The LiveLabs Testbed & Mobile Sensing-based Applications.

Archan Misra,
Singapore Management University
archanm@smu.edu.sg

Feb 17, 2012

Talk Outline

- **LiveLabs: A Mobile Behavioral Experimentation Analogue of PlanetLab**

- **Energy-Efficient Context Acquisition**
  - A3R: Adaptive Accelerometer-based Activity Recognition
  - ACQUA and Distributed Analytics

- **Using Rich, Individual Context**
  - Context-Driven Real-time Femtocell Adaptation
  - CAMEO: Predicting Context for Better Mobile Advertising

SMU
SINGAPORE MANAGEMENT UNIVERSITY

LiveLabs

Globally-unique lifestyle R&D

1. Network technologies for advanced broadband wireless infrastructure.

2. An automated service that lets consumer companies easily run lifestyle experiments.

3. A participant base of 30,000 consumers in 3 key public space (SMU, Malls, Sentosa)

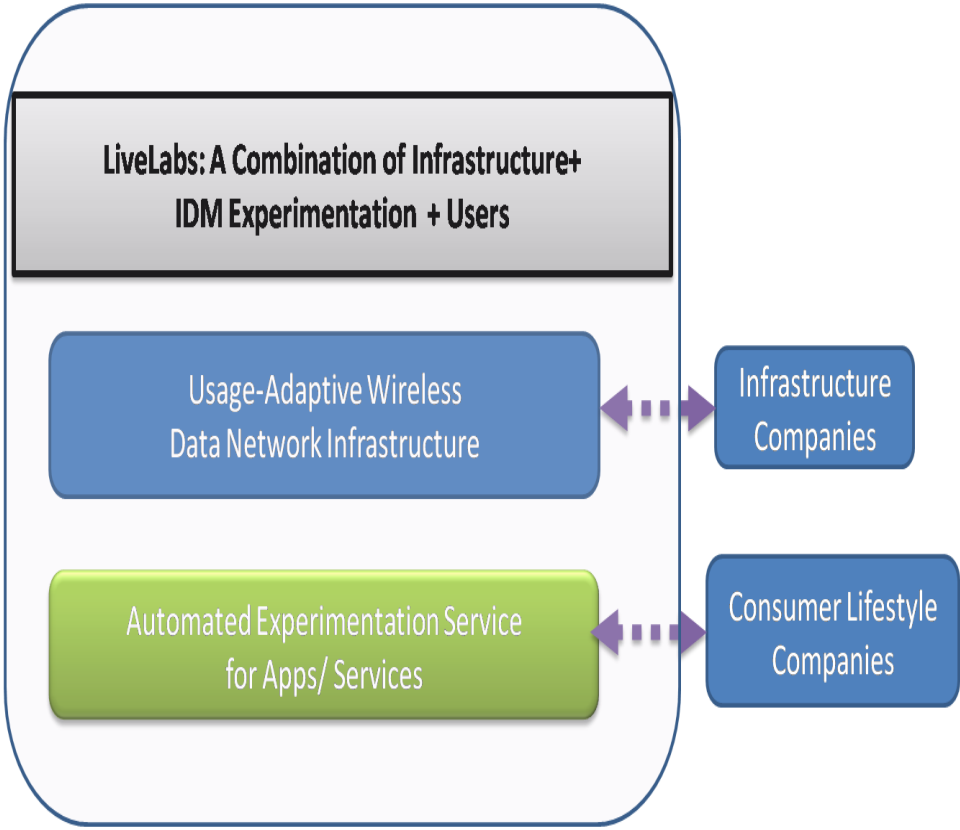

Globally-First Automated Behavioral Experimentation Service

LiveLabs

Globally-Unique Individual and Aggregate Usage-Adaptive Wireless Network

StarHub QUALCOMM DBS sentosa Microsoft Research CapitaMalls Asia ShowNearby accenture

SINGAPORE MANAGEMENT UNIVERSITY | Information Systems

LiveLabs – Downtown Lifestyle Sensing

LiveLabs@Sentosa

sentosa

Tourism & Hospitality

- Crowd Behavior & Movement Optimization
- Personalized Recommendations for Leisure and F&B

QUALCOMM

LiveLabs@ Clarke Quay

LiveLabs@ Plaza Sing

Microsoft Research

Downtown Lifestyle Sensing Testbed:

- **Wireless infrastructure that adapts to real-time usage & hotspots**
- **Behavioral experimentation software**

Retail & Consumption Lifestyle

- Real-time Insight into In-Store+ Online Purchase and Visit Behavior
- Mall Visitor Experience Optimization

Telco & Digital Media

- Quad-Play Offerings Target to Gen-Z Participants
- Rich IDM Delivery under Dynamic Indoor-Outdoor Hotspots

StarHub

DBS

SMU
SINGAPORE MANAGEMENT UNIVERSITY

CapitaMalls Asia

LiveLabs@ SMU

LiveLabs Ecosystem

Examples of Expected Future Users

Committed Users

Globally-First Automated Behavioral Experimentation Service

LiveLabs

Globally-Unique Individual and Aggregate Usage-Adaptive Wireless Network

Key Technology Providers & Users

Key R&D Challenges and Outcomes

- **Challenge 1: Deep, continuous, context collection**
  - Year 1: Collect context from network traces only
  - Year 2: Collect some context from cell phones
  - Year 3: Energy-efficient deep context (cell phones + network)
- **Challenge 2: Fine grained indoor localization**
  - Year 1: 5 to 10m resolution
  - Year 2: 2 to 5m resolution
  - Year 3: <= 1m resolution
- **Challenge 3: Handle transient network traffic loads**
  - Year 1: Offload pre-determined network loads to wired backbone
  - Year 2: Offload network loads to wireless backbones
  - Year 3: Offload traffic based on dynamic traffic patterns
- **Challenge 4: Run automated social experiments on cell phones**
  - Year 1: Build basic framework to run experiments
  - Year 2: Integrate mechanisms to control participant selection
  - Year 3: Integrate end-to-end tools to allow 3rd party developers to use LiveLabs experimentation service
- **Challenge 5: Support privacy preferences of users at runtime**
  - Year 2: Build in mobile device support for privacy enforcement
  - Year 3: Dynamic App checking to enforce context-sensitive privacy

LiveLabs Architecture

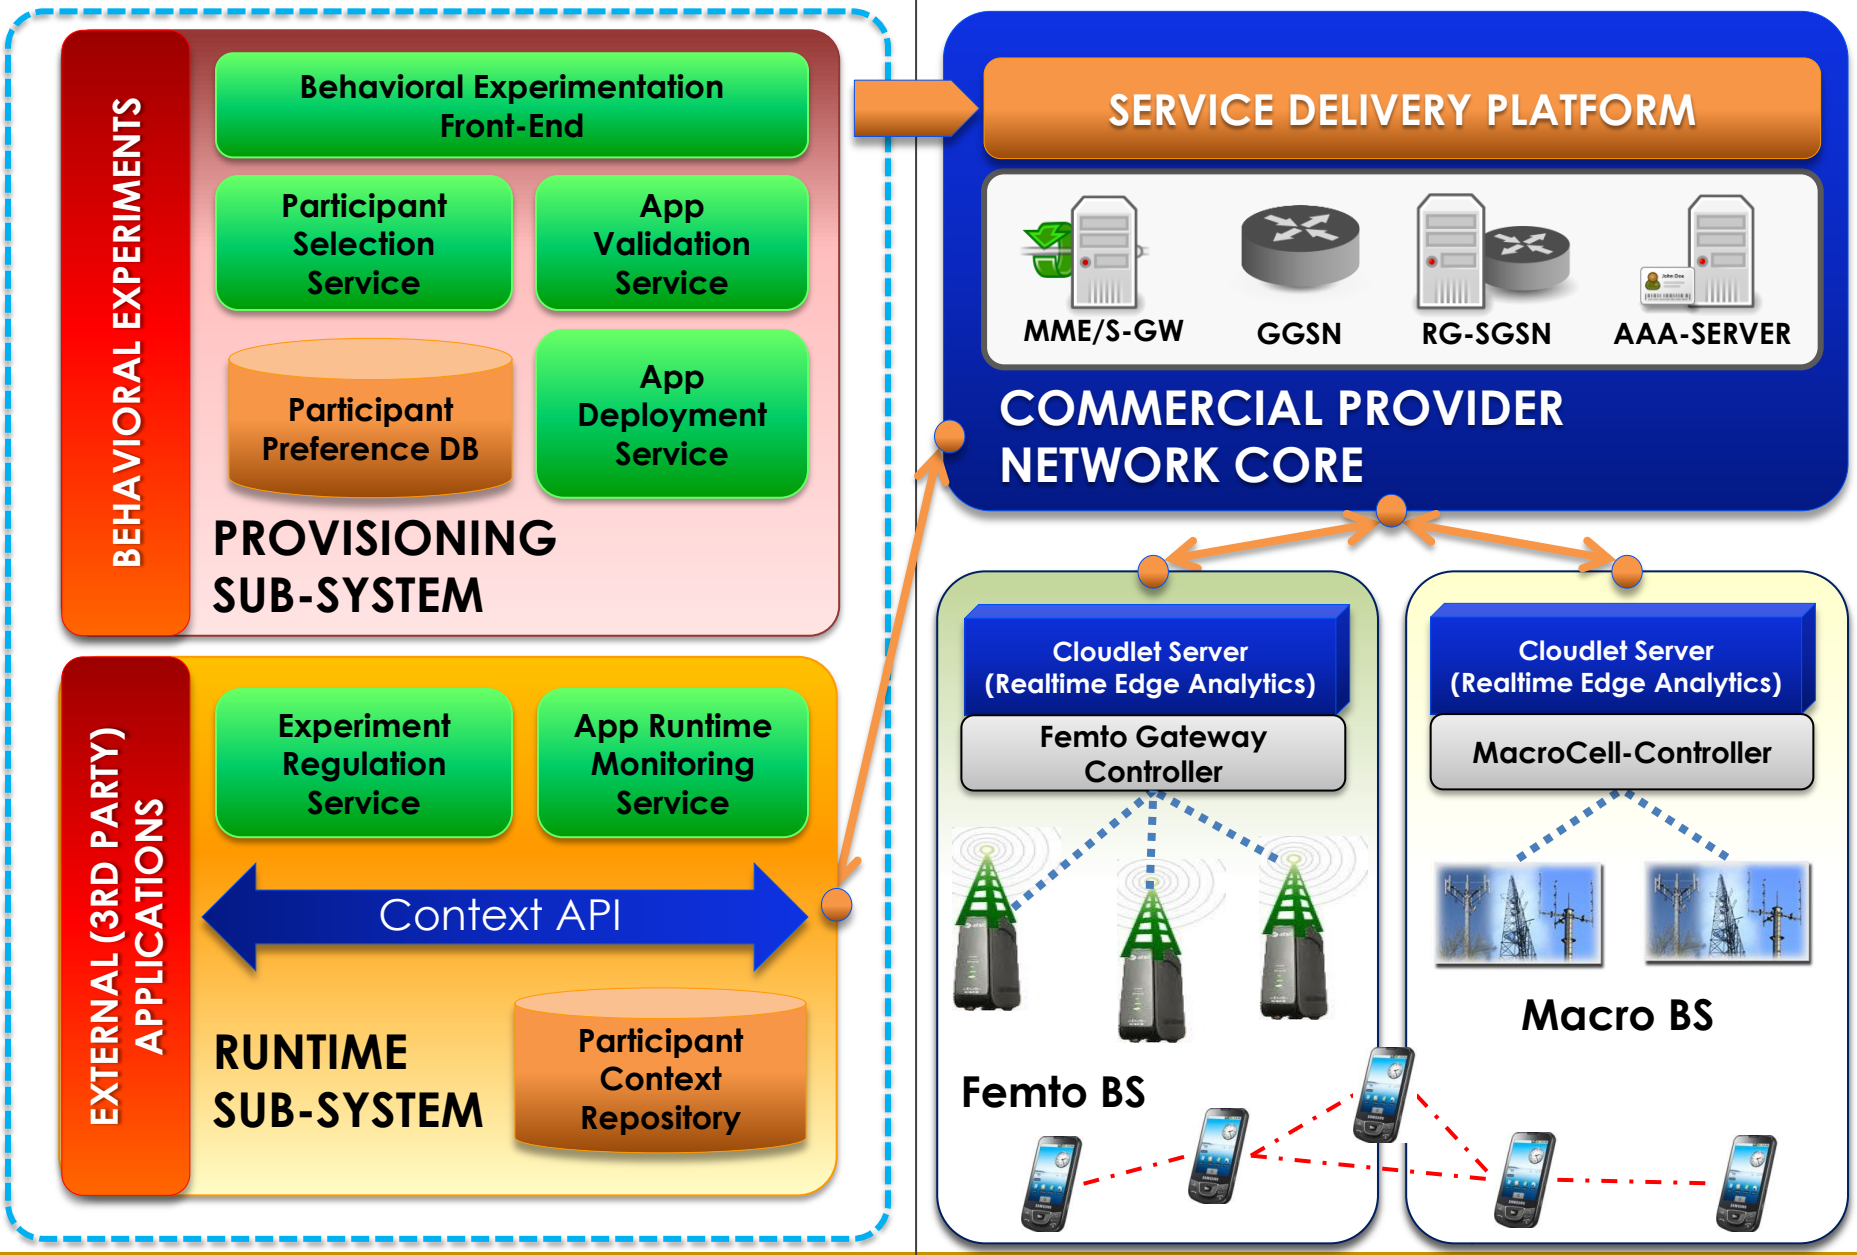

Talk Outline

- **LiveLabs: A Mobile Behavioral Experimentation Analogue of PlanetLab**

- **Energy-Efficient Context Acquisition**
  - A3R: Adaptive Accelerometer-based Activity Recognition
  - ACQUA and Distributed Analytics

- **Using Rich, Individual Context**
  - Context-Driven Real-time Femtocell Adaptation
  - CAMEO: Predicting Context for Better Mobile Advertising

SMU
SINGAPORE MANAGEMENT UNIVERSITY

A3R: Adaptive Accelerometer-based Activity Recognition

- Key Idea: Adjust accelerometer "parameters" based on the current activity of the individual.

- Two parameters:
  - Sampling frequency of accelerometer stream (sf)
  - Features Used for Activity Classification (F)

- Goal: reduce energy overhead of activity recognition without sacrificing accuracy

Energy Overhead Variation

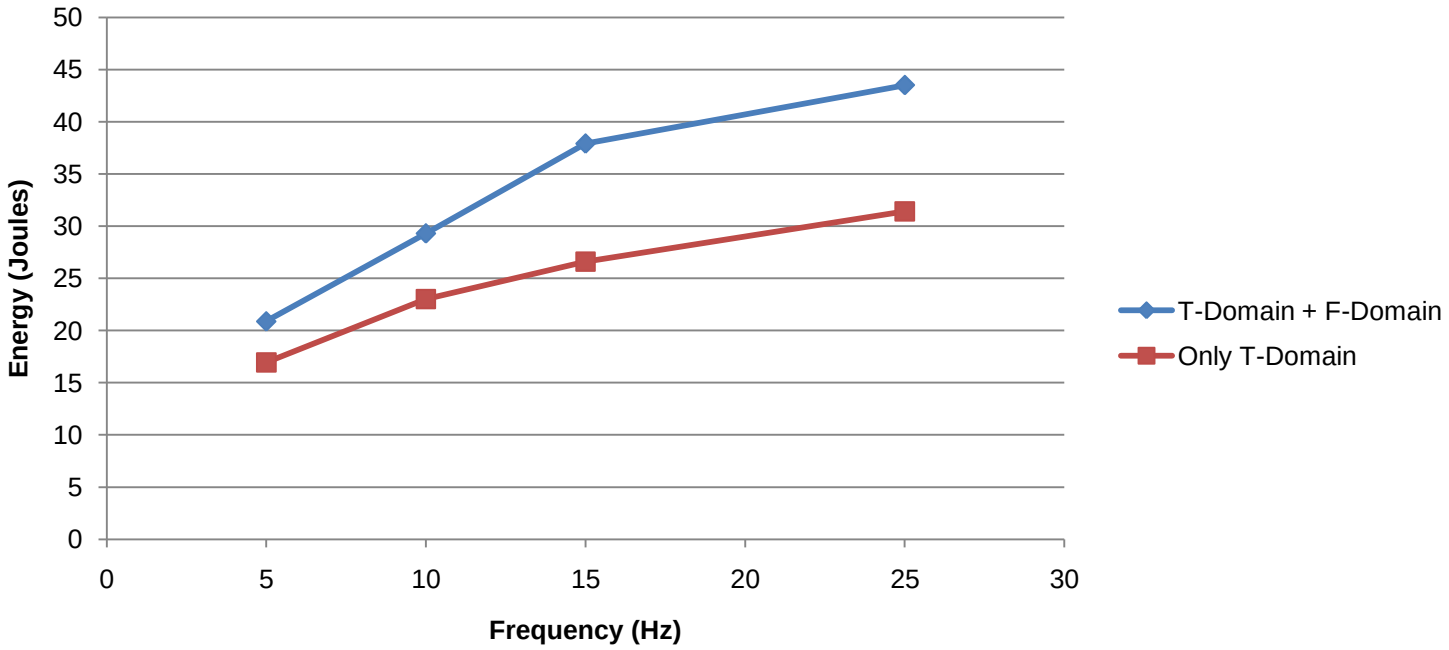

- Energy overhead increases with sf.
- Non-linear increase when frequency-domain features are selected along with time-domain features.

Classification Accuracy Variation

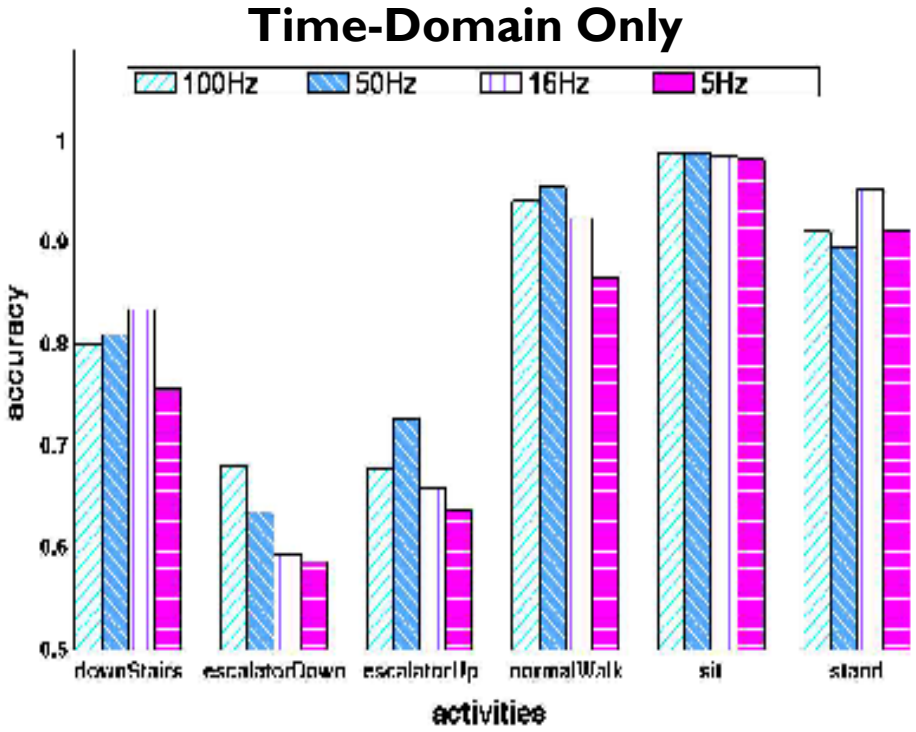

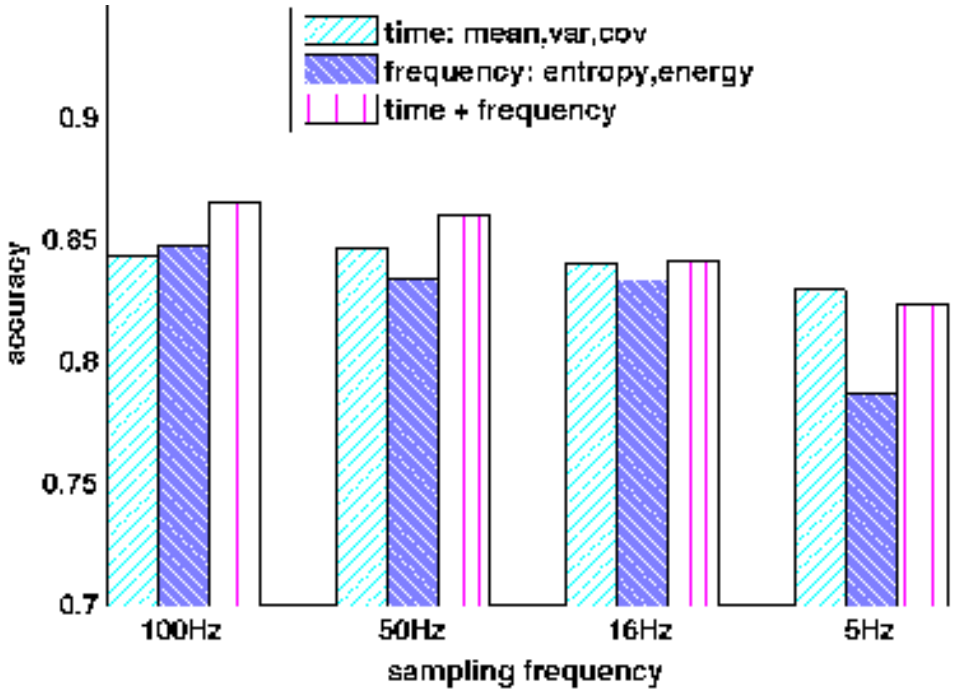

- Most 'stationary' activities (e.g., sit, stand) OK with only sf (1/0.5 Hz).
- Selected activities (e.g., climbing stairs) require (time,frequency) features

SMU Singapore Management University | School of Information Systems

A3R: Results on Real User Behavior

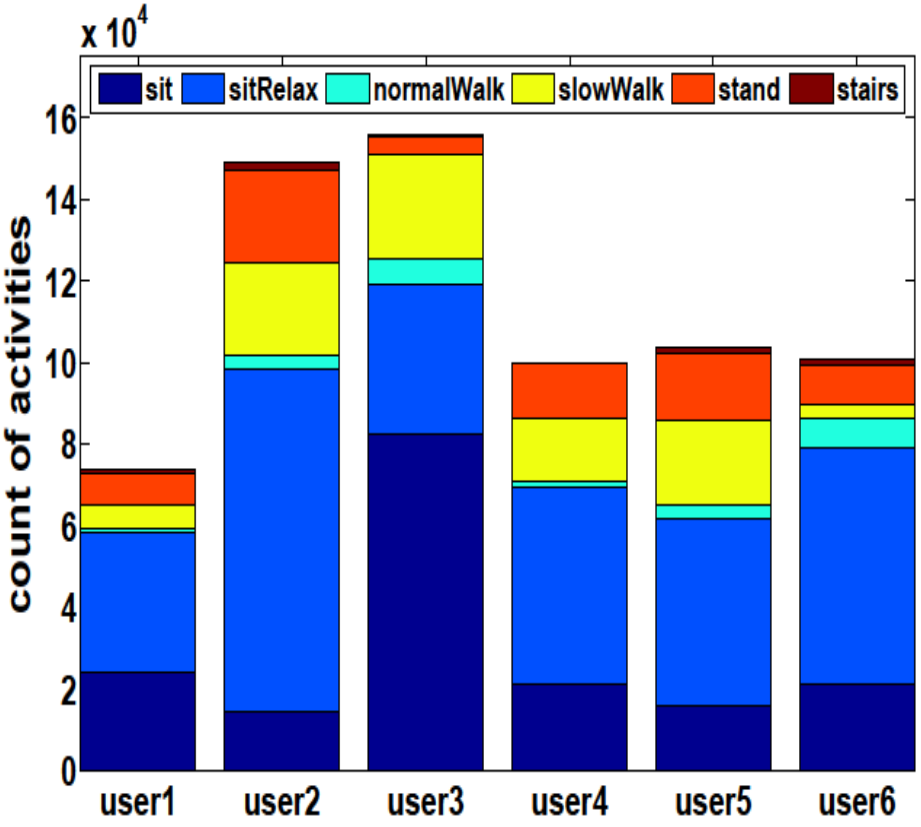

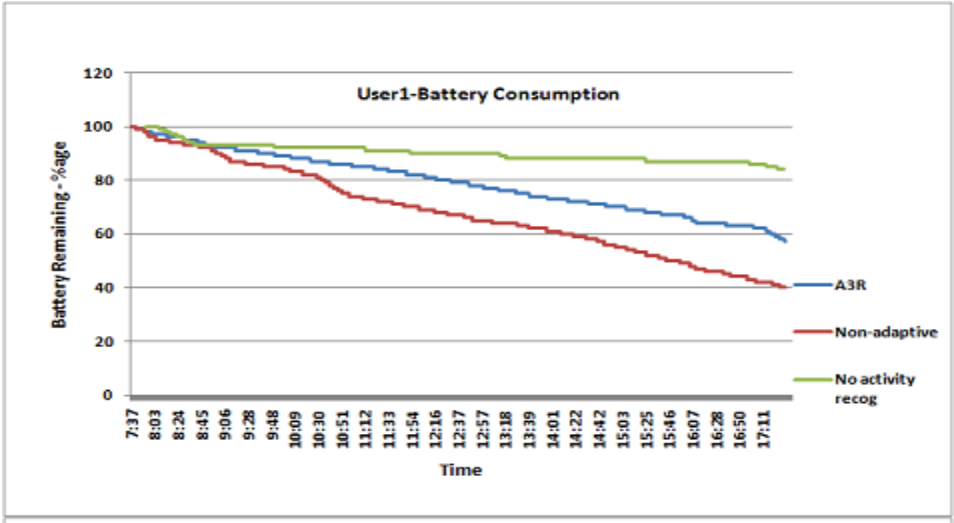

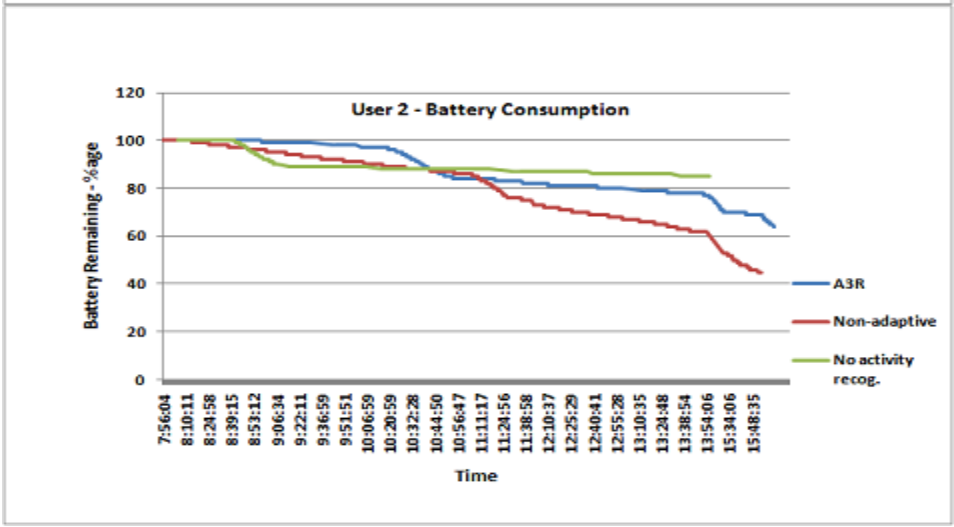

- Over 30% savings in energy under "regular" lifestyle

SMU Singapore Management University | School of Information Systems

ACQUA (Acquisition Cost-Aware Query Adaptation) Scenario

Context deduced from wirelessly connected sensors+ sensors on tother phones

Phone runs a complex event processing (CEP) engine with rules for alerts

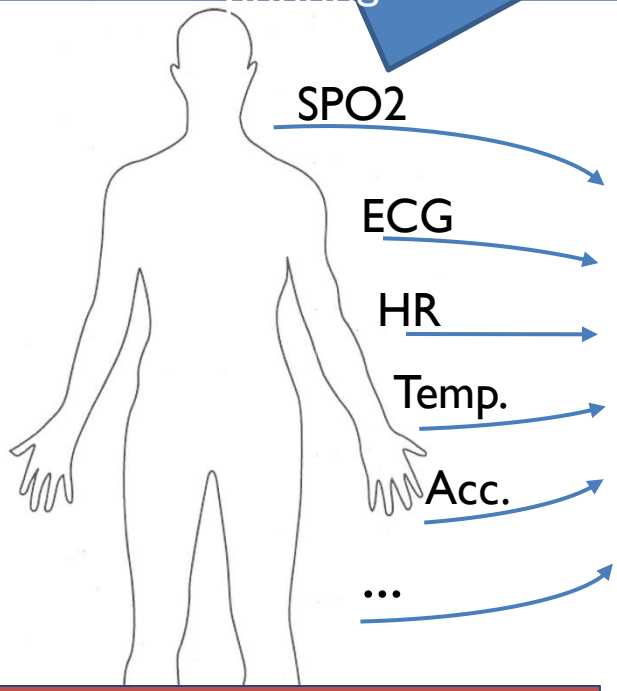

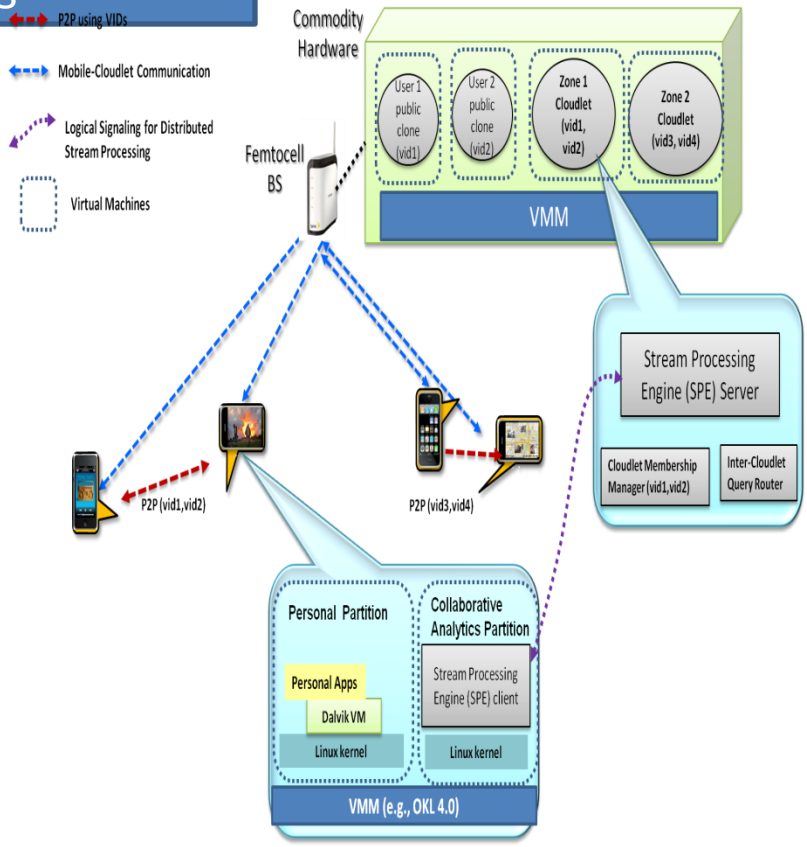

IF Avg(Window(HR)) > 100
AND Avg(Window(Acc)) < 2
AND AVG(Window(Temp))>80F
THEN SMS(caregiver)

ACQUA Architecture

Query Logic Specification Module

- Stream-SQL based specification of query syntax

Cost Modeler

- External specification of sensor-specific trx. Cost model
- Dynamic evaluation of stream selectivity

Normalized Query Syntax

C(.); P(.)

Dynamic Query Evaluation Optimizer

- Signals sensors to adjust push-vs-pull mode
- Determines retrieval sequence for sensor streams

Asynchronous Event Engine

- Maintains partial query evaluation state

Push/Pull, Batch commands

Dynamic Sensor Control (DSC)

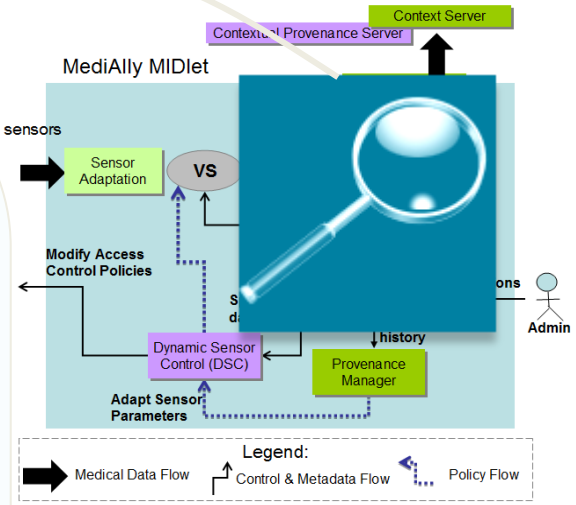

Acquiring N Data-Tuples from Sensor

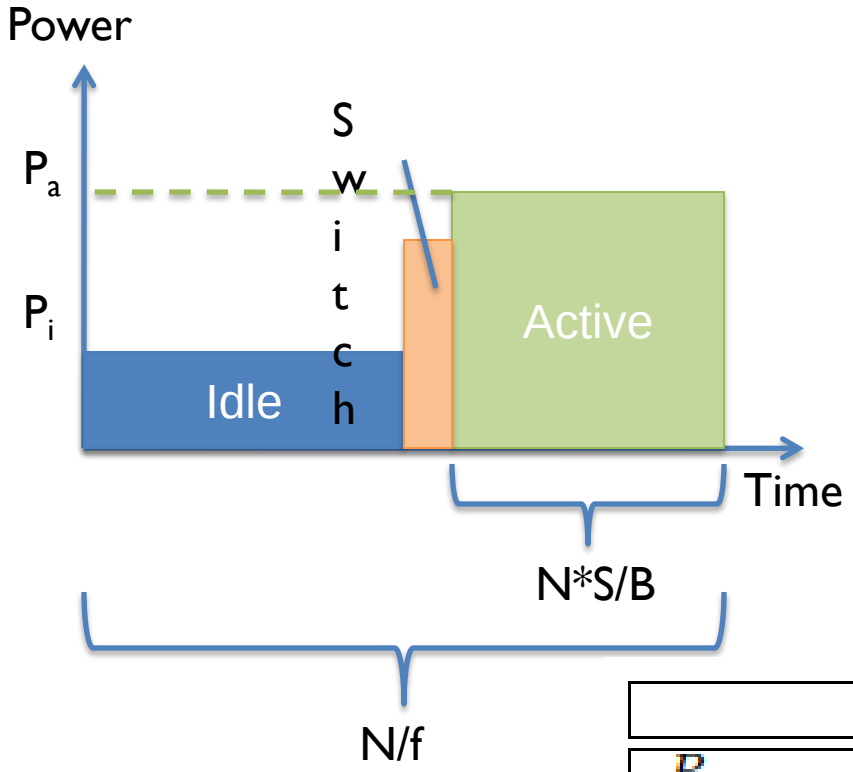

- Idle mode consumes P_i mW
- Active mode consumes P_a mW
- Sensor rate is f Hz
- A tuple is S bits
- Bandwidth is B Mbps

	IEEE 802.11	Bluetooth 2.0+EDR
P_a	947 mW	60mW
P_i	231 mW	5 mW
B	54 Mbps	1 Mbps
E_{switch}	14 μJoule	–
Th_{idle}	100 ms	–
T_{switch}	–	6 msec

Enhanced Evaluation Order

if Avg(S2, 5)>20 AND S1<10 AND Max(S3,10)<4 then email(doctor).

Predicate	Avg(S2, 5)>20	S1<10	Max(S3,10)<4
Acquisition	5 * .02 = 0.1 nJ	0.2 nJ	10 * .01 = 0.1 nJ
Pr(false)	0.95	0.5	0.8
Acq./Pr(f)	0.1/0.95	0.2/0.5	0.1/0.8

- Evaluate predicates with lowest energy consumption first
- Evaluate predicates with highest false probability first
- Evaluate predicate with lowest normalized acquisition cost first.

Performance Results

Bluetooth

802.11

Energy

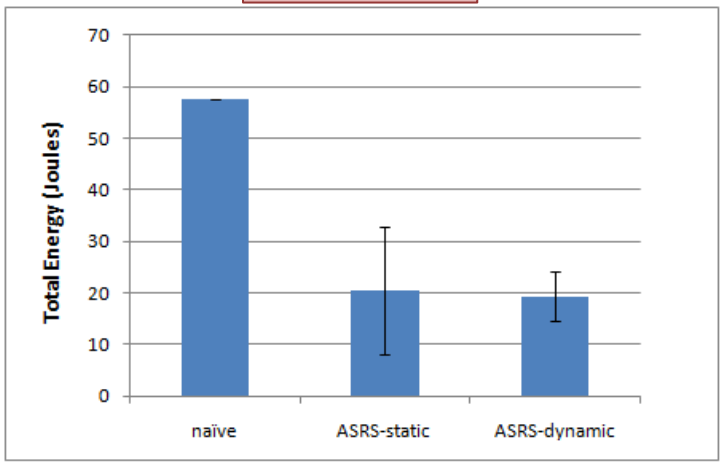

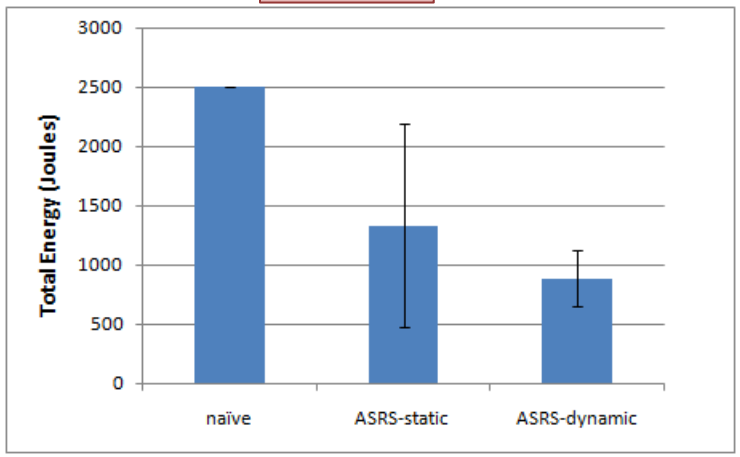

Bytes

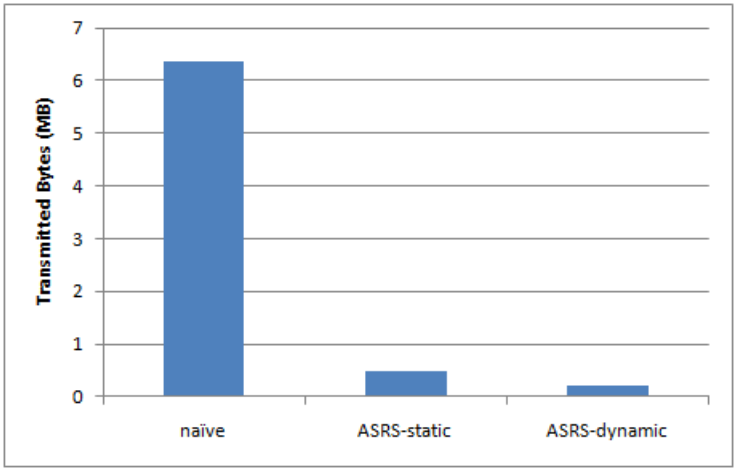

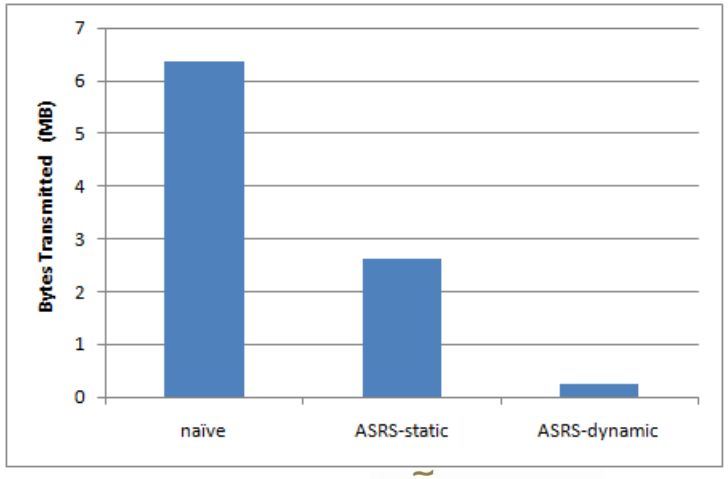

SMU Singapore Management University | School of Information Systems

ProxSense: Distributed Evaluation of CCG Graphs

if Avg(A, 5)<70 AND (C<3 OR Max(B,4)>100 then transmit(location).

if Avg(C, 10)<50 AND (AVG(D,5)>3 AND Max(B,4)>100 then transmit(LocomotionState)

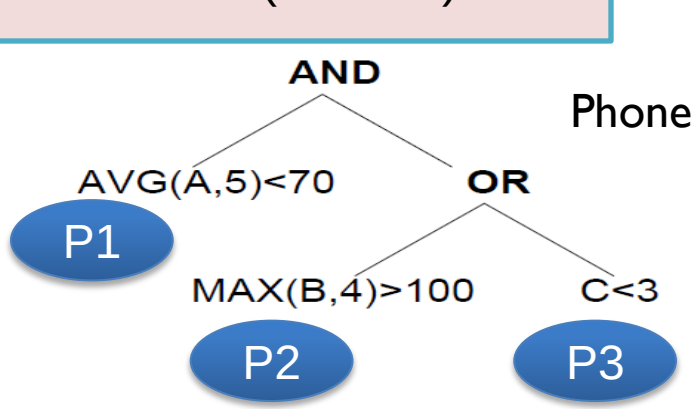

Phone 1

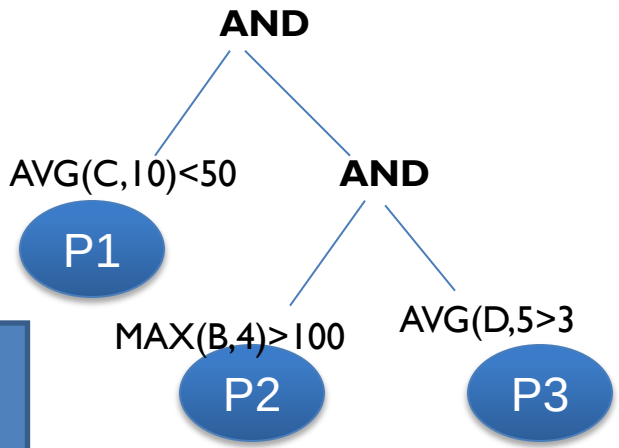

Phone 2

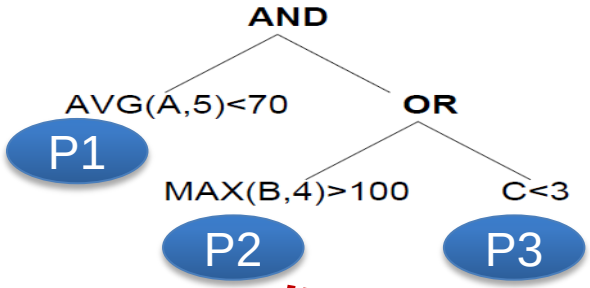

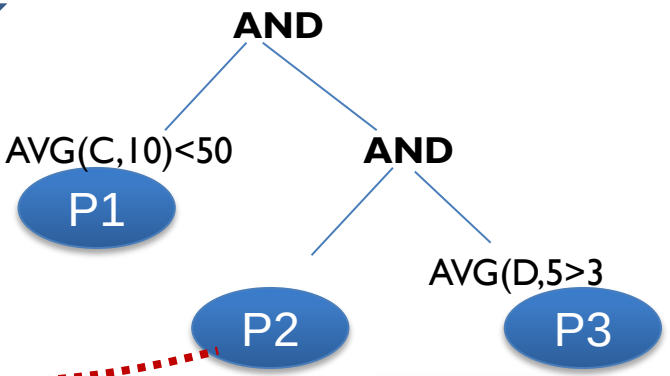

Remote binding and networked commn.

SMU Singapore Management University | School of Information Systems

Talk Outline

- **LiveLabs: A Mobile Behavioral Experimentation Analogue of PlanetLab**

- **Energy-Efficient Context Acquisition**
  - A3R: Adaptive Accelerometer-based Activity Recognition
  - ACQUA and Distributed Analytics

- **Using Rich, Individual Context**
  - Context-Driven Real-time Femtocell Adaptation
  - CAMEO: Predicting Context for Better Mobile Advertising

SMU
SINGAPORE MANAGEMENT UNIVERSITY

The Femto Problem

- Handoff when (RSSI(target)-RSSI(serving)> Th for a period of Ts)
- Fixed Th & Ts can mean:
  - High Th: Fast moving indoor users can take too long to handoff, leading to loss of signal quality and throughput at the cell edge.
  - Low Th: Slow moving indoor users will handoff too soon—random movement can lead to significantly greater ping-pong effect, especially when signal strength diffusion is not uniform.

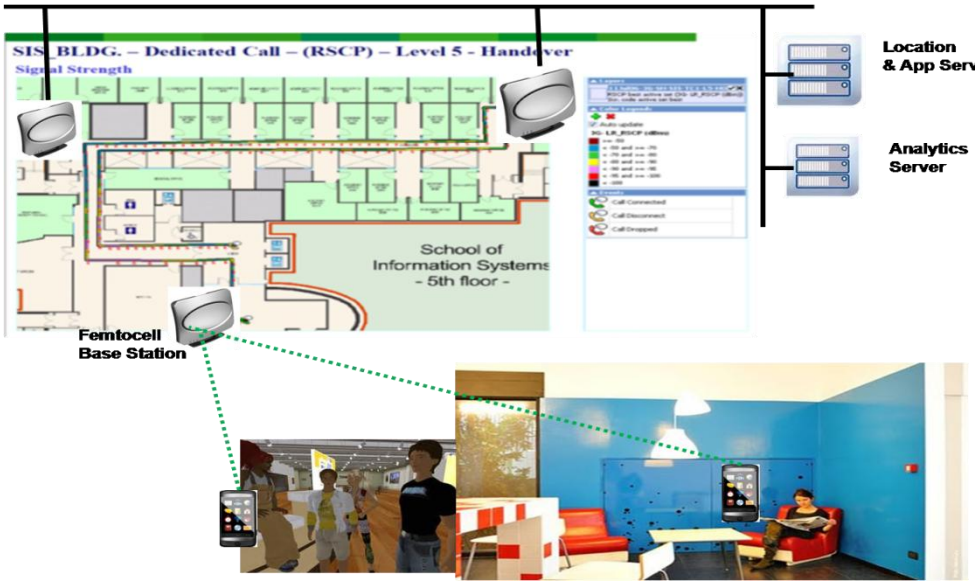

- Lot of work in simulations, but very little captures the practical challenges:
  - Time-varying, anisotropic, RF propagation.
  - Mobile device-based user speed estimation is not perfect.
  - No use of prediction of movement patterns.

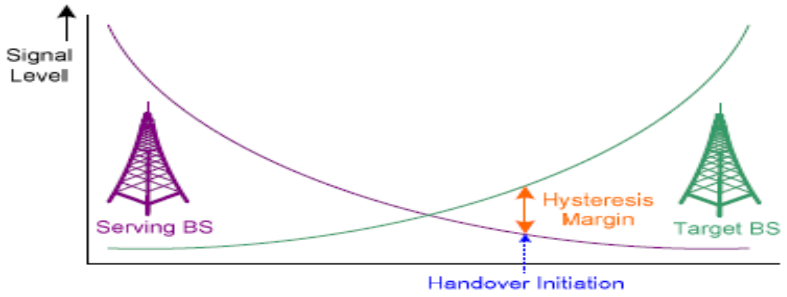

Figure 1. Principle of the conventional hysteresis margin.

SMU Singapore Management University | School of Information Systems

Adaptive High-Bandwidth Indoor Wireless Networks

- **Research Questions:**
  - How to use real-time analytics on collected context to improve future wireless network ability to handle traffic loads?

- **Technical novelty:**
  - *Combine network* (RF) context + mobile-device *user* (RF+sensors+ applications) context to predict network conditions.
  - Dynamically use such *current+ predictive group context* to adapt network parameters

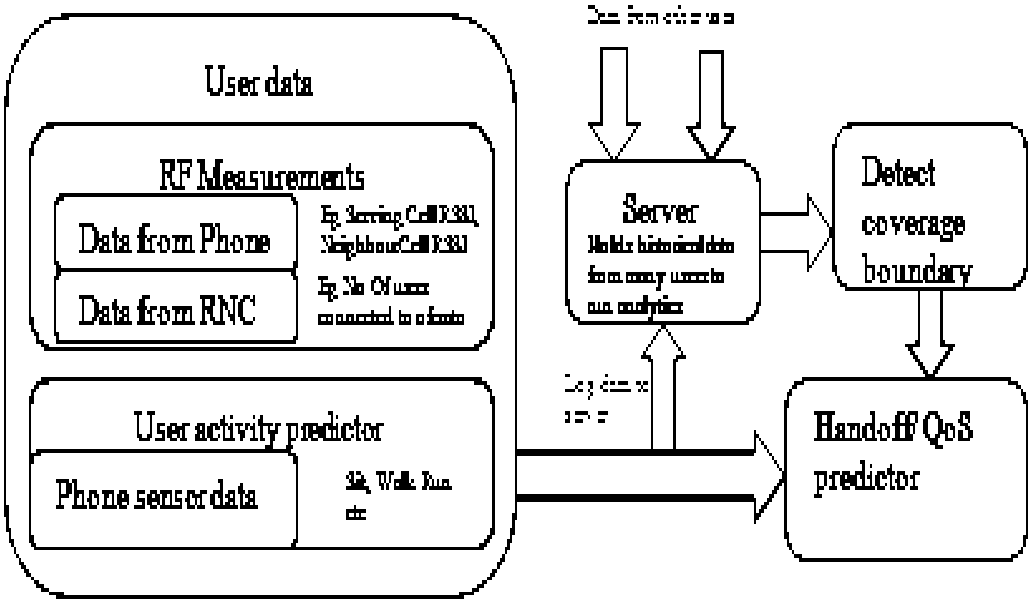

The Real-Time Closed-Loop Context Sensing & Adaptation Framework

SMU Singapore Management University | School of Information Systems

Adaptive Wireless Networks…Progress So Far

- **Deployment:**
  - 6 Femtocell APs deployed on 2 Floors of SIS Building (level 5 and level 3)
- **Emprical Data Collection**
  - Network conditions and parameters collected longitudinally
- **Research Insights:**
  - User movement speed strongly influences network behavior (e.g., handoffs)
  - Indoor environments require different analytics than outdoors.
  - Two new features provide good prediction:
    - No of "DL Power Up" Signals & BLER

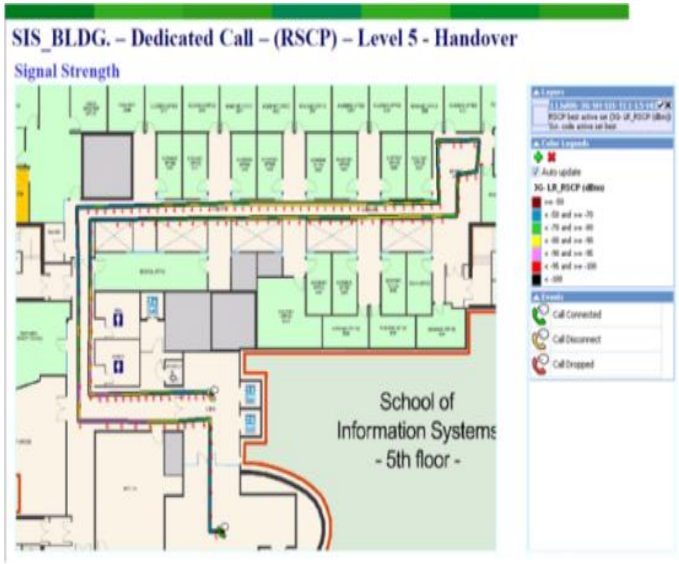

RF Map of Femto Deployment

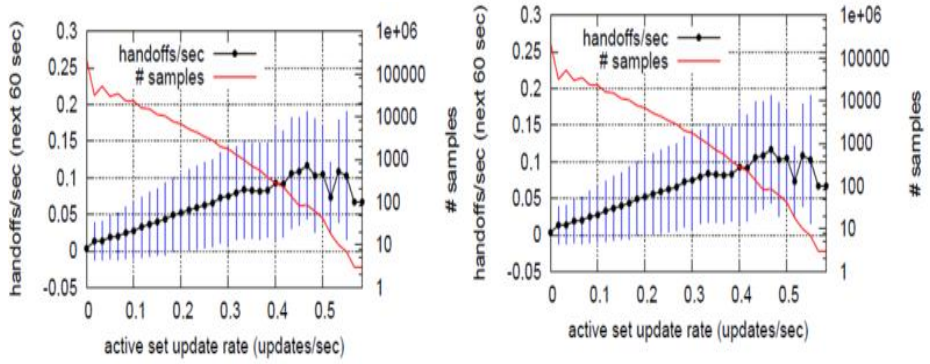

Active Set Update Helps Predict Handoffs (Outdoors)

CAMEO: Optimizing Mobile Advertising

Motivation:
- Ad supported free Apps are very popular.
- Telco providers increasingly moving to metered data plans.

"Free" is not really free!

Rank	Name	Rate of ads (kbps)	Rate of app (kbps)
6	Angry Birds	4.15	0.15
10	Angry Birds Rio	4.03	0.14
10	Hanging with Friends	1.76	5.86
11	Talking Tom 2	1.71	0.02
14	Words with Friends	0.89	2.44
15	Angry Birds Seasons	3.59	0.17
19	Dictionary.com	3.69	9.74

Key Research Idea:
Prefetch Ad Content during Cheap Connectivity (e.g., WiFi@Home) and serve from local cache on phone

Context & Mechanics of Mobile Advertising

Context	Parameters	Presence
Location	latitude and longitude, country, localization, IP address	Always
Device	OS, OS version, model, make, firmware, current time, screen density, screen size, storage size, hardware id, subscriber id, user agent String	Always
Network	telecom carrier, radio, speed, IP address	Always
User	age, gender, subscriber id, language, user-id	Mostly
Others	application id, application version, banner size, format	Mostly

Typical Context Fields in Ad Network SDK

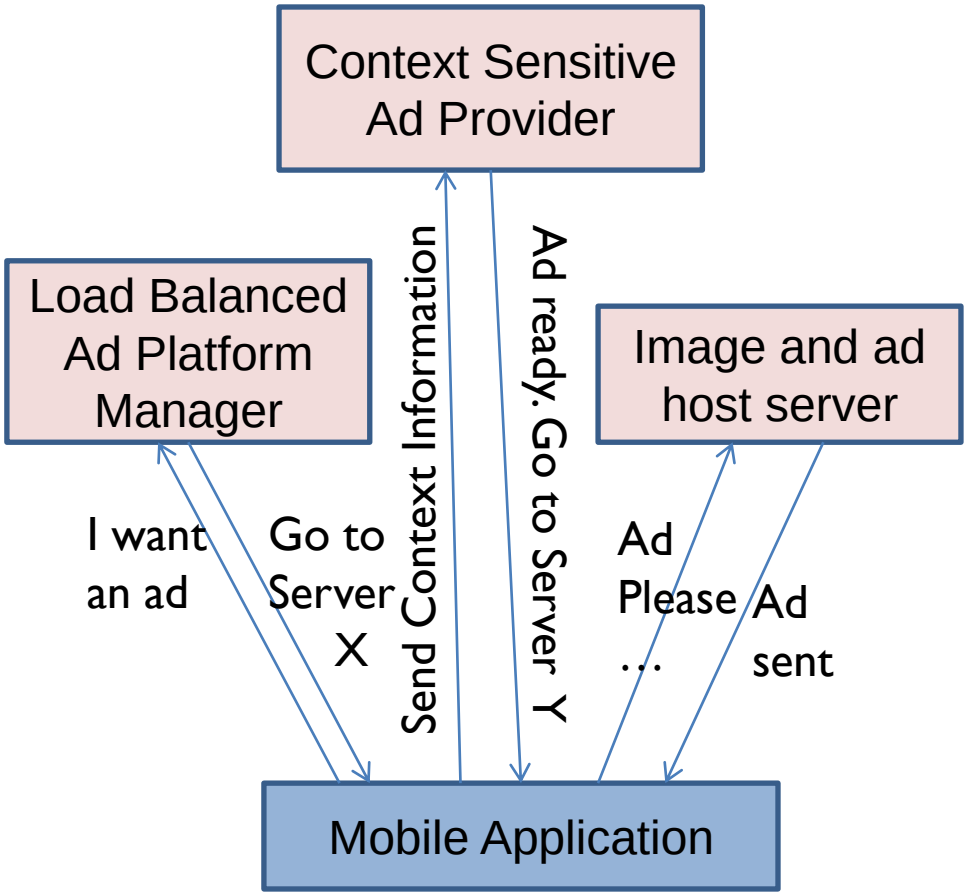

CAMEO: Approach & Architecture

- CAMEO 'predicts' user context (location, application use, etc.)
- CAMEO pre-fetches and caches ads locally using predicted context, when connected via "cheap" networks.
- Ads served locally when application is invoked.

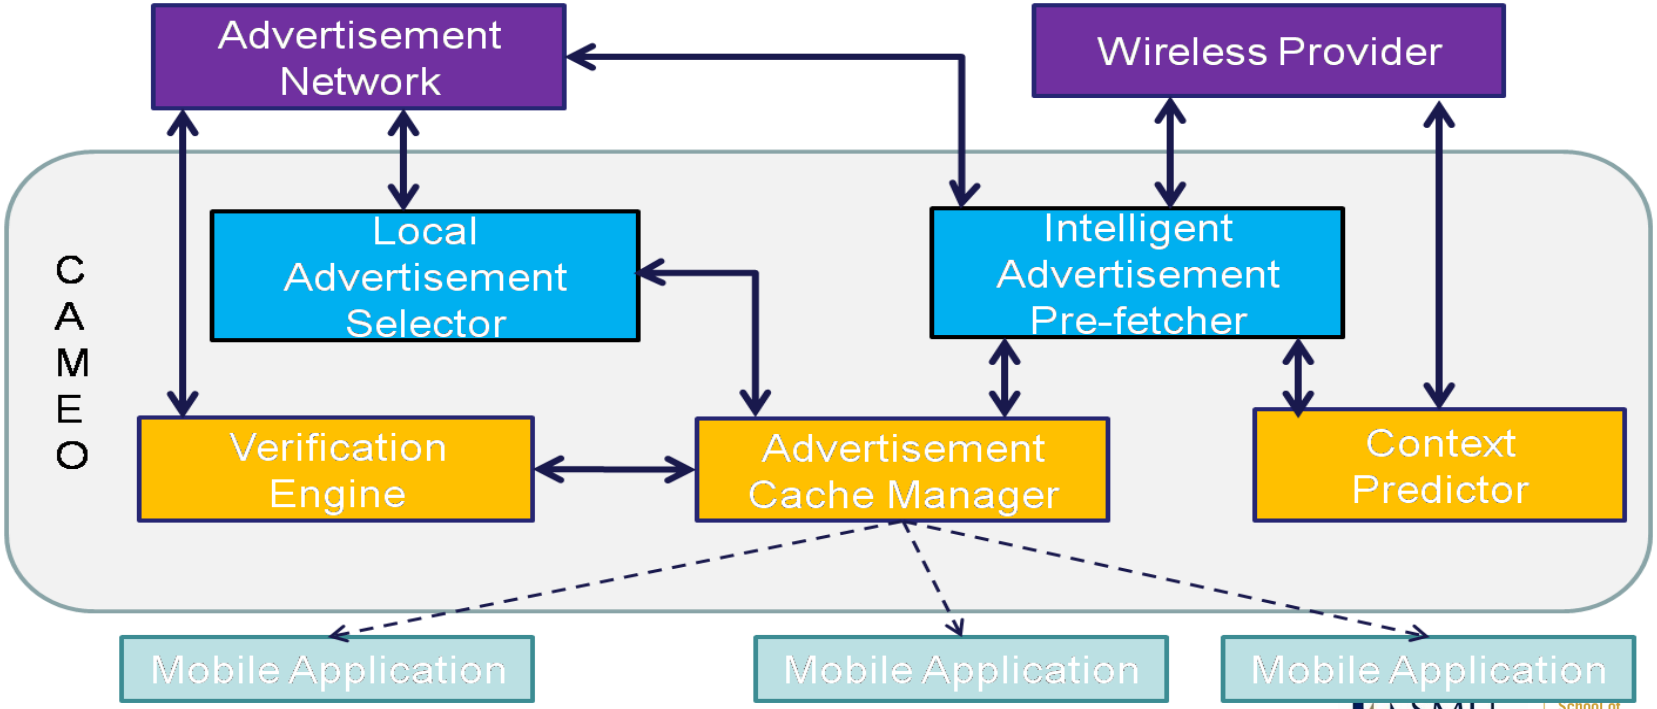

SMU Singapore Management University | School of Information Systems

Conclusions

- LiveLabs is a large-scale testbed for
  - R&D into adaptive context-driven wireless networks and mobile applications
  - Easy experimentation with new services over real users in real indoor/outdoor public spaces.

- **Key technical challenges/advances include:**
  - Energy-efficient sensing and collaboration fusion of activity context
  - Using such context to build usage-adaptive heterogeneous access networks (Macro, femto, WiFi)
  - Using such context to optimize application and content delivery architectures.

- **Other ongoing projects (not covered here):**
  - Accurate (<1m) indoor localization without fingerprinting.
  - Recognition of semantic activities based on low-level sensor-based signatures.

SMU SINGAPORE MANAGEMENT UNIVERSITY | School of Information Systems

Three Key Trends in Mobile Computing

Increased **sensor-richness and display capabilities** in mobile devices

- Tablet sales to eclipse laptop sales by 2012.
- Embedded sensor market doubling each yer 2012-2014

Emergence of proximal P2P Among Mobile Devices

- FlashLinQ radios from Qualcomm
- SocialWiFi proposal from WiFi Forum

Processing at the Edge of the Network (Better responsiveness & scaling)

- Linux-based processors on gen-2 femtos
- VM-based Cloudlet for personalized offloading

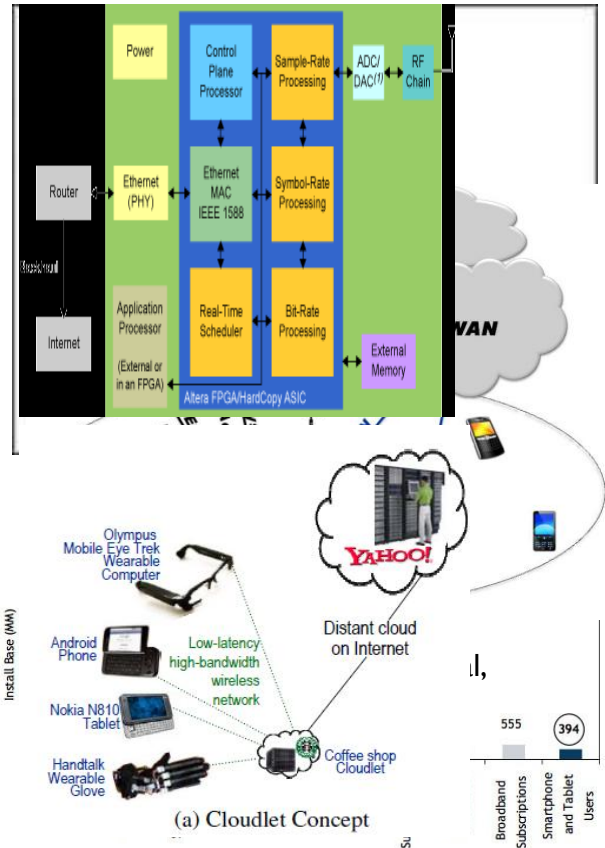

Source: M. Satyanarayana, et al, The Case for VM-based Cloudlets in Mobile Computing, IEEE Pervasive Computing, 2009.

SMU Singapore Management University | School of Information Systems

Acknowledgements & References

Project Name	Collaborators	Publication Reference
ACQUA	Lipyeow Lim (Univ. of Hawaii)	A. Misra and L. Lim, “Optimizing Sensor Data Acquisition for Energy-Efficient Smartphone-based Continuous Event Processing”, IEEE MDM, 06/2011
A3R	Zhixian Yan, Dipanjan Chakraborty and Karl Aberer (EPFL), Vigneshwaran Subbaraju (SMU)	Under submission.
Femtocell Adaptation	Srini Seshan (CMU), Vigneshwaran Subbaraju (SMU)	---In Preparation.
CAMEO	Srini Seshan (CMU) Azeem Khan (SMU), Vigneshwaran Subbaraju (SMU)	HotMobile 2012

SMU Singapore Management University | School of Information Systems

Collaborative Sensing & the "Mobile 3.0" Computing Architecture

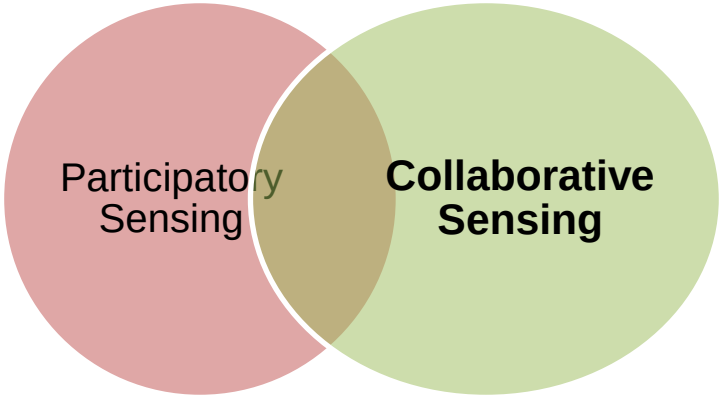

- Collaborative for individual benefit
- Near-real-time analytics
- Localized in space & time

SMU Singapore Management University | School of Information Systems