4.62 – 4.69
Over 30°C	>0	Medium	Automatic	4.59 – 4.70
Over 30°C	>0	Medium	Manual	4.28 – 4.39
Over 30°C	None	Medium	Automatic	4.25 – 4.29
Over 30°C	>0	High	Manual	4.08 – 4.14

Just want to play a game?
High CPU use?

[2] Ella Peltonen, Eemil Lagerspetz, Petteri Nurmi, and Sasu Tarkoma. Energy Modeling of System Settings: A Crowdsourced Approach. PerCom '15. Best Paper Award.

Battery lifetime

→ 98% better expected battery life

Battery Temperature	Distance Traveled	CPU Use Level	Screen Brightness	Estimated Battery Life (h)
Under 30°C	>0	Low	Automatic	8.83 – 9.12
Under 30°C	>0	Low	Manual	8.49 – 8.82
Under 30°C	>0	High	Automatic	8.09 – 8.24
Under 30°C	>0	Medium	Automatic	7.65 – 7.89
Under 30°C	>0	Medium	Manual	7.34 – 7.60
Under 30°C	>0	High	Manual	7.27 – 7.41
Under 30°C	None	Medium	Automatic	6.57 – 6.64
Under 30°C	None	Low	Automatic	6.28 – 6.35
Under 30°C	None	Medium	Manual	6.13 – 6.20
Under 30°C	None	Low	Manual	5.88 – 5.96
Under 30°C	None	High	Automatic	5.78 – 5.82
Over 30°C	>0	Low	Automatic	5.08 – 5.22
Under 30°C	None	High	Manual	5.00 – 5.04
Over 30°C	>0	Low	Manual	4.73 – 4.88
Over 30°C	>0	High	Automatic	4.62 – 4.69
Over 30°C	>0	Medium	Automatic	4.59 – 4.70
Over 30°C	>0	Medium	Manual	4.28 – 4.39
Over 30°C	None	Medium	Automatic	4.25 – 4.29
Over 30°C	>0	High	Manual	4.08 – 4.14

[2] Ella Peltonen, Eemil Lagerspetz, Petteri Nurmi, and Sasu Tarkoma. Energy Modeling of System Settings: A Crowdsourced Approach. PerCom '15. Best Paper Award.

The PADS Project

The PADS project will develop new privacy enhancing algorithms and methods for **Data Science**

Two methodological goals

- Develop the algorithmic framework for privacy-aware predictive modelling, and
- Scalable implementation of the framework

Two key use cases

- Large-scale genome data processing and sharing for personalized medicine, and
- Mobile and environmental sensing

Research status

Mobile sensing data analysis in progress

Full paper prepared based on questionnaire study

Privacy SDK outlined in the recent Big Data workshop paper.

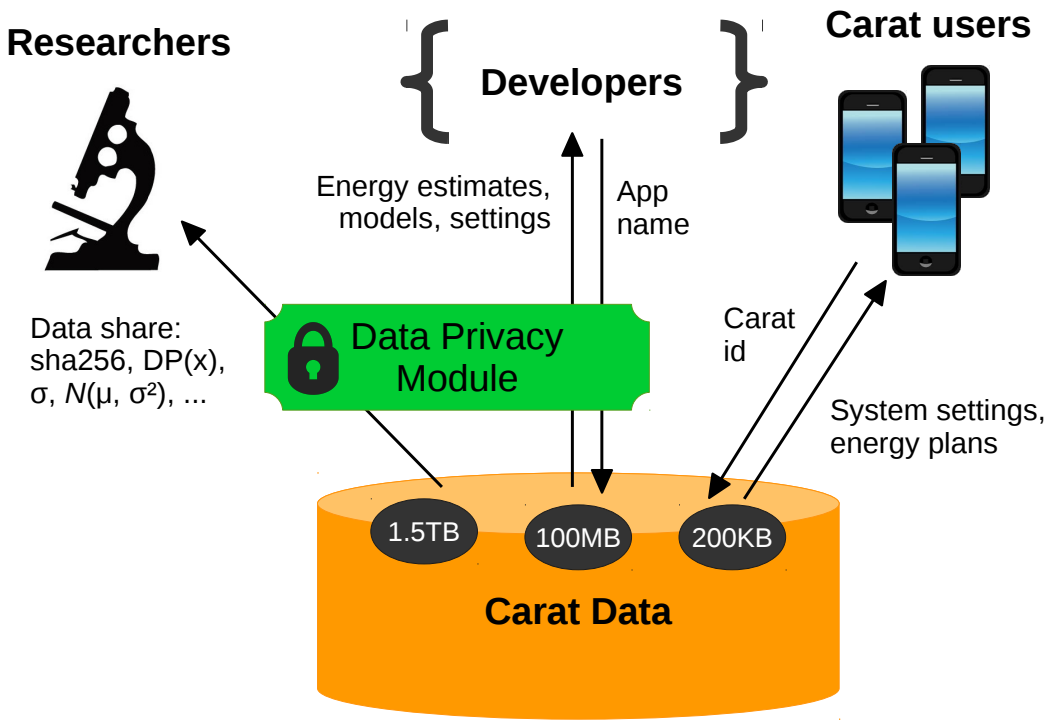

Malware Infection Rates

We studied malware based on the dataset

McAfee, Mobile Sandbox, MalGenome, …

Malware infection rates are higher than conservative estimates (0.26% of devices)

Google says 0.12% of manually installed packages are malware, not very far from this number

Lookout Antivirus predicts >1%

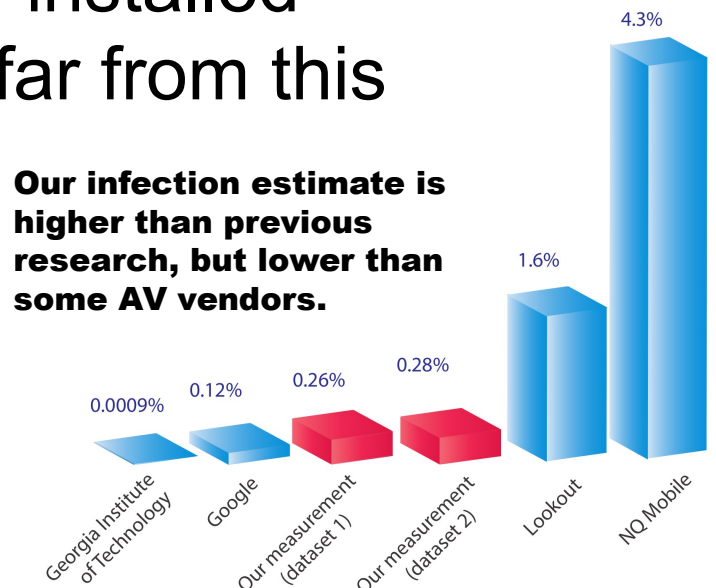

An Early Warning System for Malware

A lightweight technique for identifying devices at risk

By looking at applications that occur with malware, it is possible to predict infection 5x better than choosing devices at random

- Useful for administrators, organisations (**Bring Your Own Device** scenario)

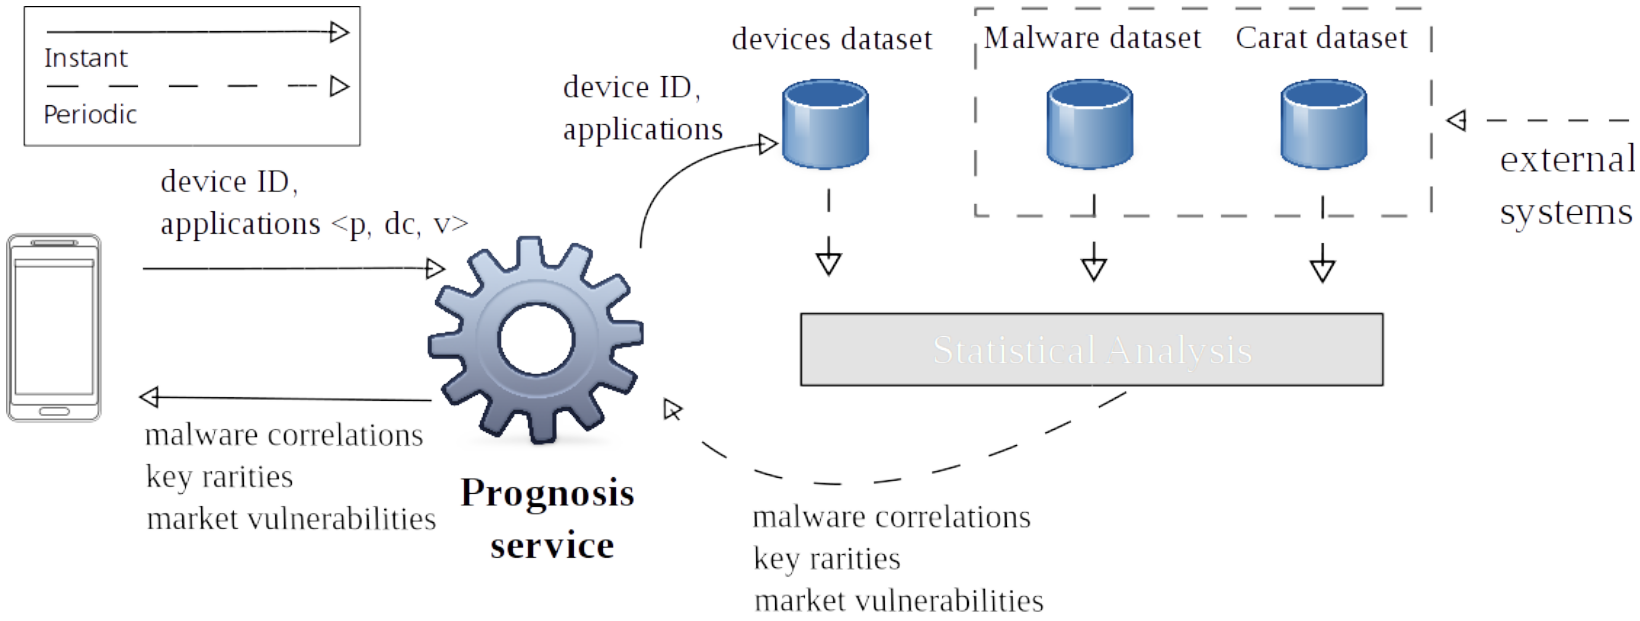

MDoctor: Increasing Awareness of Infection Vulnerability

MDoctor shows status of applications according to a malware dataset

Infection vulnerability can be seen from device health

Three metrics for application analysis: malware correlation, key rarity, and market vulnerability

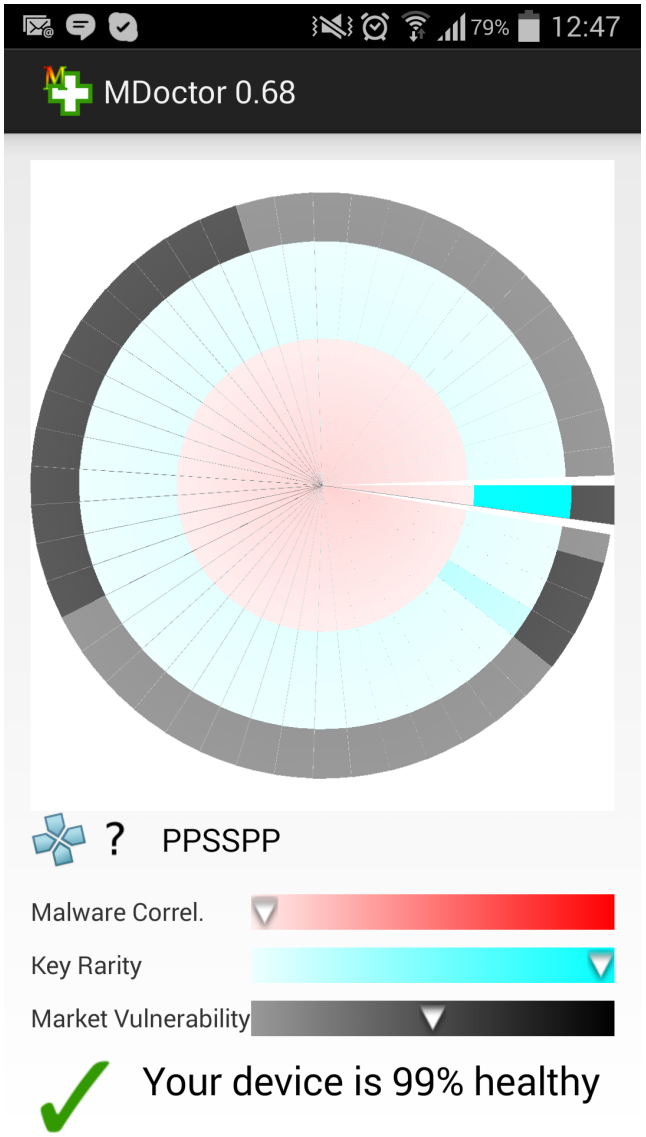

5G Research

Starting point in 2014: LTE RAN and EPC with SDN and Cloud

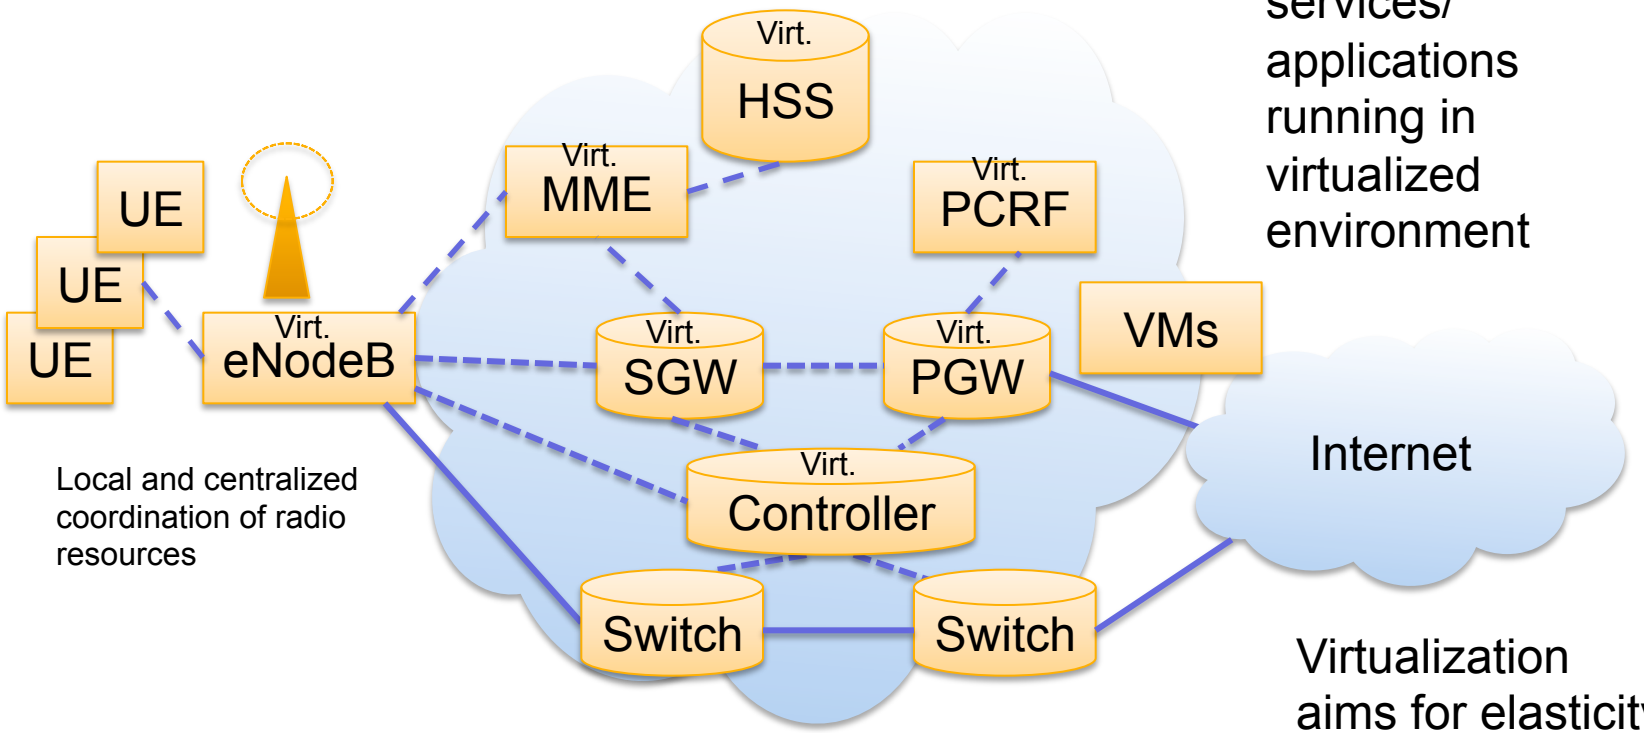

5G elements as services/ applications running in virtualized environment

Local and centralized coordination of radio resources

Virtualization aims for elasticity and runtime configuration support

NOKIA CENTER FOR ADVANCED RESEARCH (NCAR)

NOKIA

HELSINGIN YLIOPISTO
HELSINGFORS UNIVERSITET
UNIVERSITY OF HELSINKI

A?

NCAR was launched in April 2016 and is a joint research center with University of Helsinki, Aalto University, and Nokia.

To foster **wider cooperation** between the universities and Nokia to enable **cross-unit research** delivering high quality results: thesis, publications, holistic concepts and demos.

Highlights

Refactoring Approach for Optimizing Mobile Networks

Coreless Mobile Networks: A state management perspective

Off-the-Shelf Software-defined Wireless Networks

Service Function Chaining: Enabling Technologies and Protocols

Refactoring: Thin Edge

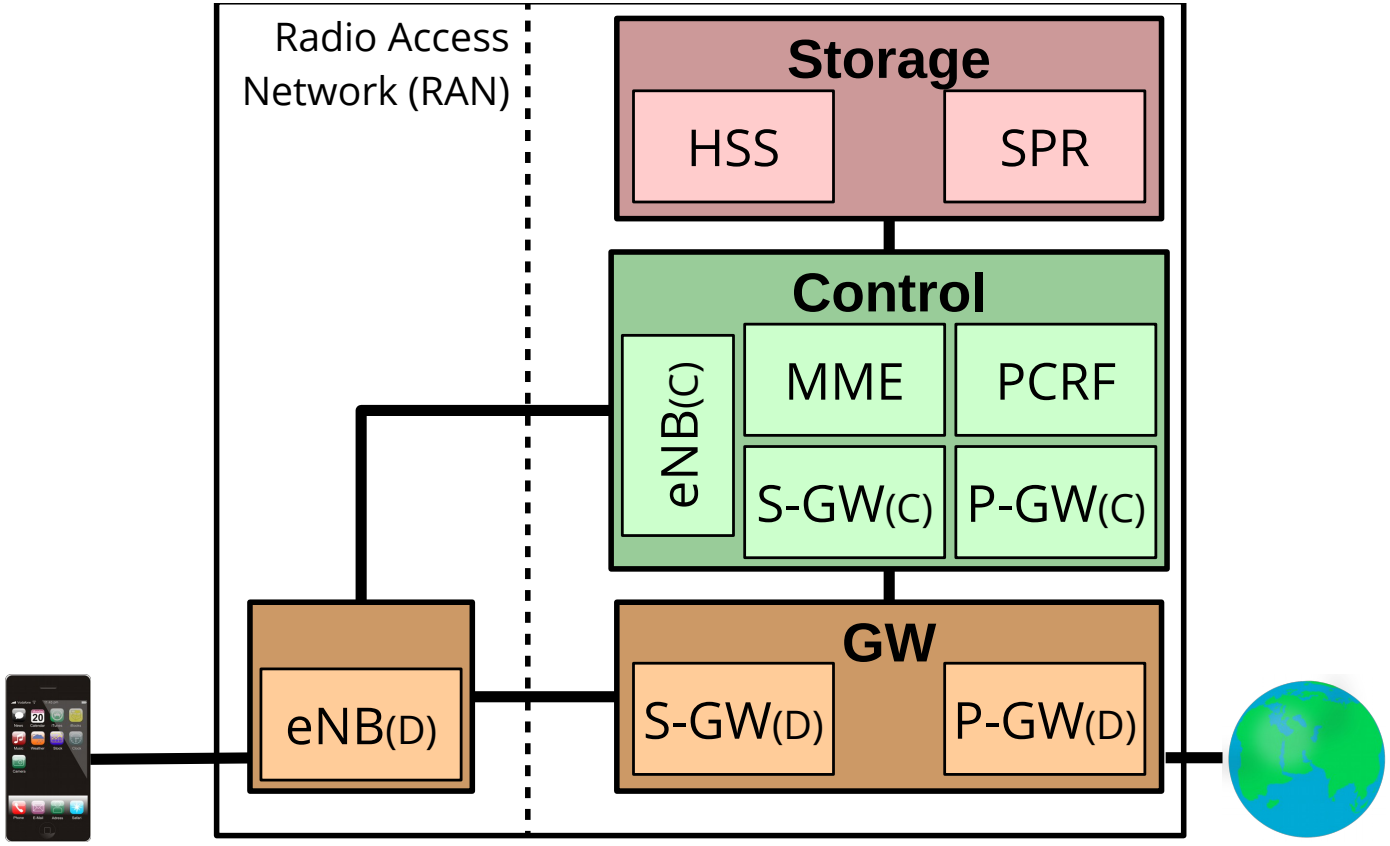

Refactoring: Intelligent Edge

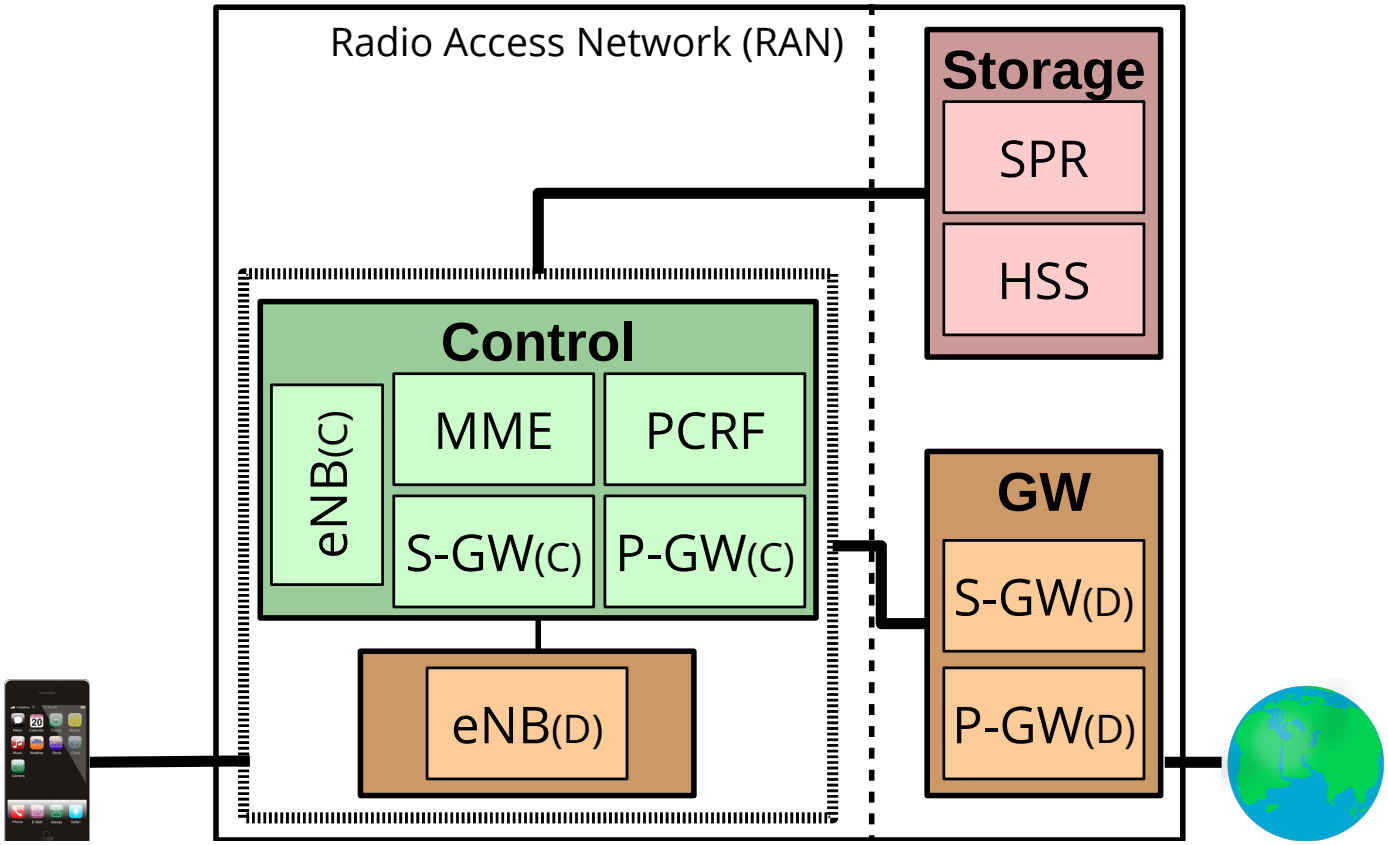

Refactoring Approach for Optimizing Mobile Networks

Implementation	Total number of signals per event: Initial Attach	Active to Idle	Idle to Active (UE)	Idle to Active (Net)	Handover (S1H)
LTE (Baseline)	*35*	*6*	*13*	*17*	*22*
Thin Edge	***24***	***6***	***13***	***16***	***16***
Intelligent Edge	***17***	***3***	***10***	***12***	***12***

A Refactoring Approach for Optimizing Mobile Networks. Matteo Pozza, Ashwin Rao, Armir Abujari, Claudio Pallazi, Hannu Flinck, and Sasu Tarkoma. *Paper in IEEE ICC 2017*

Coreless Mobile Networks: A state management perspective

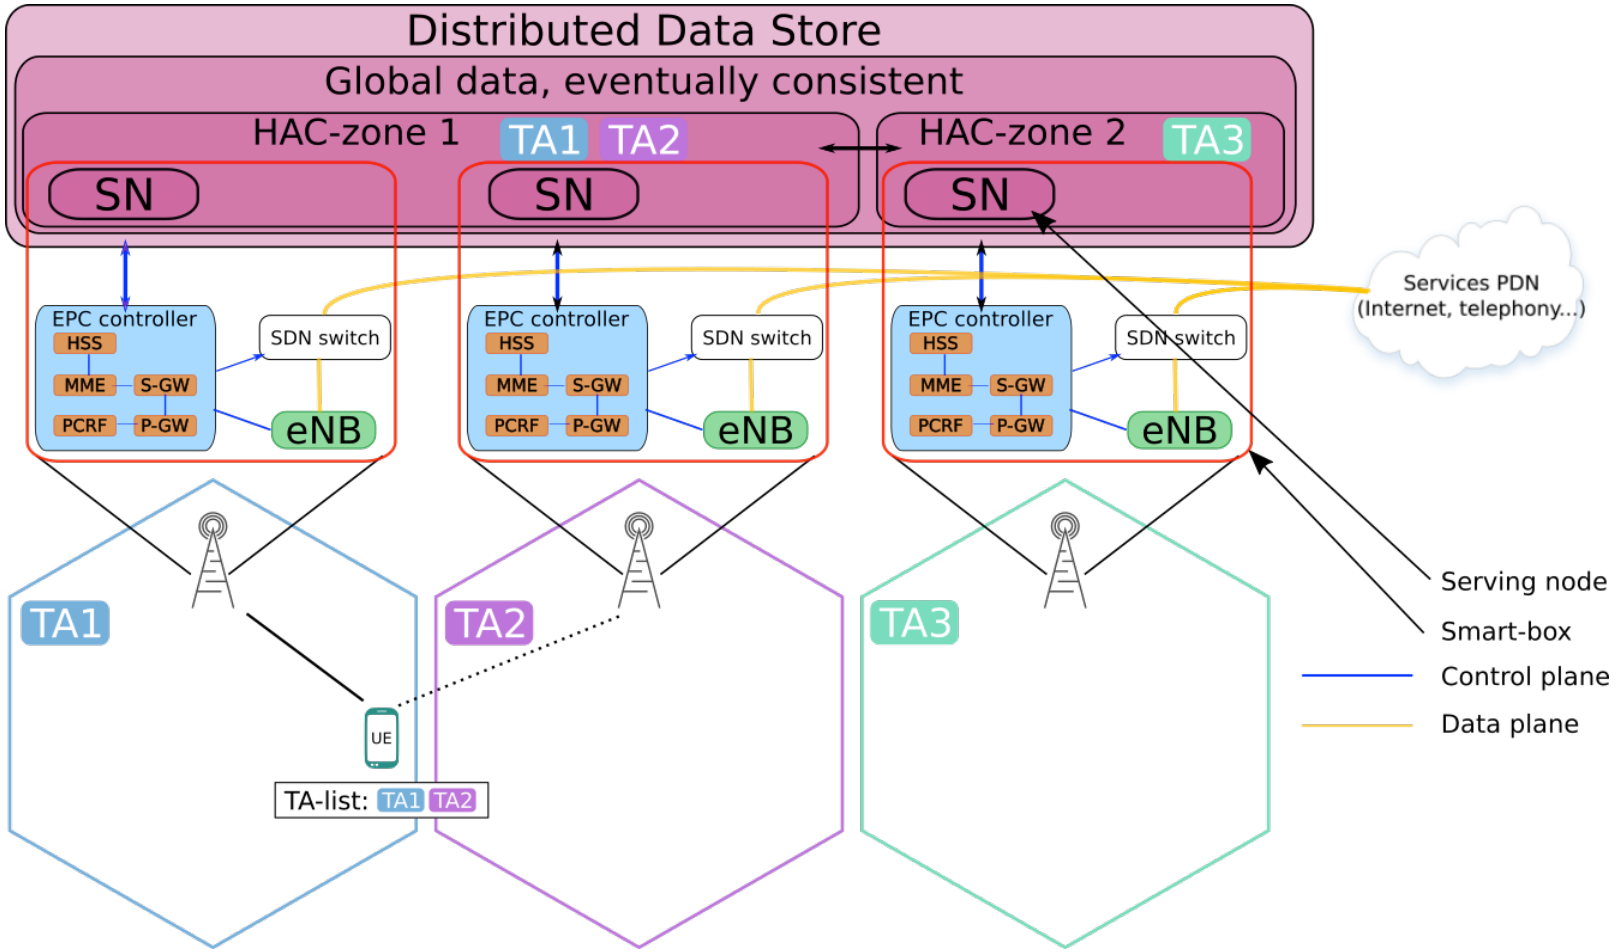

Frans Ojala, 2016

Implications

In theory, if the data store is the bottleneck, our results indicate the following numbers for a simulation of 15 eNB:

Current deployments are seeing a maximum of 1000 UE / eNB

UE per area increases depending on configuration: ~84 - 740 x

5G prospects 100 - 1000 x

Off-the-Shelf Software-defined Wireless Networks

- OVS in base station
- Use Wireless Isolation to force flows to OVS.
  - Either all packets or jus ARP queries depening on implementation.
- Can be used with OpenWrt capable APs or with (at leas some) enterprise APs
- Two approaches, Intelligent and Thin Edge
- Thin edge:
  - Traffic is forced to flow through external host.

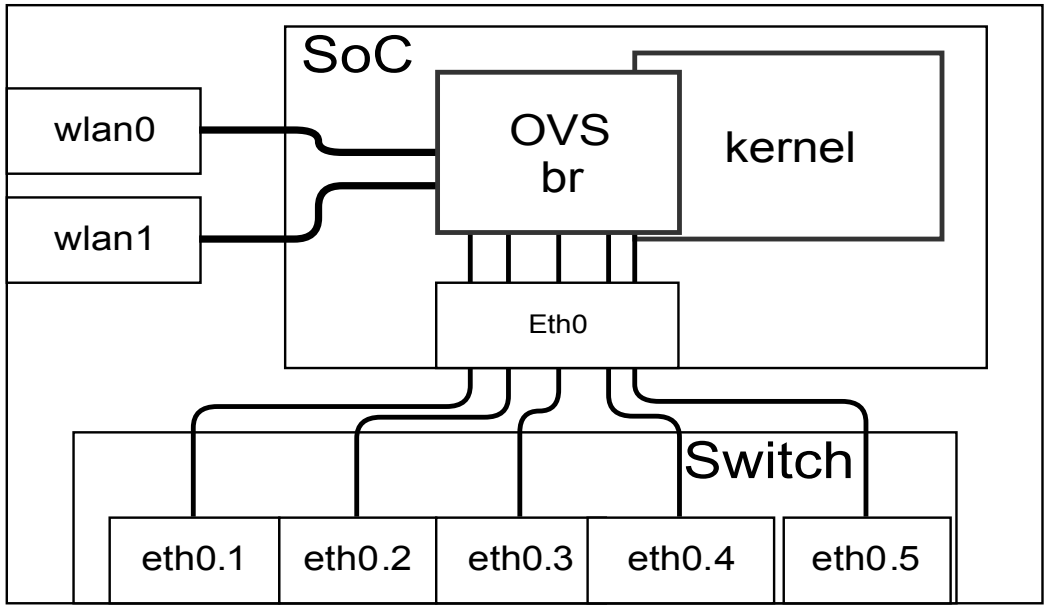

Seppo Hätönen, Petri Savolainen, Ashwin Rao, Hannu Flinck, and Sasu Tarkoma. SIGCOMM 2016 demo.

Deployable on Off-the-Shelf Devices

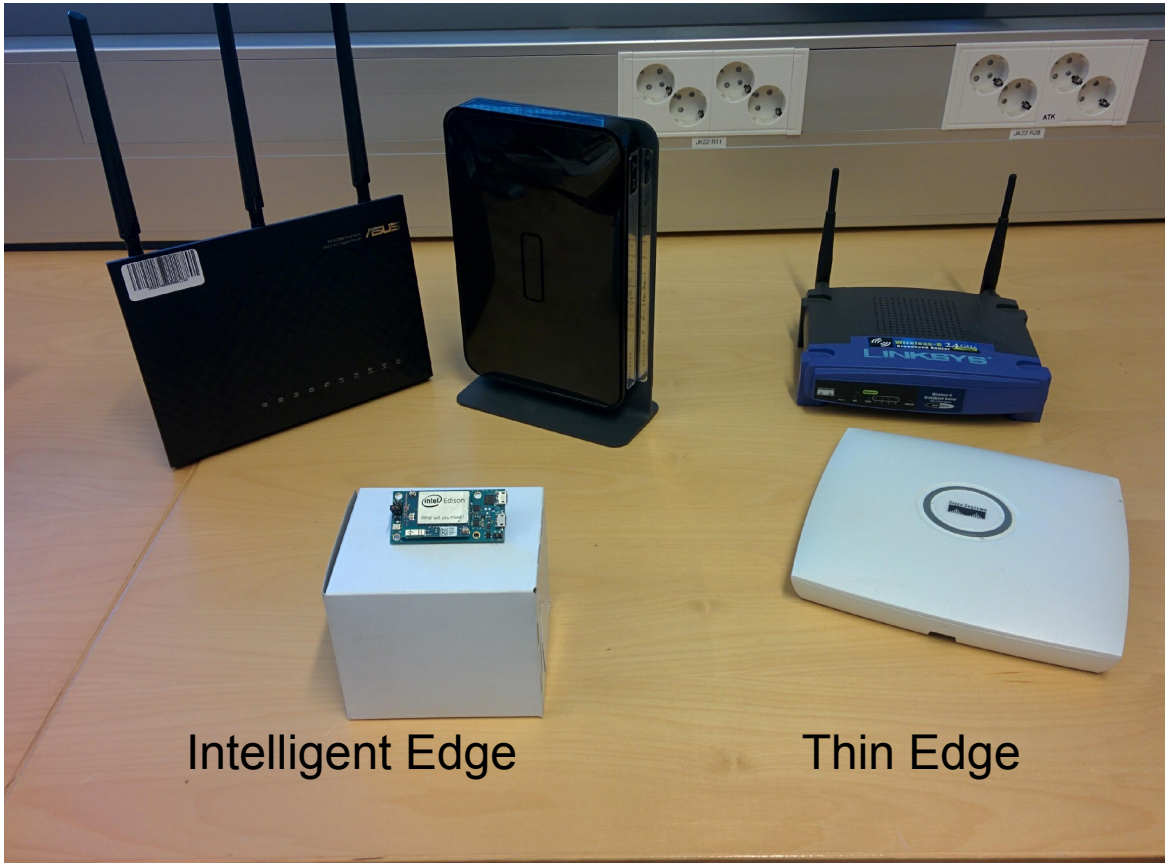

Unified Mobile Edge for IoT Devices

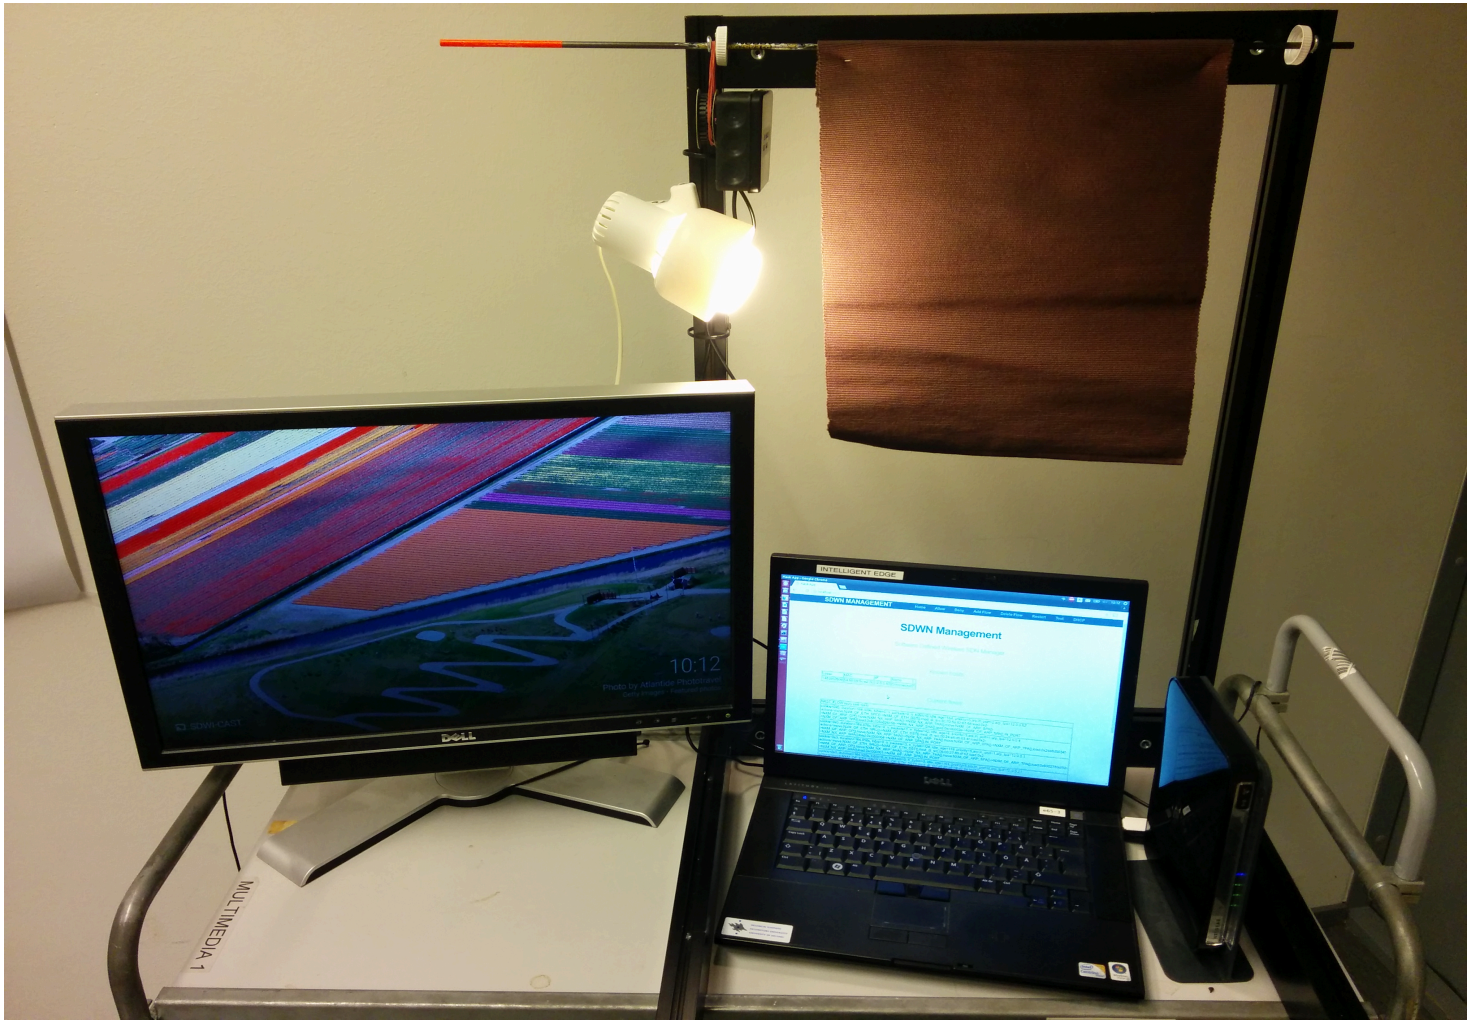

IoT hub running at the edge as an SFC service

Thank You!

www.cs.helsinki.fi
carat.cs.helsinki.fi
ncar.cs.helsinki.fi