

Contributions to the Autonomy of Ubiquitous Software Systems

Romain Rouvoy

Habilitation defense (HDR)

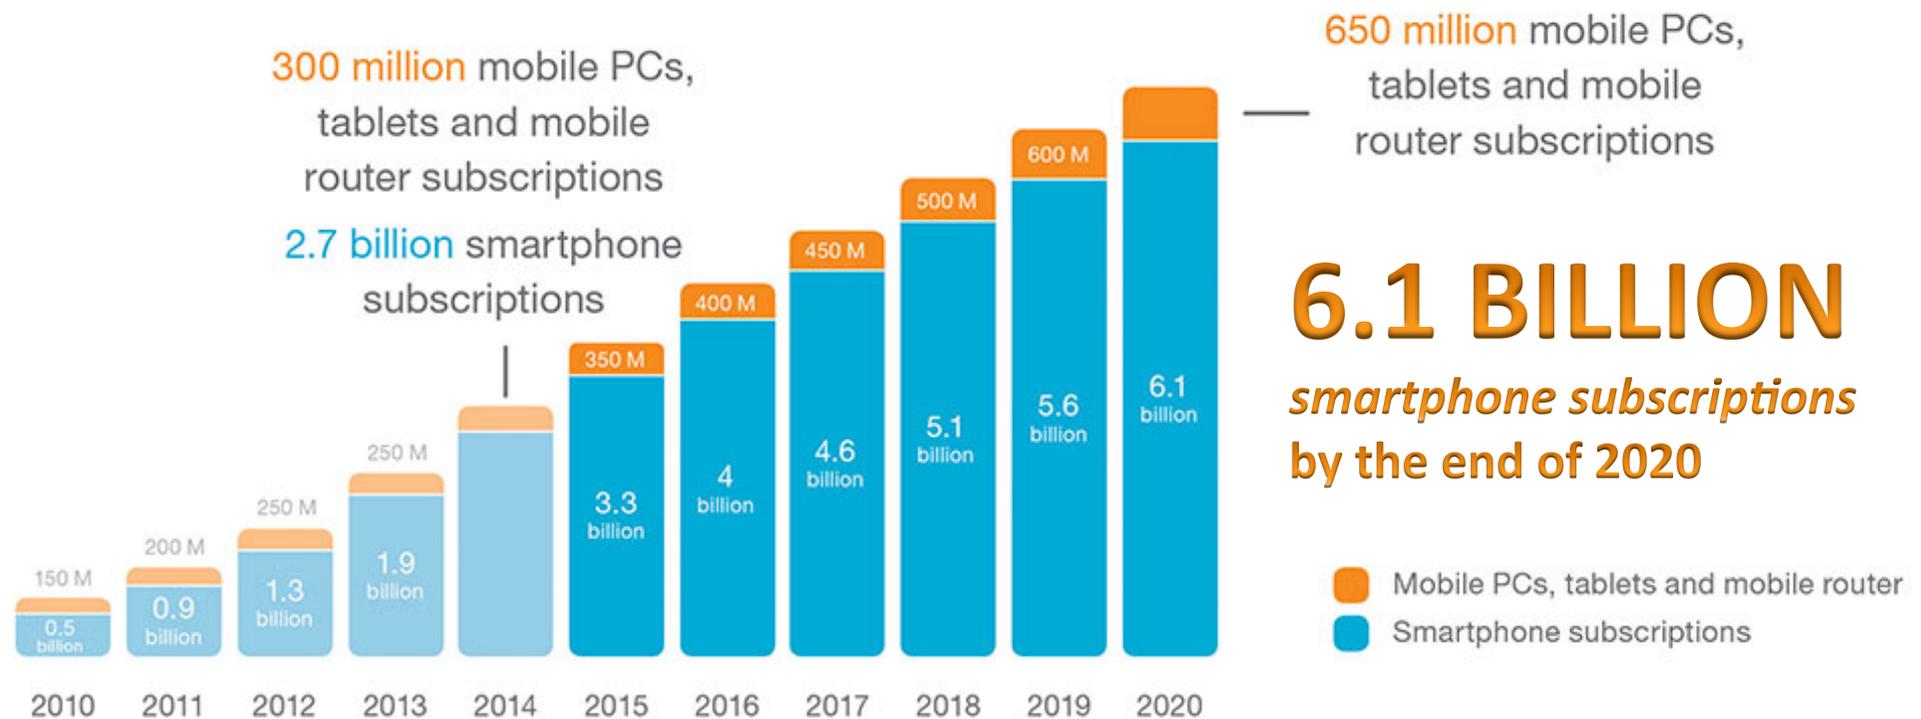
Friday 5th December 2014

A few facts about me

- **2003–06: Ph.D. thesis** from USTL / INRIA
 - *Supervision: P. Merle & J.M. Geib*
 - *Topic: Middleware / components / transaction*
- **2007–08: Postdoc researcher @ Univ. of Oslo**
 - *Member of the ND group (IFI)*
 - *Contributions to the EU FP6 MUSIC R&D project*
- **2008– : Associate professor @ Univ. of Lille 1**
 - *Member of the Spirals team (Lille 1 & Inria)*
 - *Head of the CS Master on software engineering*
 - *Distributed systems & software engineering*
 - *Middleware for ubiquitous computing*

Evolution of connected devices

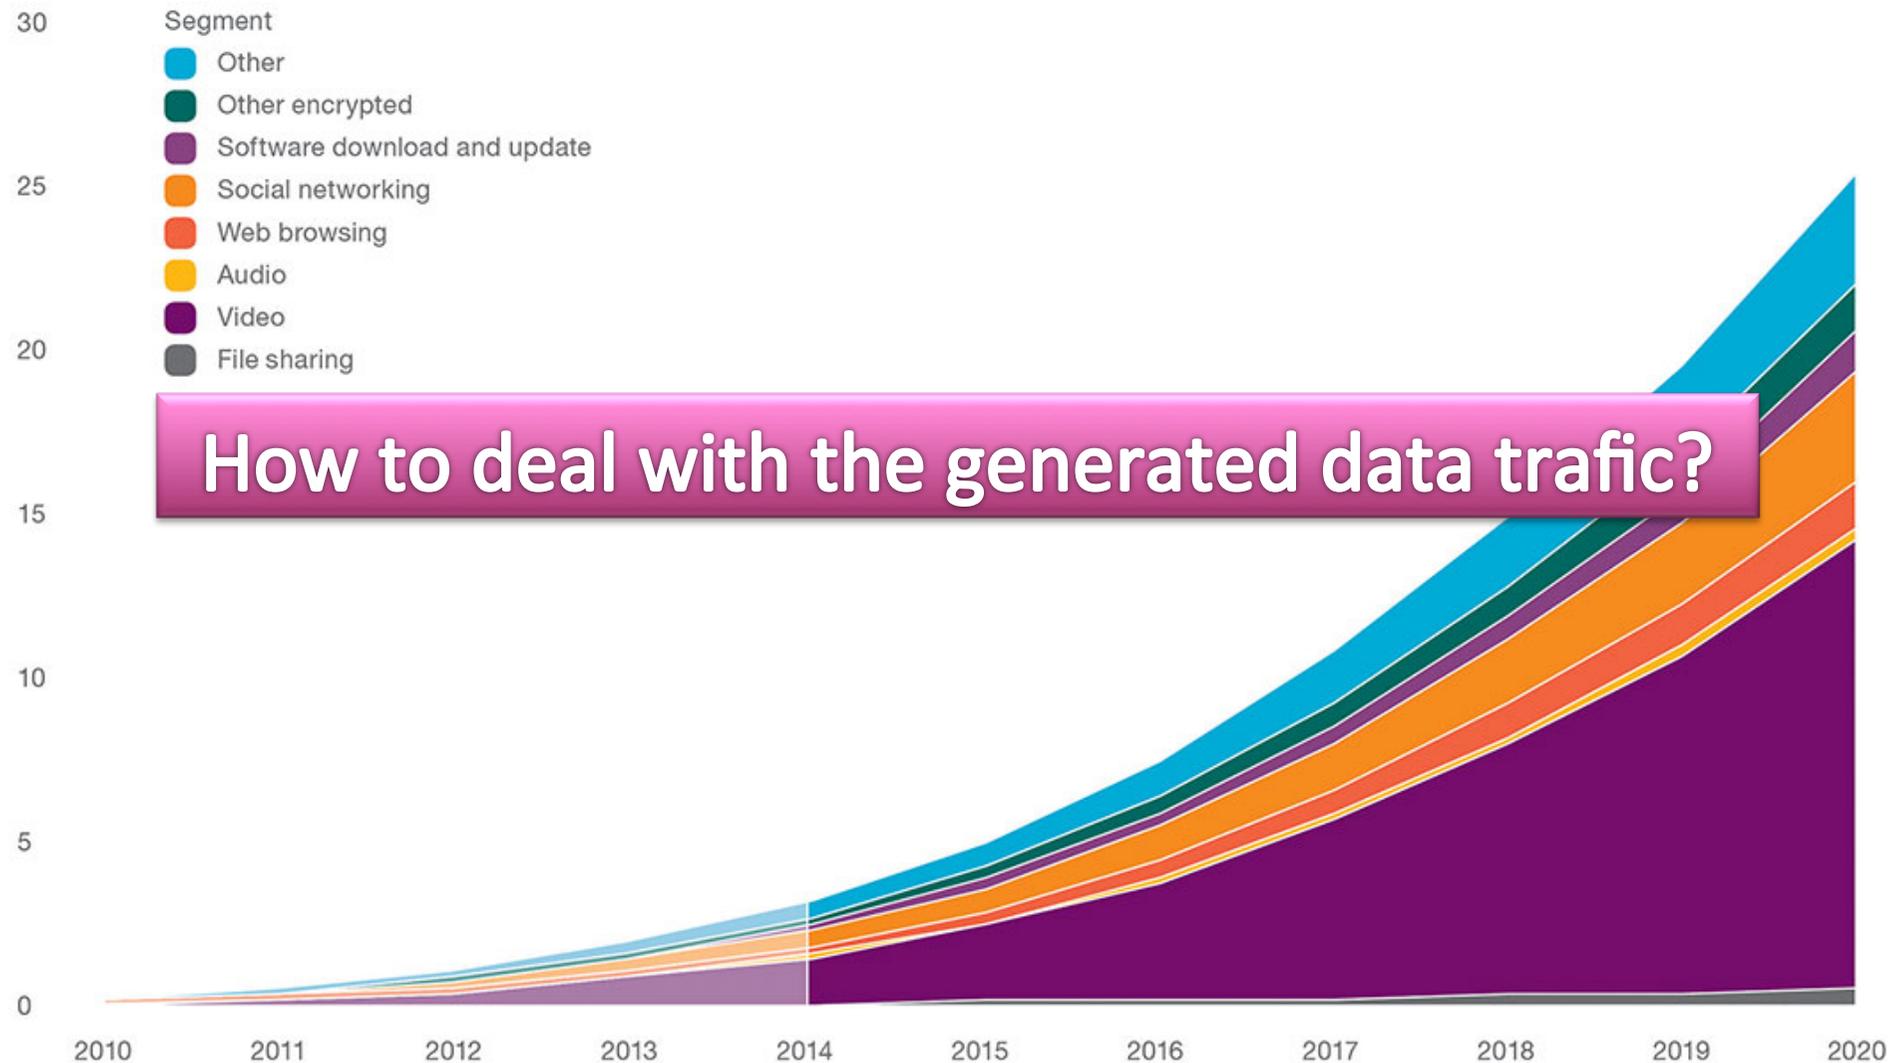
Smartphones, mobile PCs, tablets and mobile routers with a cellular connection



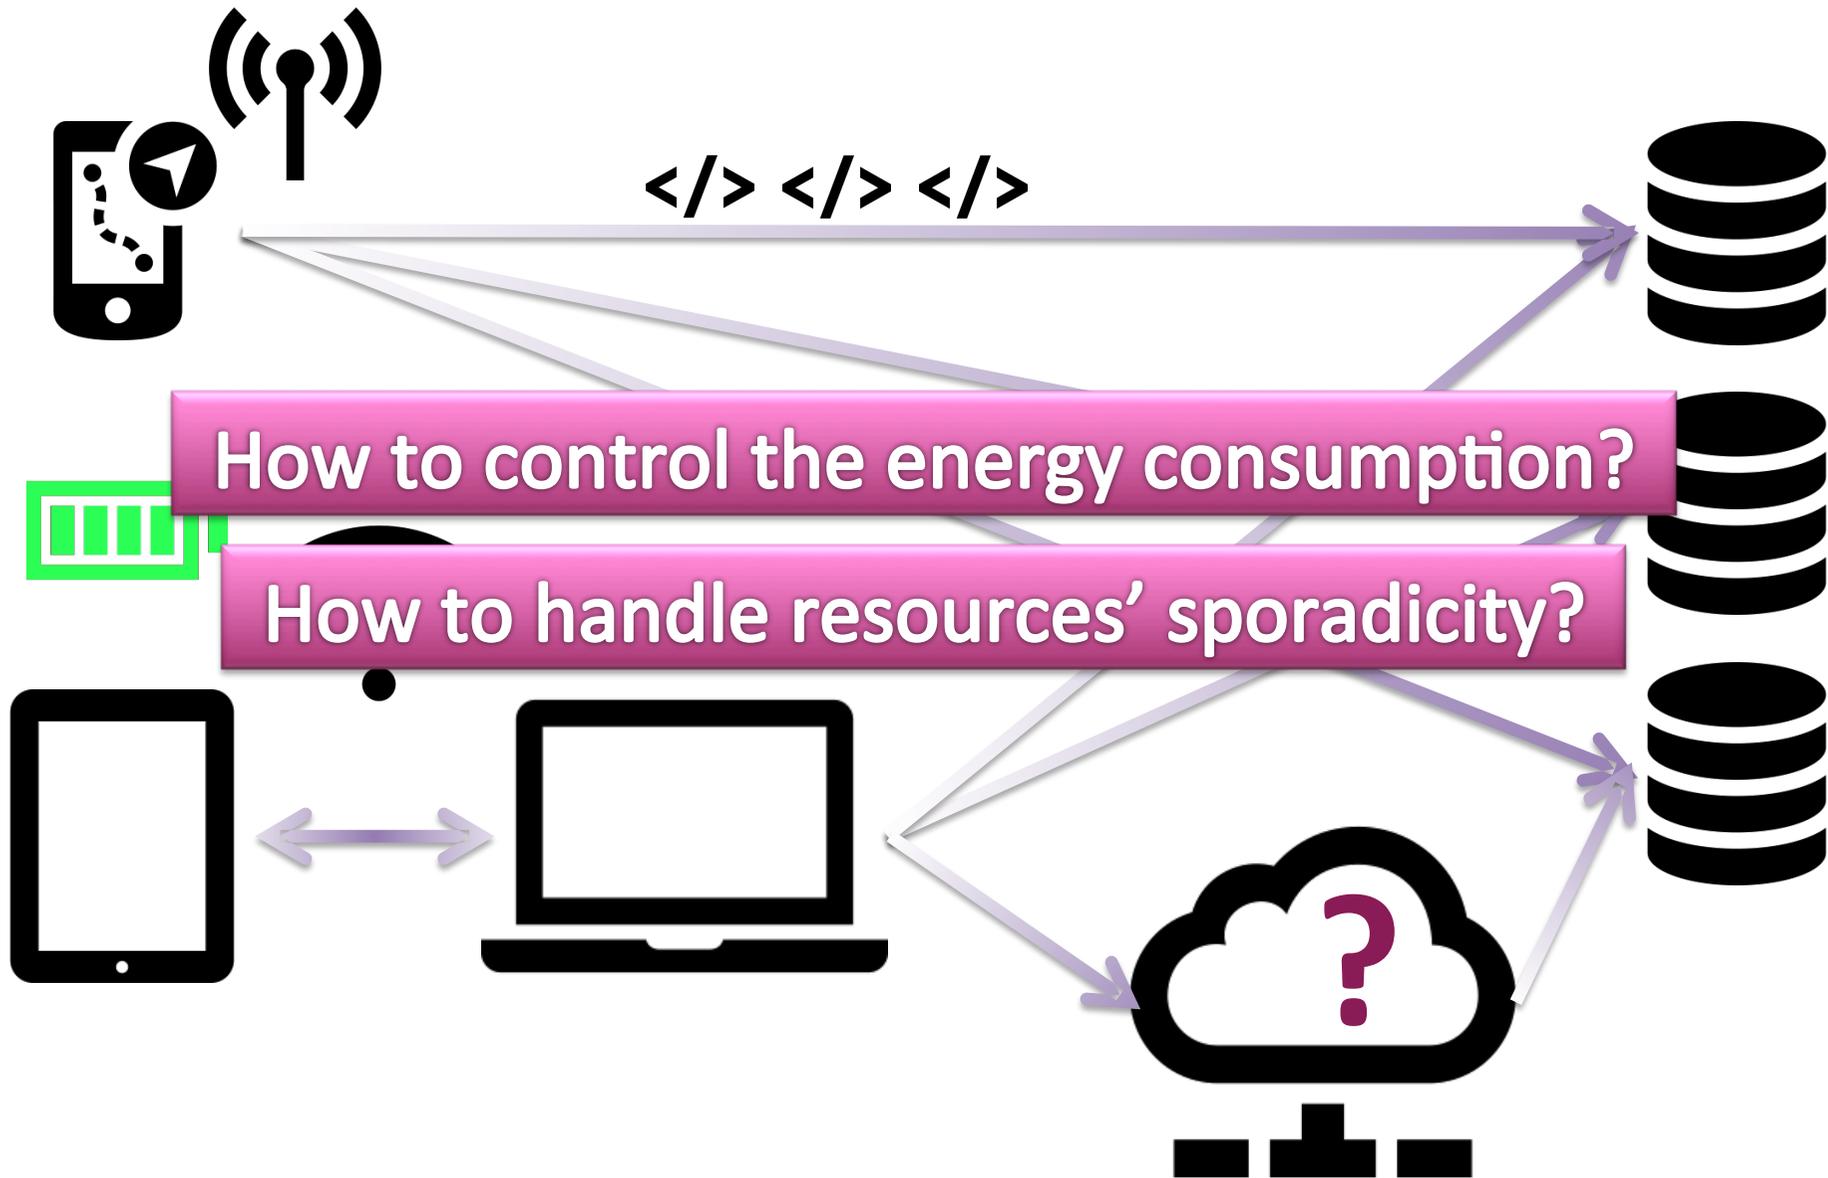
How to program applications for billions of connected devices?

Evolution of application usages

ExaBytes/month

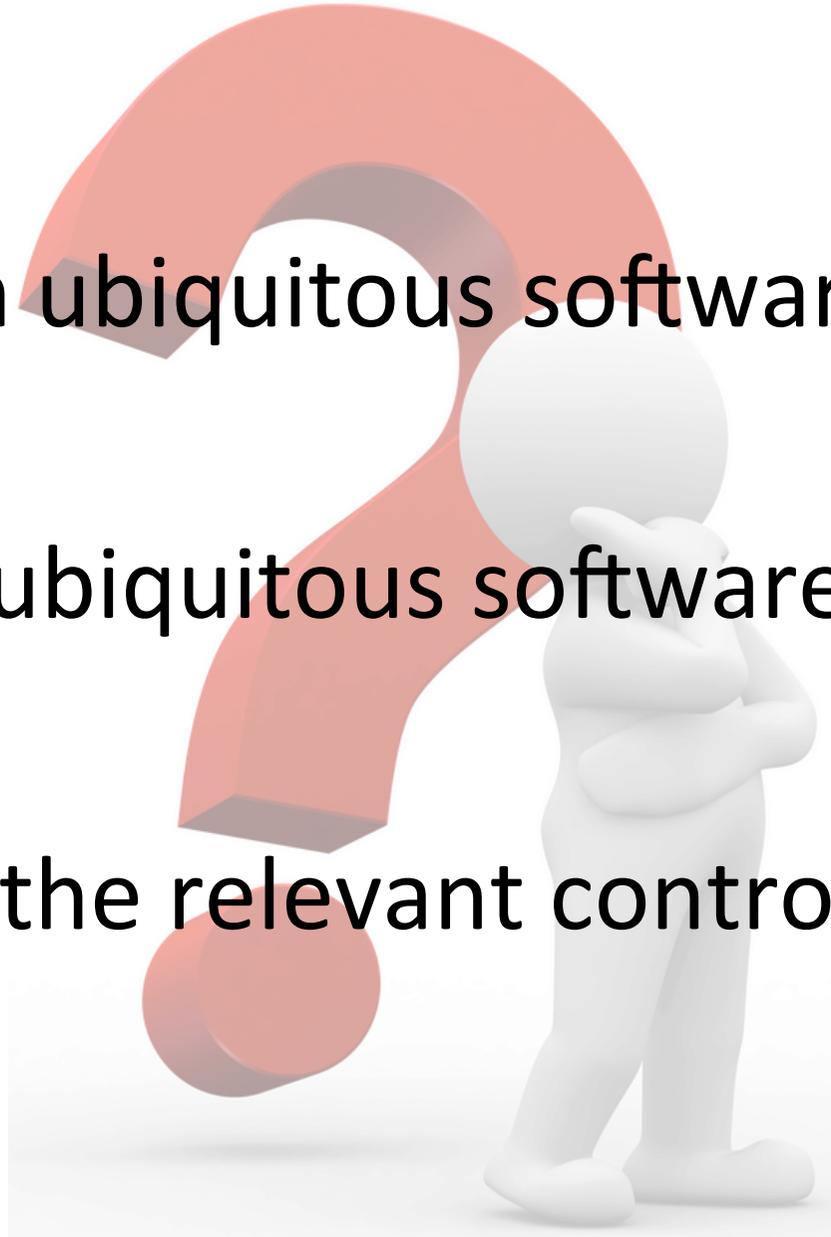


Ubiquitous software systems

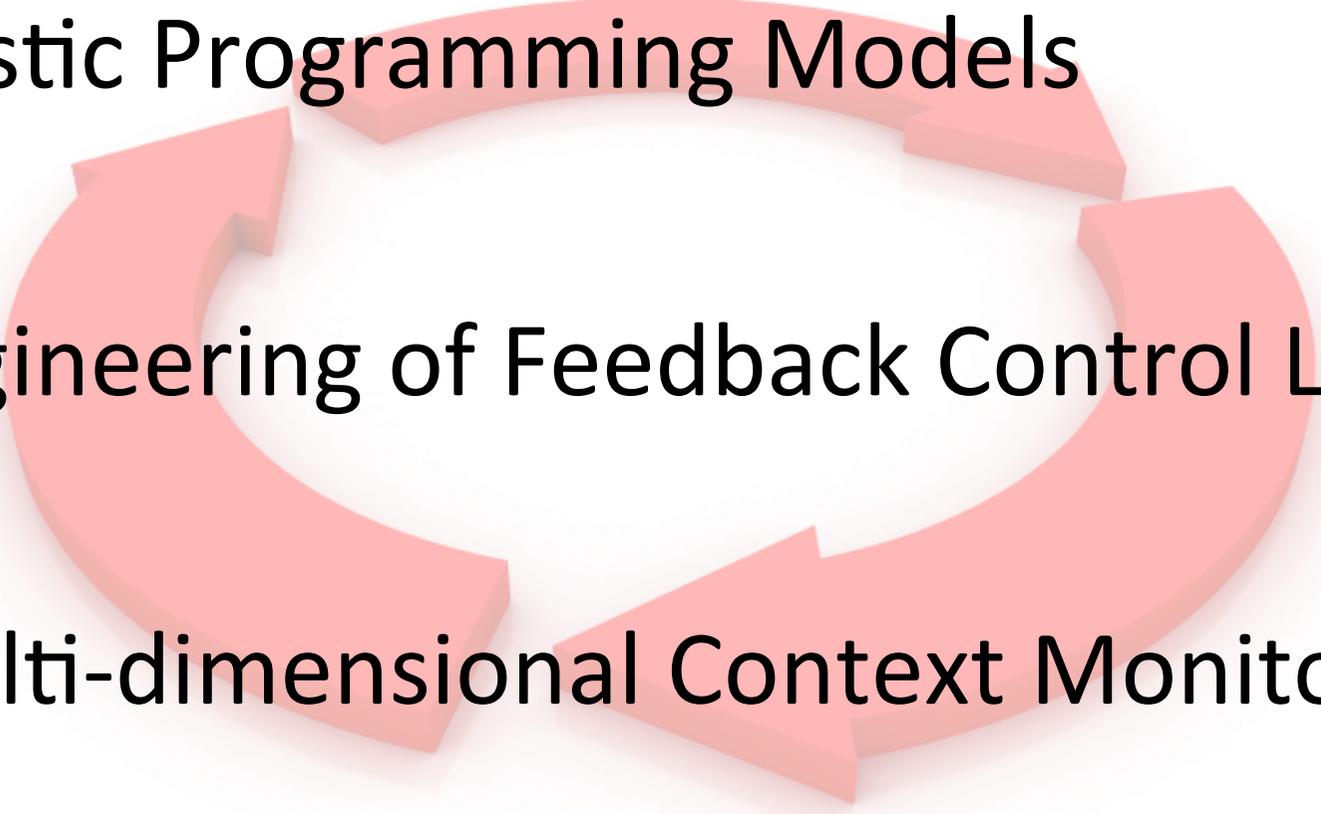


Challenges: How to...

- **Program** ubiquitous software systems?
- **Control** ubiquitous software systems?
- **Identify** the relevant control features?



Our approach: To support...

- 
1. Elastic Programming Models
 2. Engineering of Feedback Control Loops
 3. Multi-dimensional Context Monitoring



2011–2013 macchiato

Enabling opportunistic code offloading

ELASTIC PROGRAMMING MODELS

Ubiquitous software systems require

- **Discovery** of resources and capabilities
 - *Including local and remote discovery*
 - *Cyber foraging* [Satya'01, Goyal'04]
- **Optimisation** of the computing resources
 - *Including energy and network consumptions*
 - *Code offloading* [MAUI'10, Cuckoo'12]
- Leverage of interoperability and **serendipity**
 - *Exploiting unexpected opportunities*

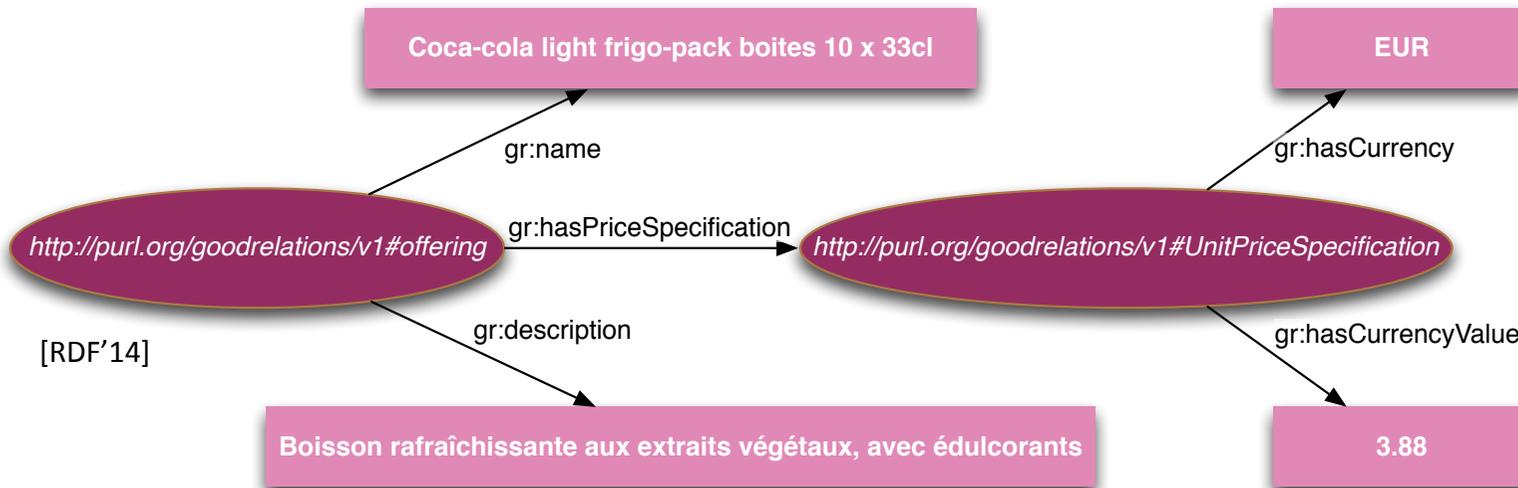
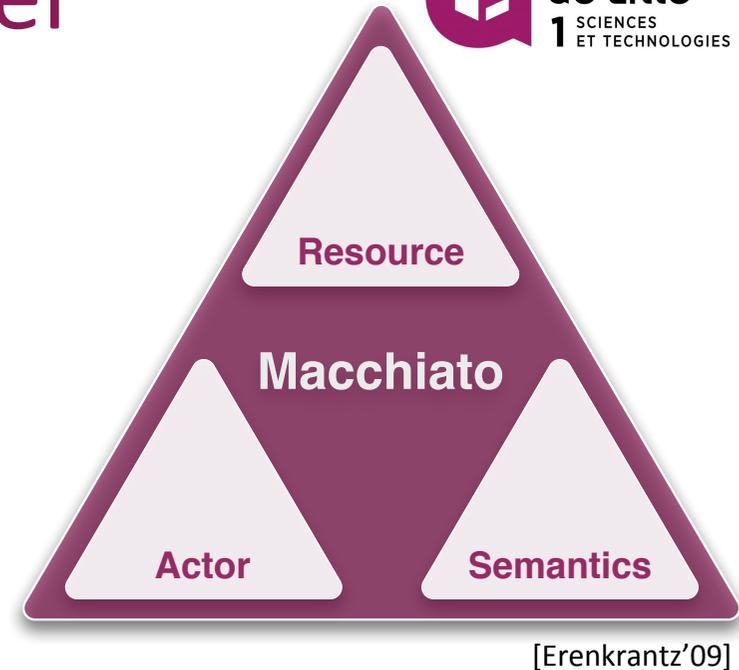
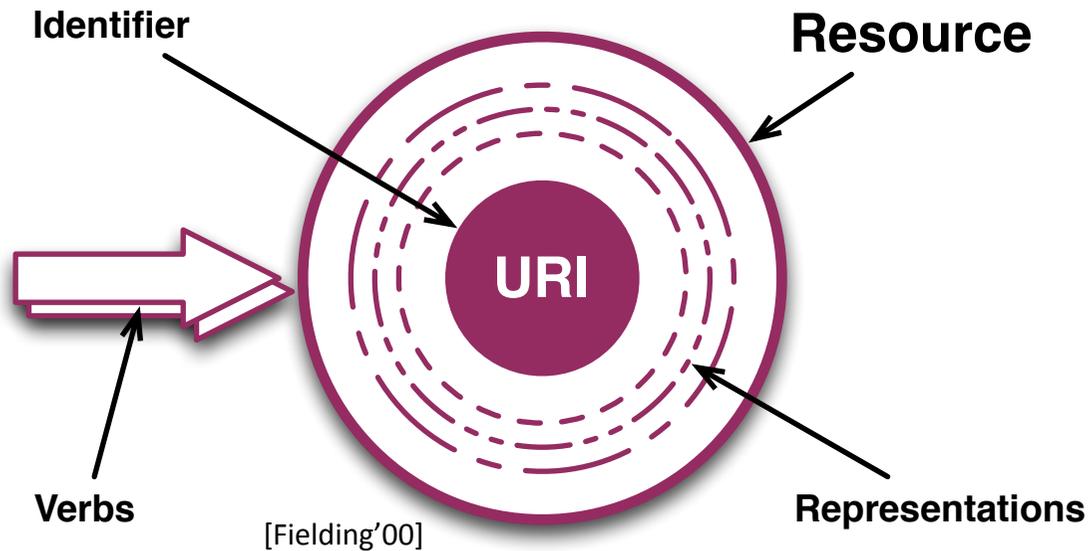
Satya'01. M. Satyanarayanan. *Pervasive Computing: Vision and Challenges*. IEEE Personal Communications. 2001

Goyal'04. S. Goyal et al. *A lightweight secure cyber foraging infrastructure for resource-constrained devices*. WMCSA'04

MAUI'10. E. Cuervo et al. *MAUI: making smartphones last longer with code offload*. MobiSys'10

Cucloo'12. R. Kemp et al. *Cuckoo: a computation offloading framework for smartphones*. MobiCASE'12.

Elastic Programming Model



Fielding'00. R. Fielding. *Architectural Styles and the Design of Network-based Software Architectures*. PhD. UC Irvine. 2000

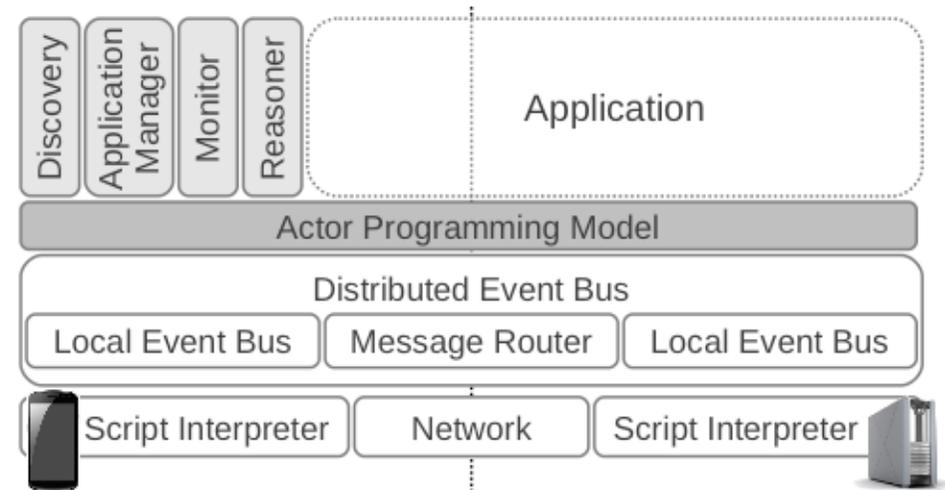
RDF'14. W3C. *Resource Description Framework (RDF)*. Feb. 2014

Erenkrantz'09. J.R. Erenkrantz. *Computational REST: A New Model for Decentralized, Internet-Scale Applications*. PhD. UC Irvine. 2009

Macchiato ubiquitous middleware

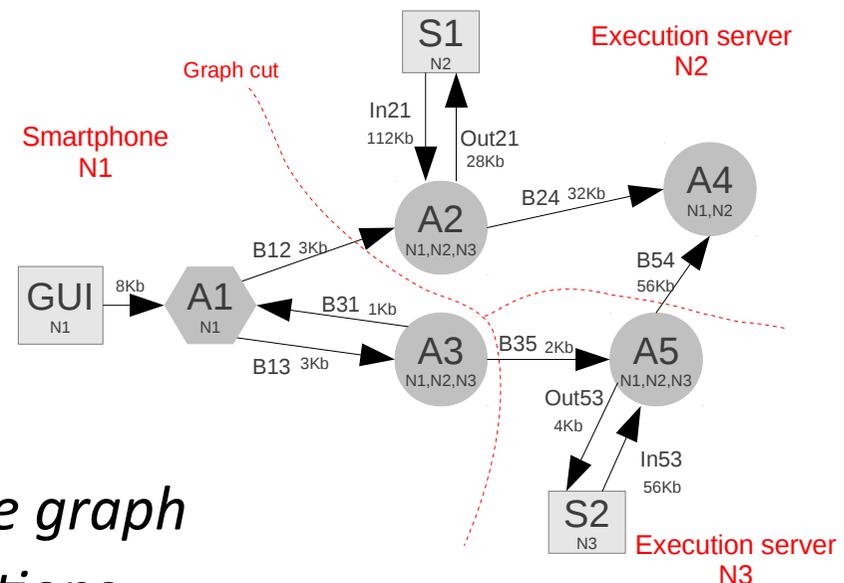
- **Distributed** middleware
 - *Based on JavaScript*
- **Actor-based** model
 - *RDF messages*
- Support for **discovery**
 - *Local discovery (UPnP)*
 - *Location-based discovery*
- Support for **offloading**
 - *Remote actor deployment*
 - *Deployment optimisation*

```
macchiato.create("search-actor").as(  
  function(msg, context) {  
    context.http.get(  
      "http://api.developer.sears.com/v1/products"  
      + "?apikey=" + context.searsapikey  
      + "&store=Sears&searchType=keyword&keyword=" +  
      msg.search, function(result) {  
        context.send(context.parseJSON(result));  
      });  
  });
```



Enabling elasticity in Macchiato

- Actor graph → SAT problem
- Discovery → constraints
- Objective function
 ↳ network exchanges
- Reasoning engine
 - *Triggers upon change in the graph*
 - *Considers sub-optimal solutions*
 - *Queues pending messages*
 - *Keeps a local copy of actors*



Cumulated energy consumption

Cumulated energy consumption



Limitations of the approach

- User-centric process
 - *Lack of coordination*
- Embedded heuristics
 - *Lack of customization*
- Cost of heuristics integration
 - *Lack of extensibility*
- Confidence in heuristics
 - *Lack of feedback*

How to address control as a crosscutting concern?

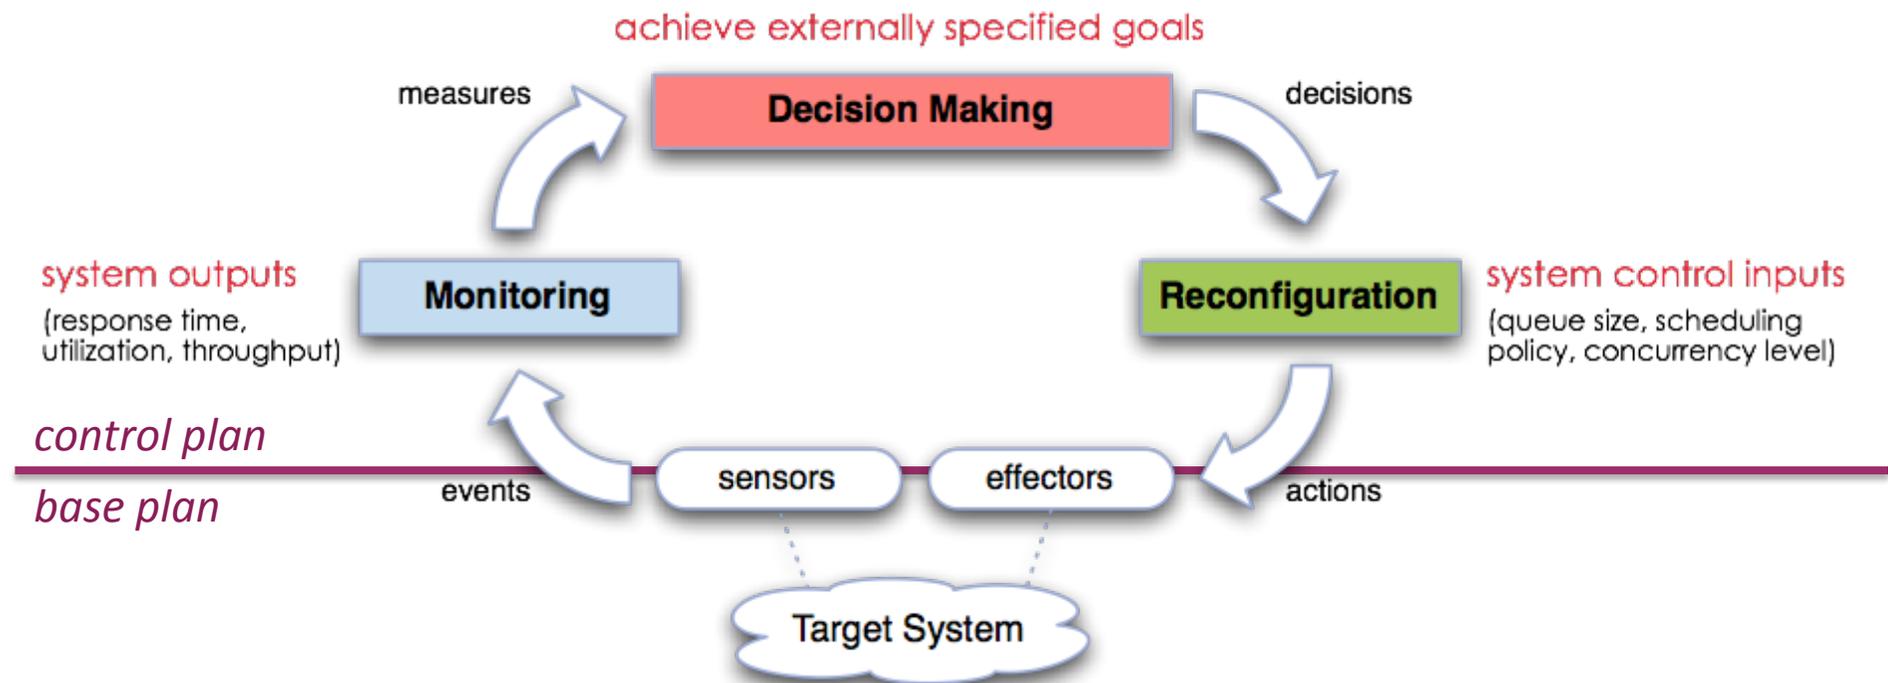


2009–2012

Self-Adaptive very Large disTributed sYstems

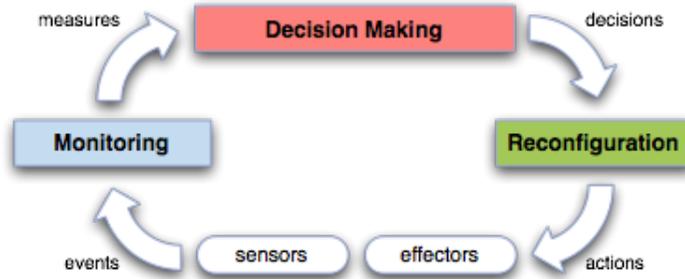
ENGINEERING OF FEEDBACK CONTROL LOOPS

Engineering of Feedback Control Loops



Engineering of Feedback Control Loops

1. Raise the level of abstraction
2. Fine-grained decomposition of FCL elements
3. Explicit interactions
4. Provide reflection capabilities
5. Embed remoting



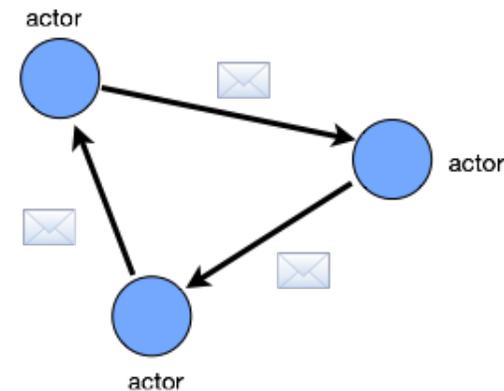
Feedback Control Loop

- Sequence of interconnected processes
- Input \times State \rightarrow Output
- Reactive
- Concurrent
- Dynamic



Domain-Specific Modeling

- Abstraction
- Automation
- Analysis



The Actor Model

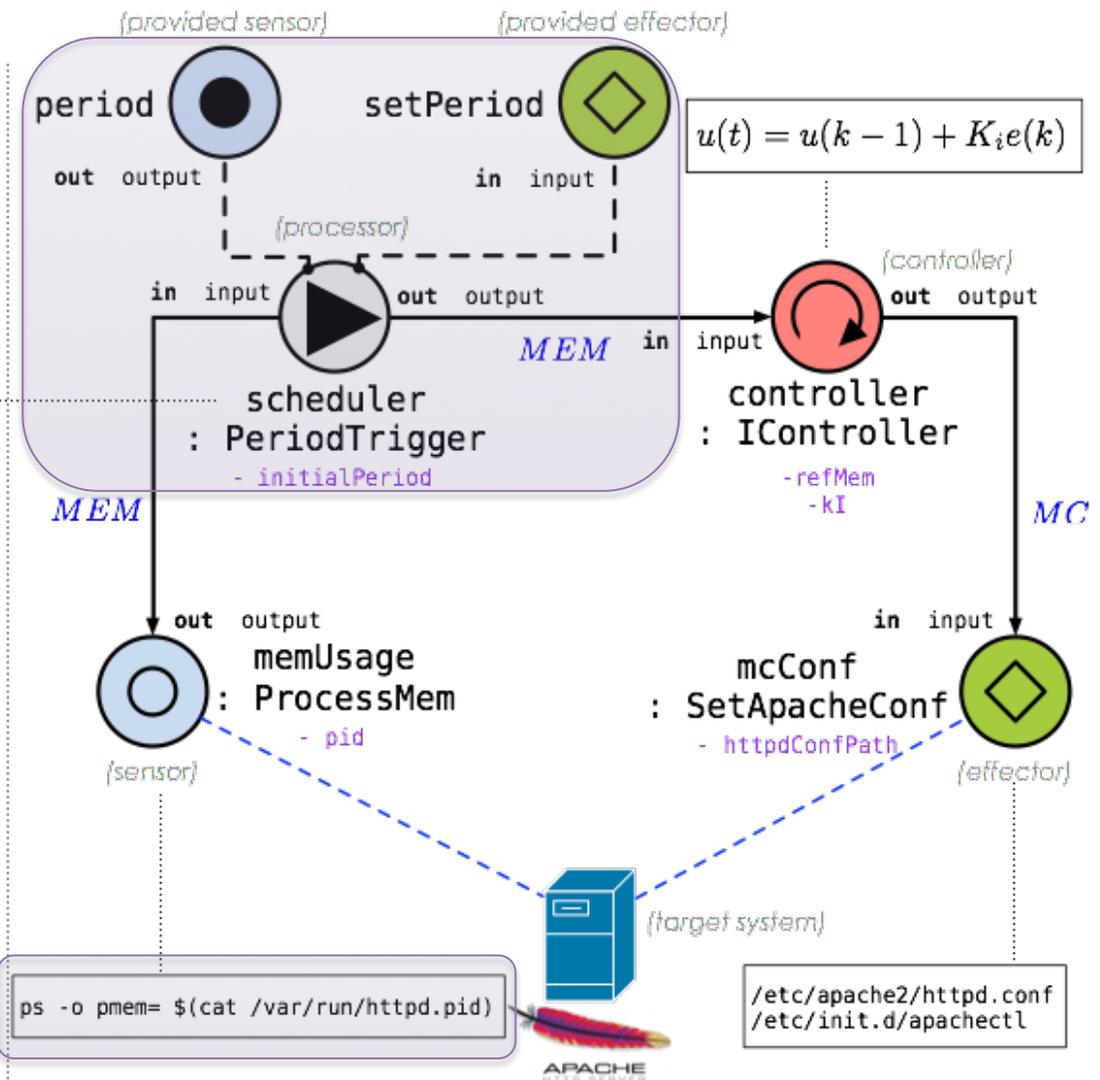
- Message passing actor networks
- Message \times State \rightarrow Message(s)
- Reactive
- Concurrent
- Dynamic
- Scalable
- Remoting through location transparency

FCDL in a nutshell

Adaptive Element

- Actor-like component
- Sensors
- Effectors
- Processors
- Controllers
- Input/output ports & properties
- Active / passive
- Implementation
 - Imperative code (e.g. Java)
 - CEP Rules (e.g. Drools)
 - STM (e.g. bzt/heptagon)
 - Matlab
 - BASH
 - ...
- Interaction contracts

$\alpha = \langle self; \downarrow (\text{input}); \uparrow (\text{output?}) \rangle \parallel \langle \uparrow (\text{setPeriod}); \emptyset; \uparrow (\text{period}) \rangle$



Ubiquitous computing – key results

- **Publications (3 PhD thesis / 8 journals / 15 conferences)**

- Daniel Romero. **Context as a Resource: A Service-Oriented Approach for Context-Awareness.** PhD thesis. Université Lille 1, Sciences et Technologies, July 2011.
- Eli Gjørven. **Enabling Self-Adaptation by Applying a Technology Agnostic Middleware with Support for Integration.** PhD thesis. University of Oslo, Dec. 2011.
- Russel Nzekwa. **Building Manageable Autonomic Control Loops for Large Scale Systems.** PhD thesis. Université Lille 1, Sciences et Technologies, July 2013.
- R. Rouvoy et al. *Dynamic planning and weaving of dependability concerns for self-adaptive ubiquitous services.* In: **SAC'09.**
- A. Taherkordi et al. *A Generic Component-based Approach for Programming, Composing and Tuning Sensor Software.* In: **The Computer Journal 54.2** (2011)
- N. Petitprez et al. *Connecting your Mobile Shopping Cart to the Internet-of-Things.* In: **DAIS'12**
- R. Rouvoy and P. Merle. *Rapid Prototyping of Domain-Specific Architecture Languages.* In: **CBSE'12**
- A. Taherkordi et al. *Optimizing Sensor Network Reprogramming via In-situ Reconfigurable Components.* In: **ACM Transactions on Sensor Networks 9.2** (2013)
- F. Krikava et al. *Integrating Adaptation Mechanisms Using Control Theory Centric Architecture Models: A Case Study.* In: **ICAC'14.**

- **Collaborations**

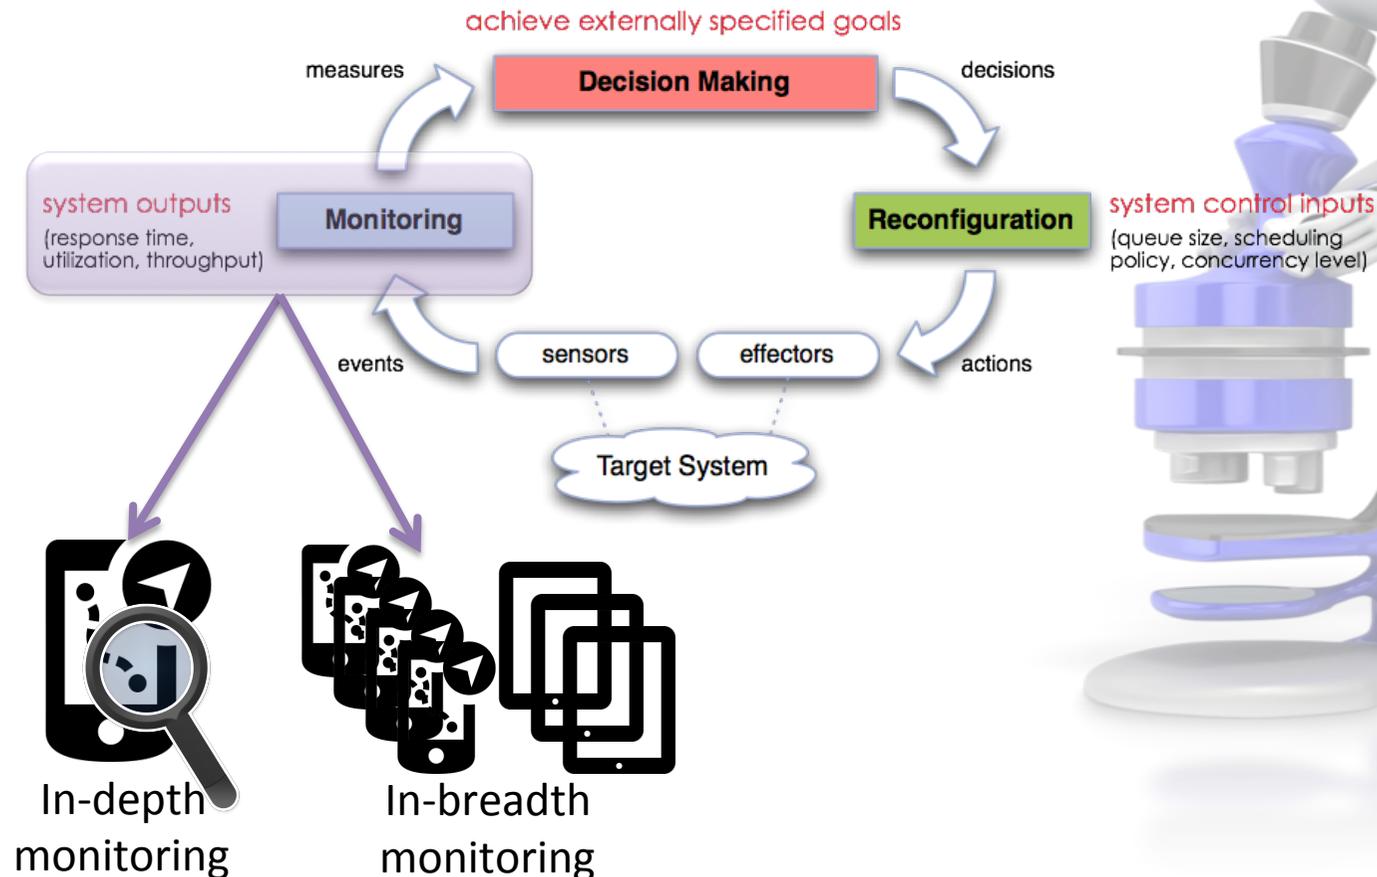
- *University of Nice (ANR SALTY, P. Collet & F. Krikava)*
- *University of Oslo (EGIDE PHC Aurora & SeaS associate team)*

- **Grants**

- ANR ARPEGE SALTY FUI Cappucino
- FUI Macchiato PIA Datalyse

A Focus on context monitoring

- A **critical task** of autonomous software systems
- Required for **system identification** & adaptation
- Produce **meaningful data** for free





2010–2014

A focus on power consumption

IN-DEPTH CONTEXT MONITORING

Context

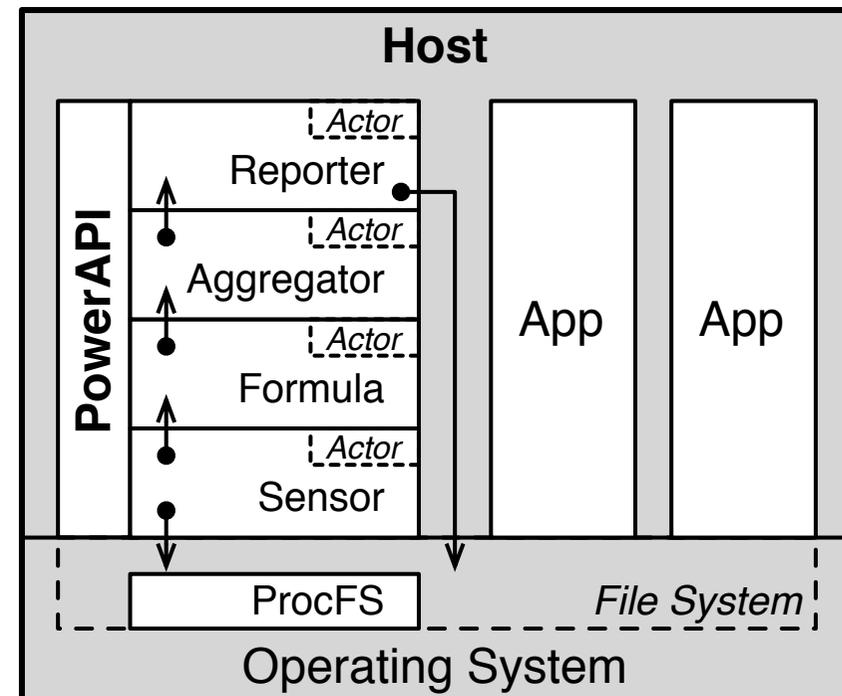
- Energy consumption of ICT (France) [TIC'09]
 - *13.5% in 2009*
 - *40% in 2019 (forecast)*
- Meanwhile
 - ↗ *IT devices & datacenters*
 - ↗ *Idle power consumption*
- Hardware optimizations
 - *Dynamic Voltage/Frequency Scaling (DVFS)*
 - *Hyper-threading*
 - *Turboboost*
 - *Core power states (C-States)*

How to identify energy hotspot in applications?

PowerAPI middleware toolkit

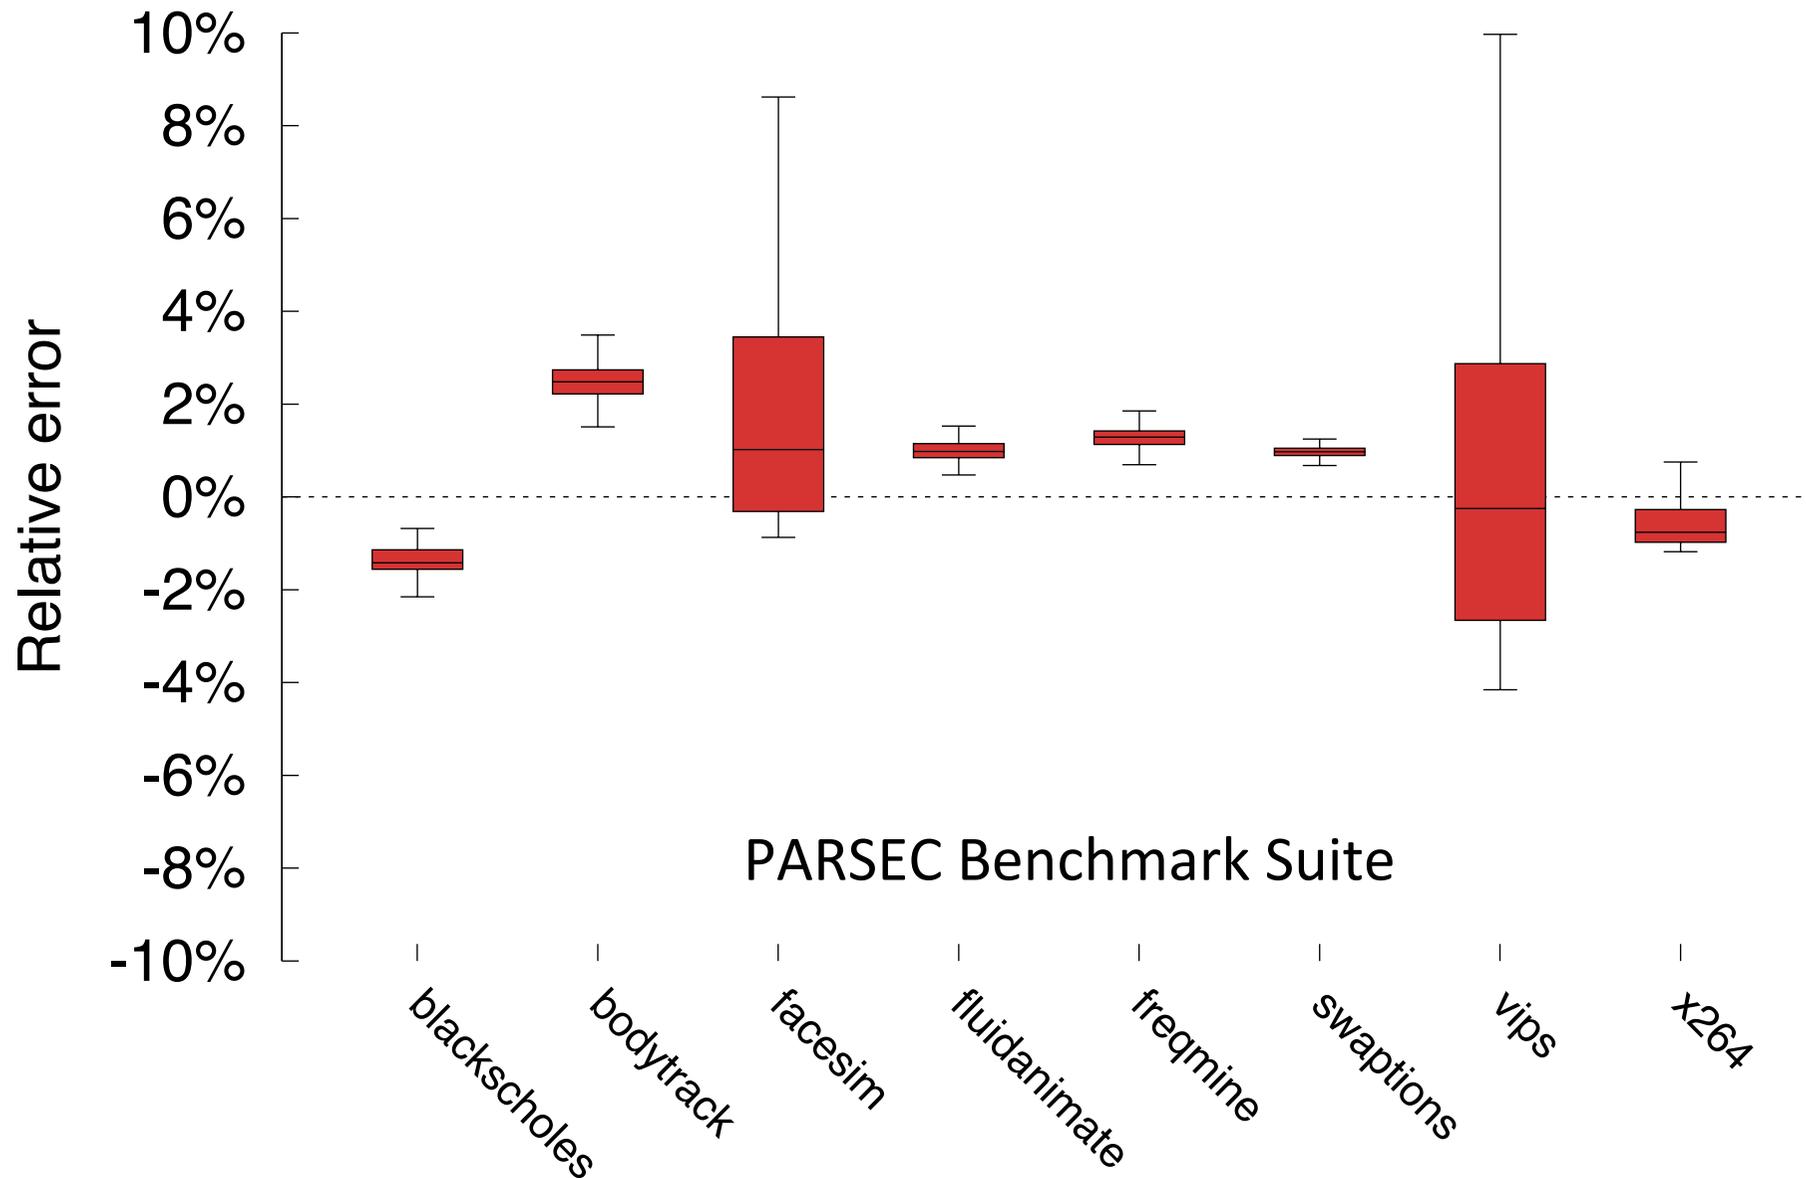
Software-defined power meters

- Internal sensors
 - *ProcFS, libpfm4*
- External sensors
 - *PowerSpy, RAPL*
- Formulas
 - *CMOS, DVFS, regression*
- Reporters
 - *GUI, Network, FUSE*

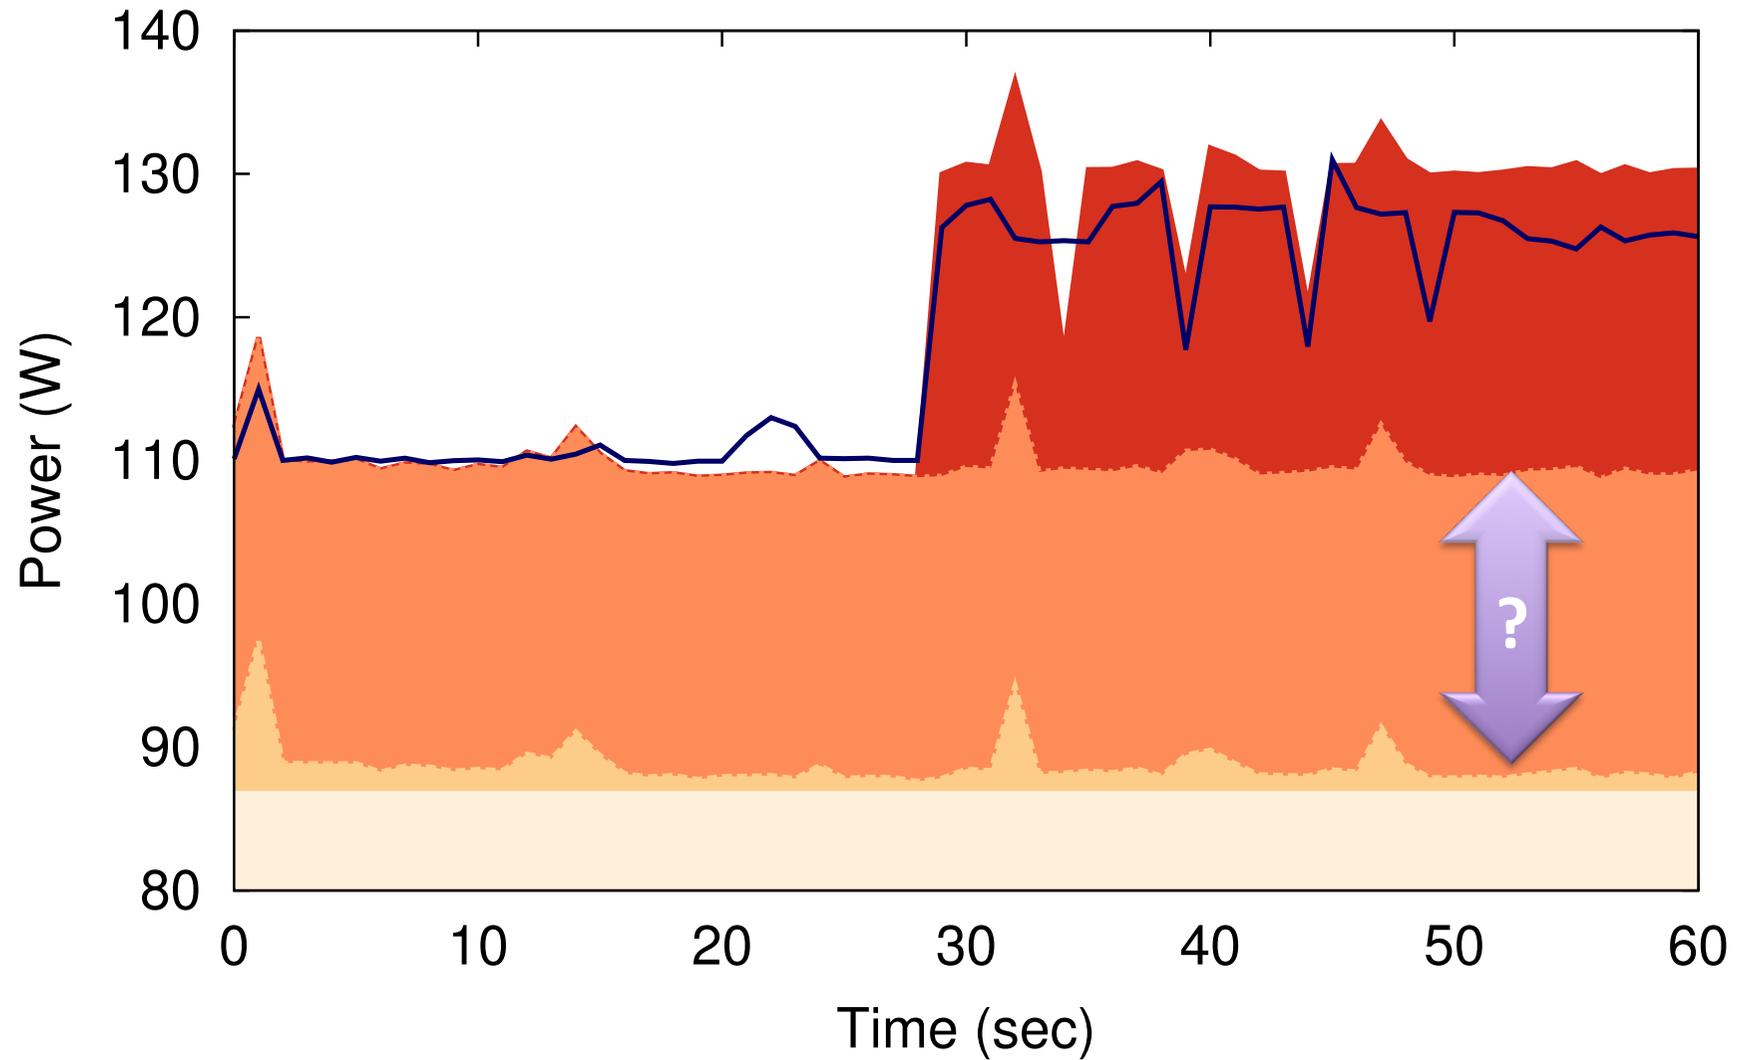


$$Power_{software}^d = CPU_{software}^{utilization} * Power_{CPU}^d$$

PowerAPI mean estimation error



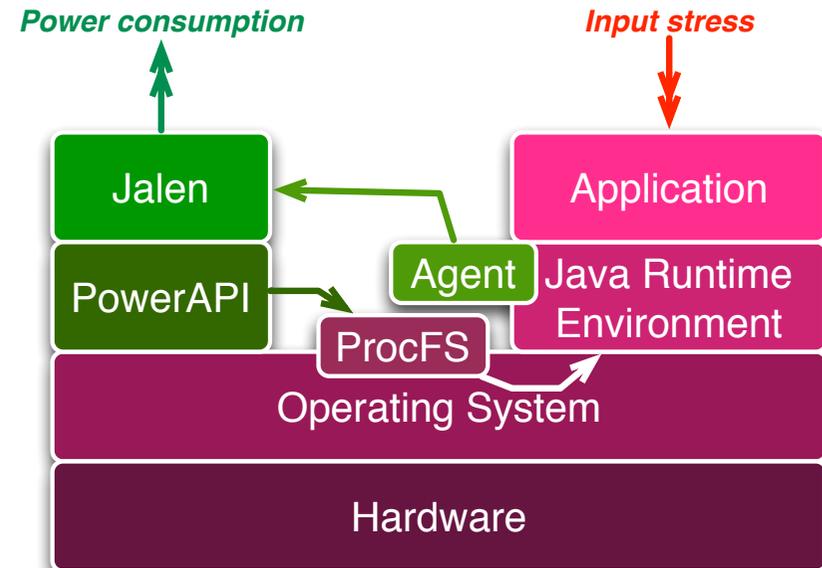
PowerAPI power estimation accuracy



x264  PowerVM  PowerSpy 
freqmine  Idle power 

Code-level power monitoring

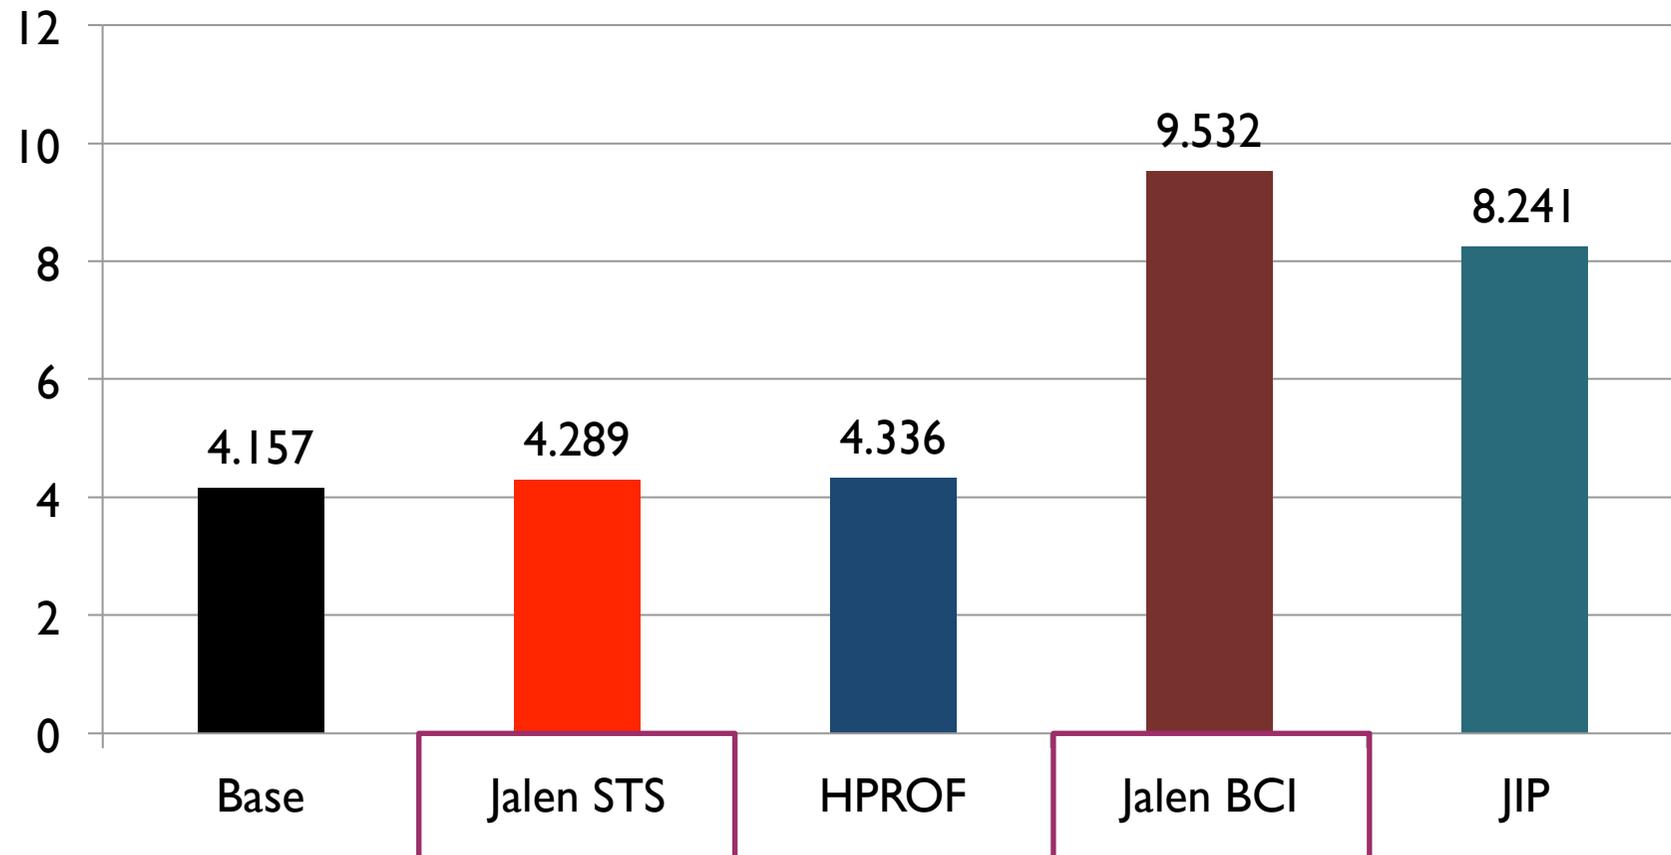
- **Jalen = PowerAPI++**
 - Exploits process estimations
 - Extracts code-level metrics
 - Monitors utilization per thread (CPU + HDD)
- 2 implementations
 - Bytecode instrumentation
 - Statistical sampling



$$\text{Power}_{\text{thread}}^{\text{CPU}} = \frac{\text{Power}_{\text{software}}^{\text{CPU}} * \text{Utilization}_{\text{thread}}^{\text{CPU}}}{\text{Duration}_{\text{cycle}}}$$

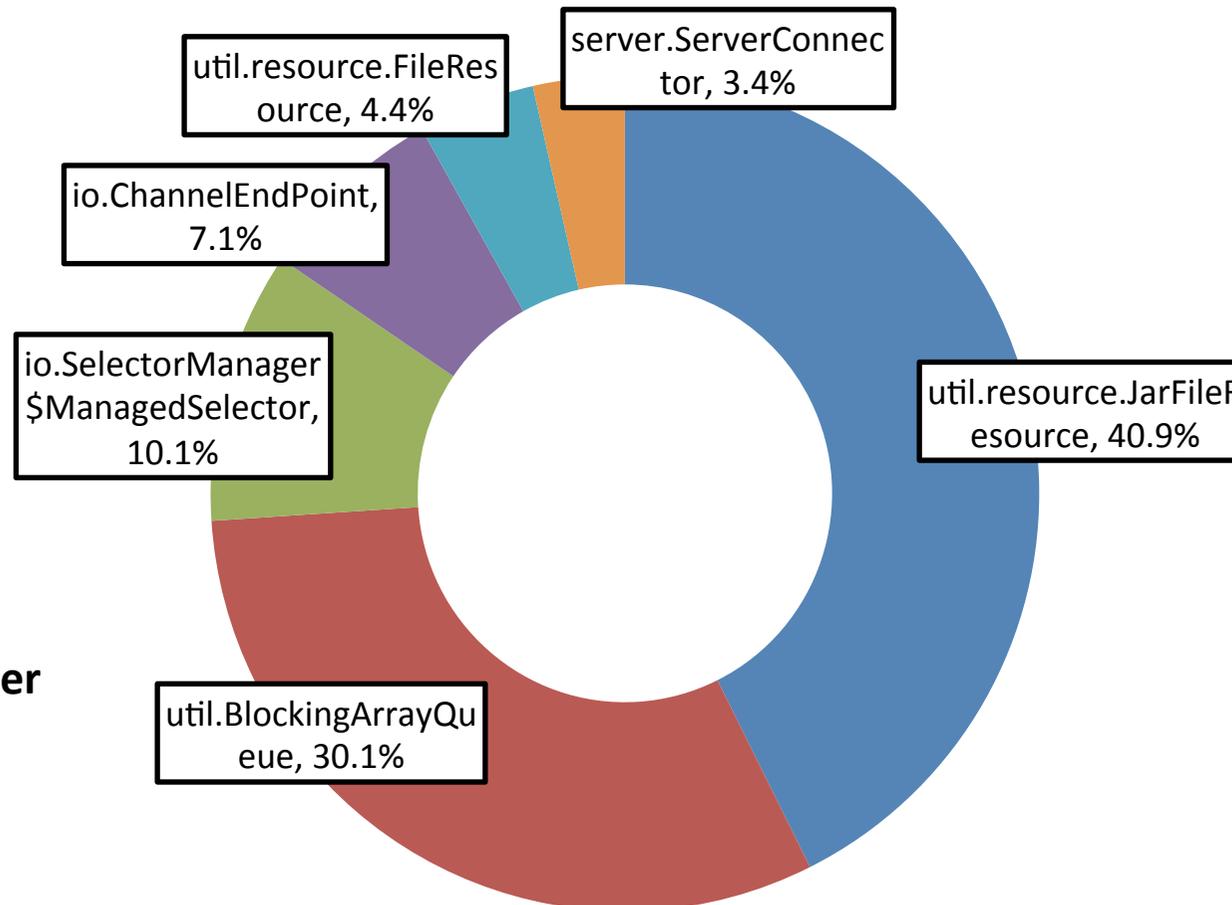
Jalen instrumentation overhead

Individual Tomcat requests (ms), sampling rate of 10 ms



Energy dissipation into classes

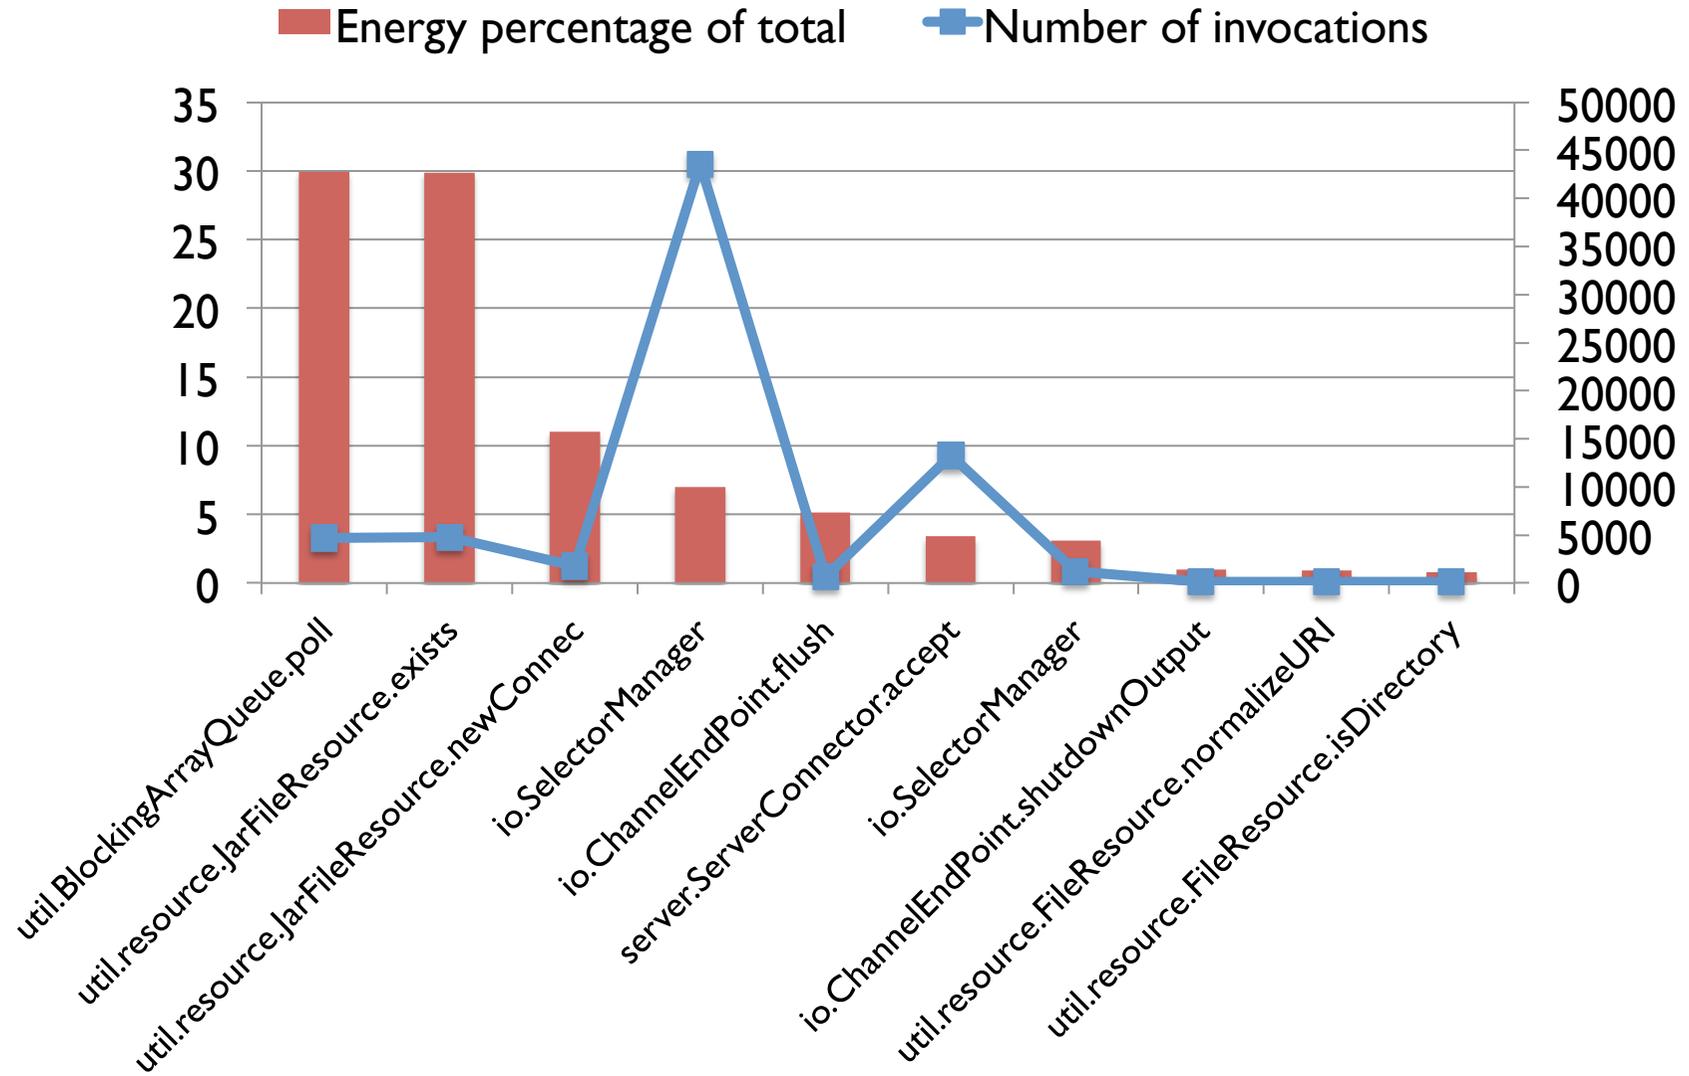
Energy percentage of total



Jetty web server
> 105,000 LoC

Stress
25 users
100,000 requests

Energy dissipation into methods



In-depth monitoring - key results

- **Publications (1 PhD thesis, 3 journals, 2 conferences)**
 - A. Nouredine. ***Towards a Better Understanding of the Energy Consumption of Software Systems***. PhD thesis. Université Lille 1. July 2014
 - A. Nouredine et al. *Monitoring Energy Hotspots in Software*. In: ***Journal of Automated Software Engineering*** (2014).
 - A. Nouredine et al. *A review of energy measurement approaches*. In: ***ACM SIGOPS Operating Systems Review*** 47.3 (Dec. 2013).
 - A. Nouredine et al. *A Review of Middleware Approaches for Energy Management in Distributed Environments*. In: ***Software: Practice and Experience*** 43.9 (Sept. 2013).
 - A. Nouredine et al. *Runtime Monitoring of Software Energy Hotspots*. In: ***ASE'12***.
- **Collaborations**
 - *Power consumption estimation in virtualised environments (Univ. of Neuchâtel)*
 - *Reducing of the power consumption of a HPC cluster (L2EP, Univ. of Lille 1)*
- **Grants**
 - *FUI EconHome*
 - *Inria ADT 2013 (1 engineer)*
- **Transfer (<http://powerapi.org> - AGPL)**
 - *Web Energy Archive (funded by ADEME)*
 - *KaliTerre (dynamic code analysis)*



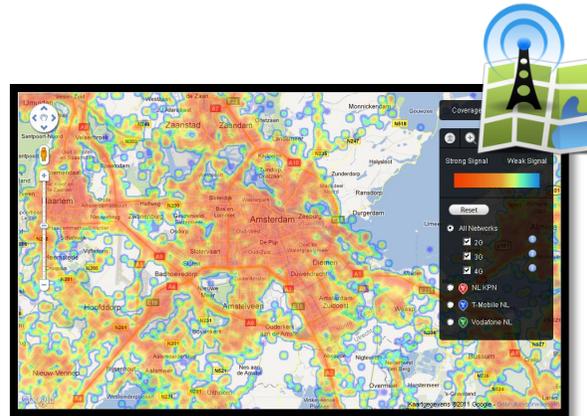
2010–2014

A focus on mobile crowd-sensing

IN-BREADTH CONTEXT MONITORING

Mobile crowd-sensing

- «*category of applications that we term mobile crowdsensing, where **individuals** with sensing and computing devices collectively **share data** and extract information **to measure and map phenomena of common interest***» [Ganti'11]



How to federate crowd-sensing tasks at large?

Research challenges

- **Flexibility** [Funf, PRISM]
 - *Supporting participative and opportunistic tasks*
- **Scalability** [MEDUSA]
 - *Supporting both individual and crowd tasks*
- **Privacy** [Anonymsense]
 - *Preserving sensitive data for the participants*
- **Energy** [Pogo]
 - *Minimizing the power consumption of tasks*
- **Incentives** [Anonymsense]
 - *Engaging participants to contribute along time*

Anonymsense. Shin et al. Anonymsense : A system for anonymous opportunistic sensing, *Pervasive and Mobile Computing*, 2011

PRISM. Das et al. PRISM : platform for remote sensing using smart-phones, *MobiSys'10*

Funf. Aharony et al. *Social fMRI: Investigating and shaping social mechanisms in the real world*, *Pervasive and Mobile Computing*, 2011

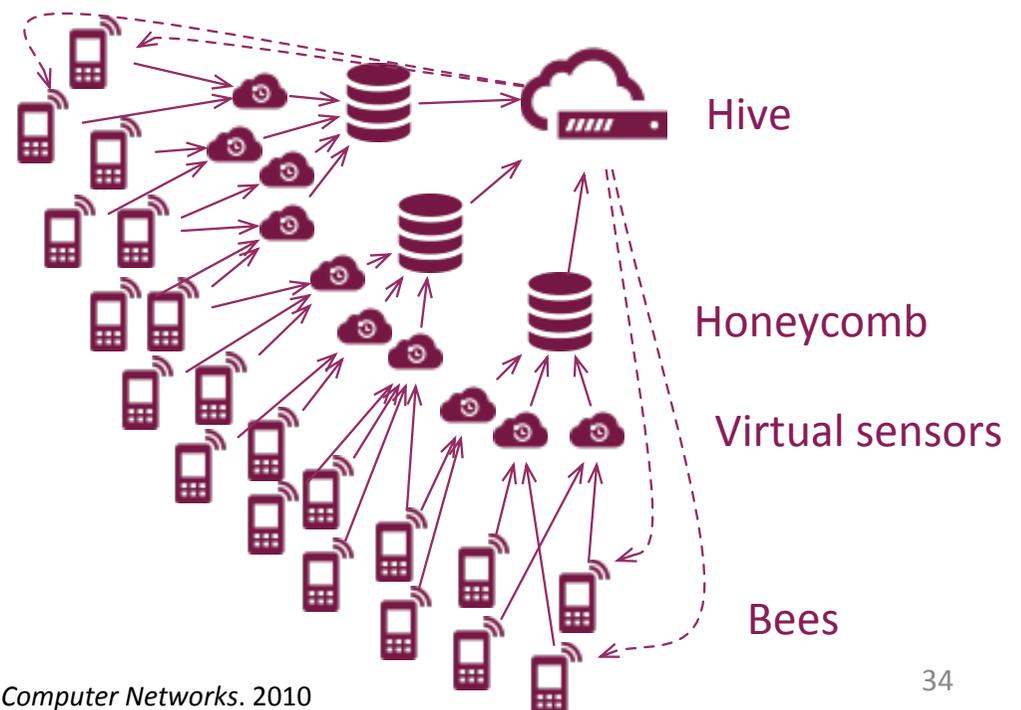
Pogo. Brouwers et al. *Pogo, a middleware for mobile phone sensing*. *Middleware'12*

Medusa. Ra et al. *Medusa: A programming framework for crowd-sensing applications*. *MobiSys'12*

APISENSE[®]: Mobile crowd-sensing

- Crowd-sensing language
 - *JavaScript extension*
 - *Injecting facades*
- Mobile infrastructure
 - *Bee SDK*
 - *Bee App*
- SaaS infrastructure
 - *Hive service*
 - *Honeycomb services*
 - *Virtual sensors* [Chowdhury10]

```
$location.onLocationChange(function(event)
{
    $trace.add({
        lat : event.latitude,
        lng : event.longitude,
        signal : $telephony.signalStrength()
    });
});
```



Orchestrating crowd-sensing tasks

sense

```
sense(function( ) { ... } );
```

recruit

```
accept(function( ) {  
    if ($network.connectionType() == 'Mobile')  
        return { battery : $battery.level() };  
});  
ranking(function(users) {  
    return users.sort('battery');  
});
```

coverage

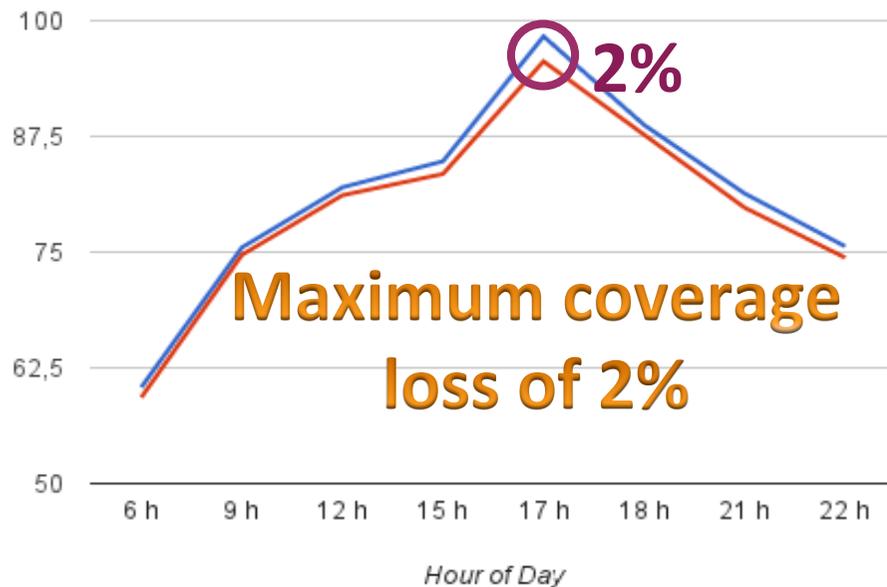
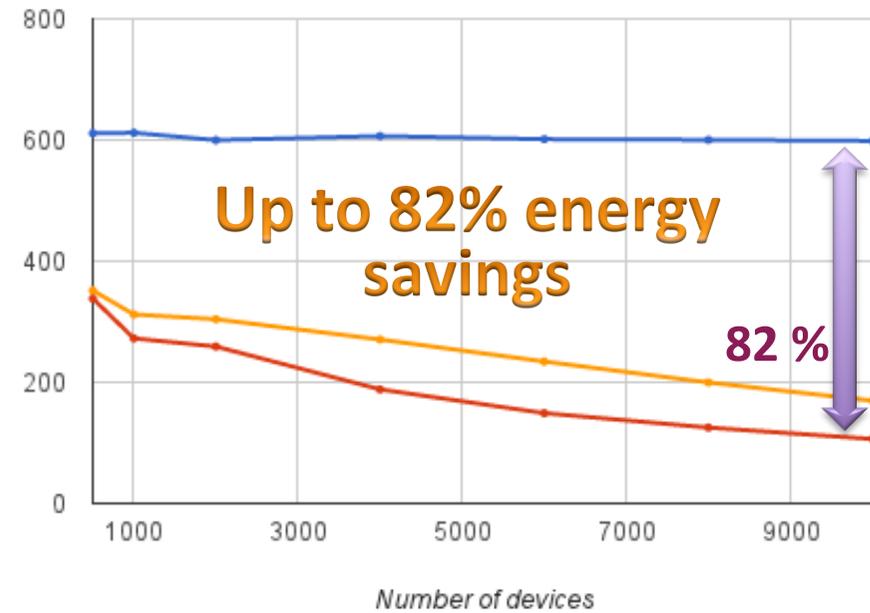
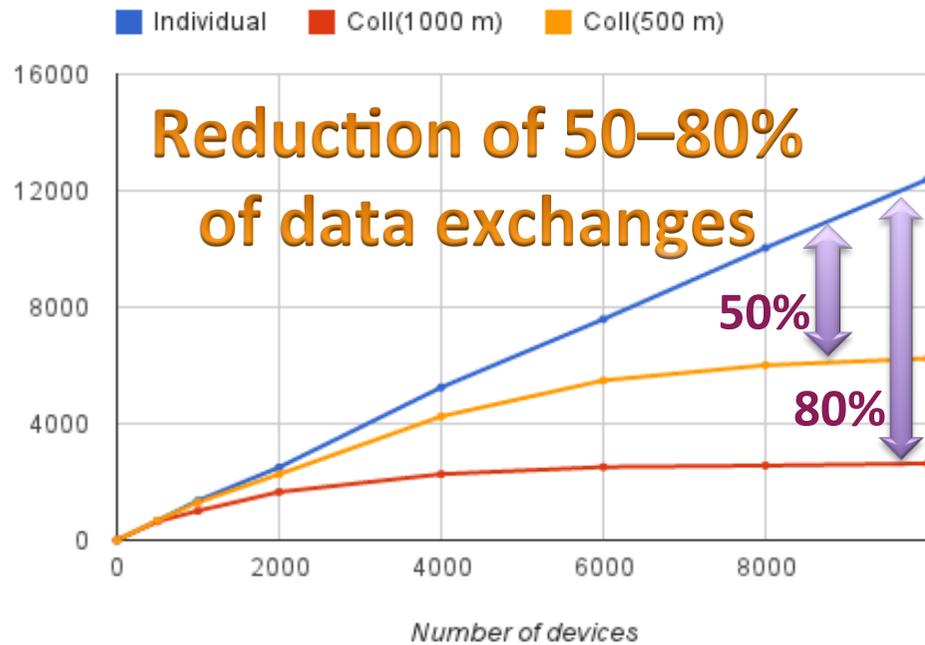
```
geoCoverage(  
    [[50.614291,3.13282],[50.604159,3.15239]],  
    '500 m');  
timeCoverage("30 min","1 H");  
duplicate(10);
```

Evaluation of APISENSE®

- Objectives
 1. *Volume of exchanged data*
 2. *Energy consumption*
 3. *Data coverage*
- Real mobility traces
 - *10 000 cabs from Beijing*
- Comparing 3 strategies
 - a. Individual: no coordination (witness)*
 - b. Coll(500): coverage objective of 500×500 meters*
 - c. Coll(1000): coverage objective of 1000×1000 meters*



Evaluation of APISENSE®

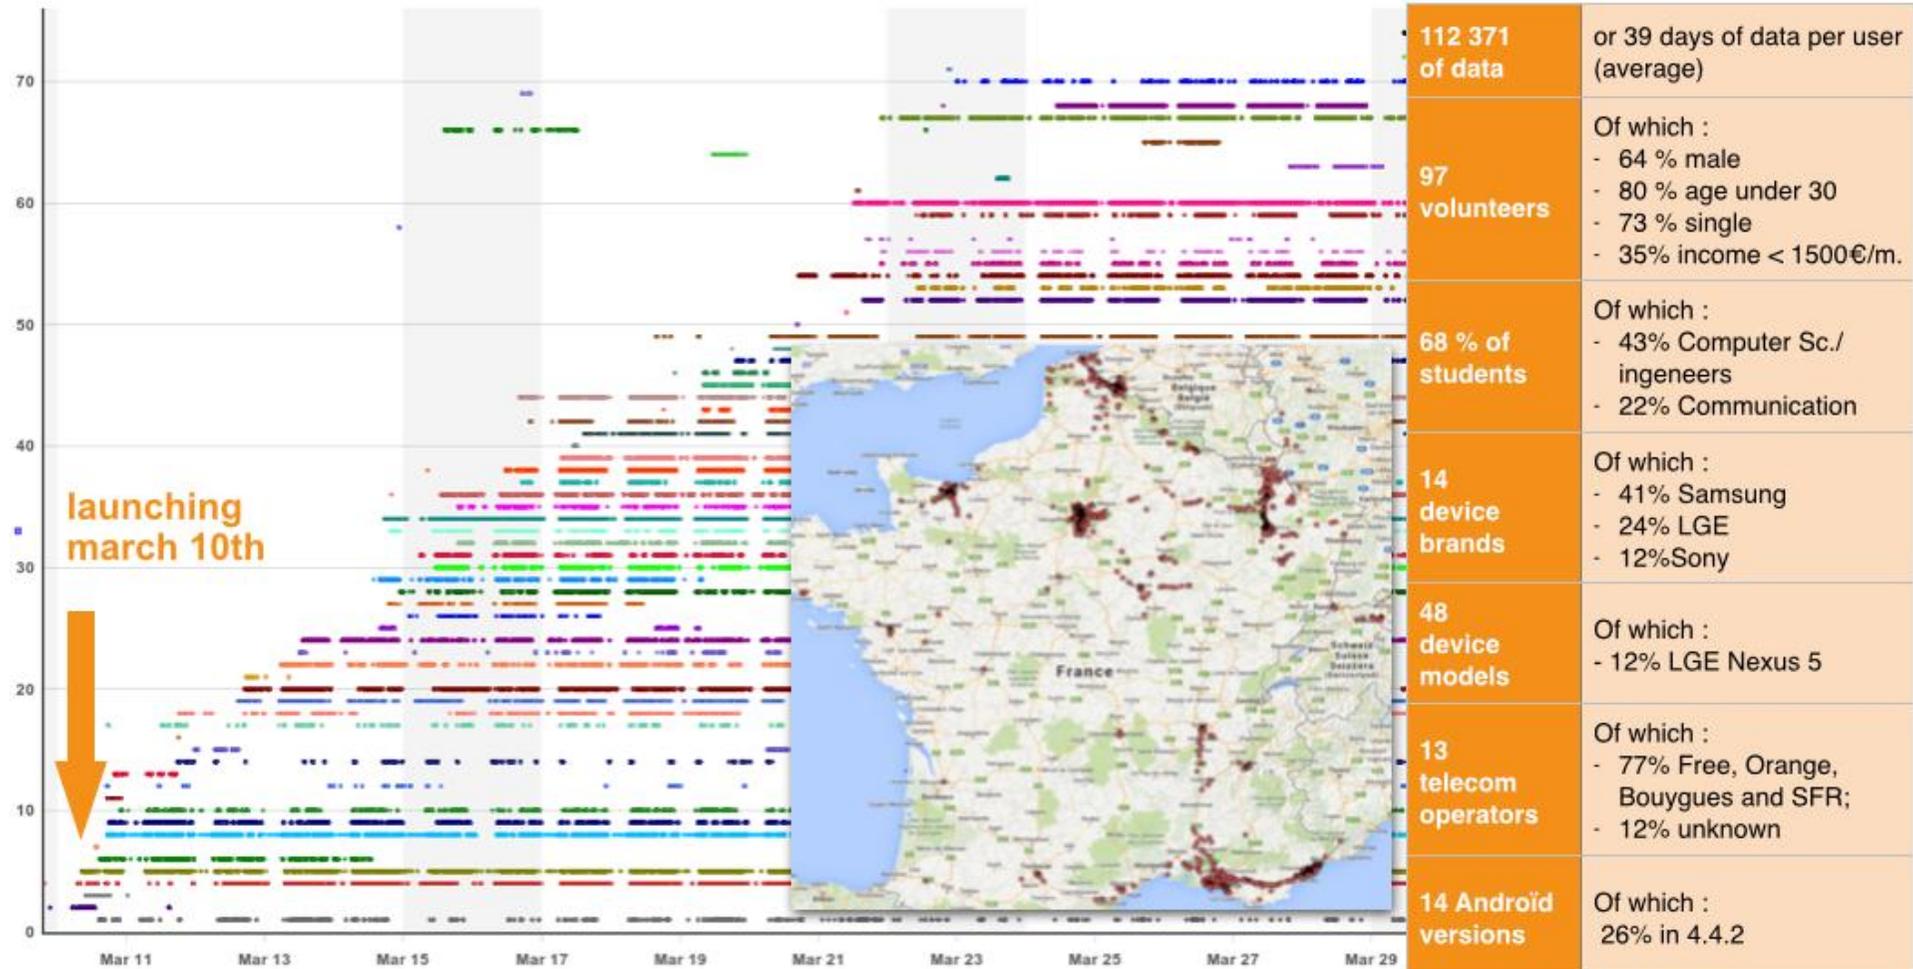


	W/B-Scanner opportunist	Citizen journalist participative
APISENSE®	4	9
Anonymsense	5	N/A
Pogo	4	N/A
MyExperience	N/A	27
Medusa	N/A	45
PRISM	??	330

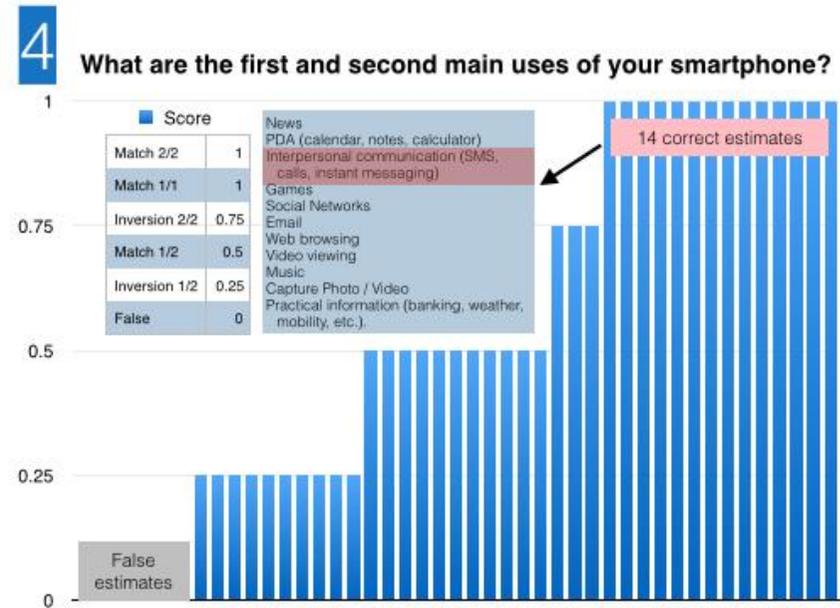
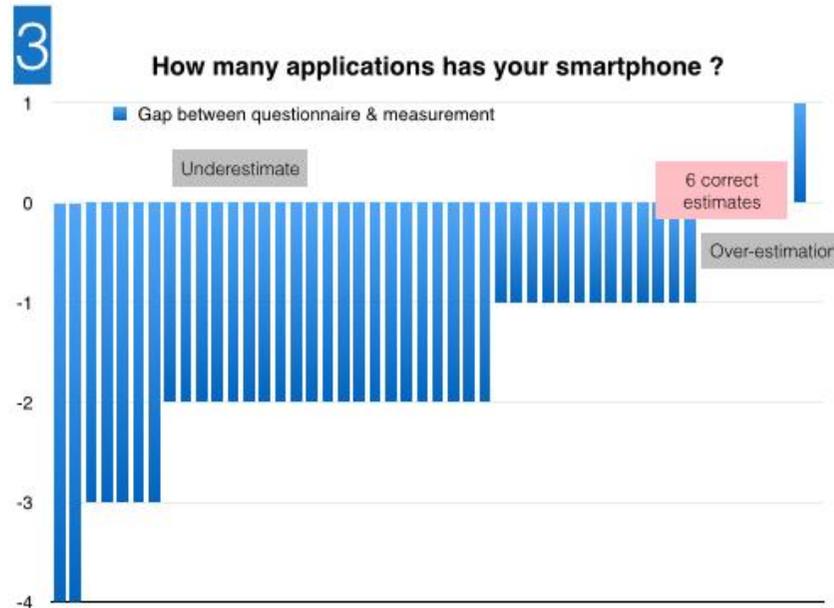
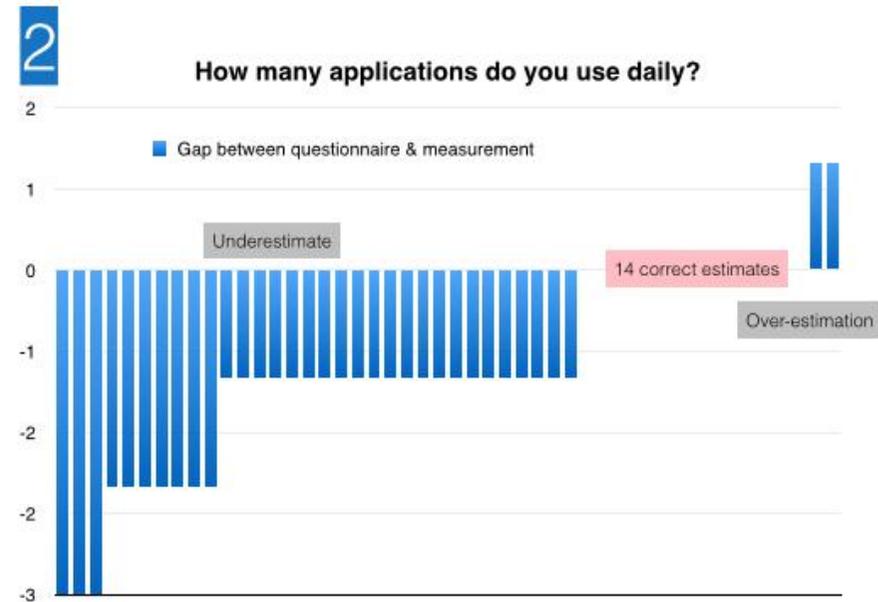
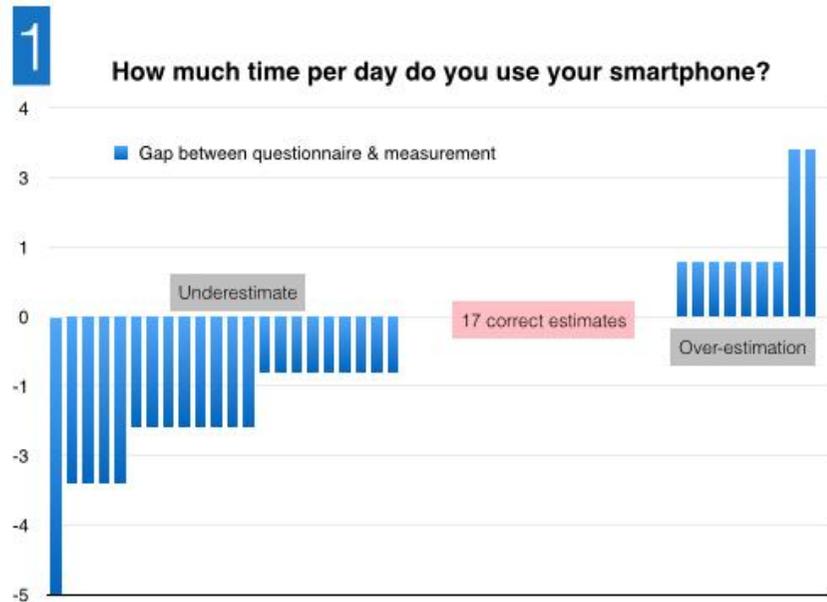
APISENSE®: Engaging participants

PRACTIC jeu-concours (march 10th - april 21th 2014)

All Campaigns (except Caen)



Interesting insights for social sciences



In-breadth monitoring – key results

- **Publications (1 PhD thesis, 1 book chapter, 2 conferences)**

- N. Haderer. **APISENSE : une plate-forme répartie pour la conception, le déploiement et l'exécution de campagnes de collectes de données sur des terminaux intelligents**. PhD thesis. Université Lille 1. Nov. 2014.
- N. Haderer, et al. **A Cloud-based Infrastructure for Crowdsourcing Data from Mobile Devices**. In: **Cloud-based Software Crowdsourcing**. 2014.
- N. Haderer, et al. **Dynamic Deployment of Sensing Experiments in the Wild Using Smartphones**. In: **DAIS'13**.
- N. Haderer et al. **A preliminary investigation of user incentives to leverage crowdsensing activities**. In: **PerHot'13**.

- **Collaborations**

- PRACTIC (Inria)
- Priva'Mov (LIRIS / CITI)
- Carbon Footprint Estimation (EQUIPPE, Univ. Lille 1)
- Indoor Air Quality (PC2A, Univ. Lille 1)
- Outdoor Air Quality (LISIC, Univ. Côte d'Opale)

- **Grants**

- Inria ADT 2012 (1 engineer)
- Inria ADT FOCUS 2014 (3 engineers)
- Microsoft Azure4Research grant

- **Transfer (<http://apisense.com>)**

- ip-label (telecom)
- Start-up (ongoing)

Conclusion

Autonomy of ubiquitous software systems

– *Elastic programming models*

- Supporting the adaptation of software systems at runtime

– *Feedback control loops*

- Promoting FCL as first class entities

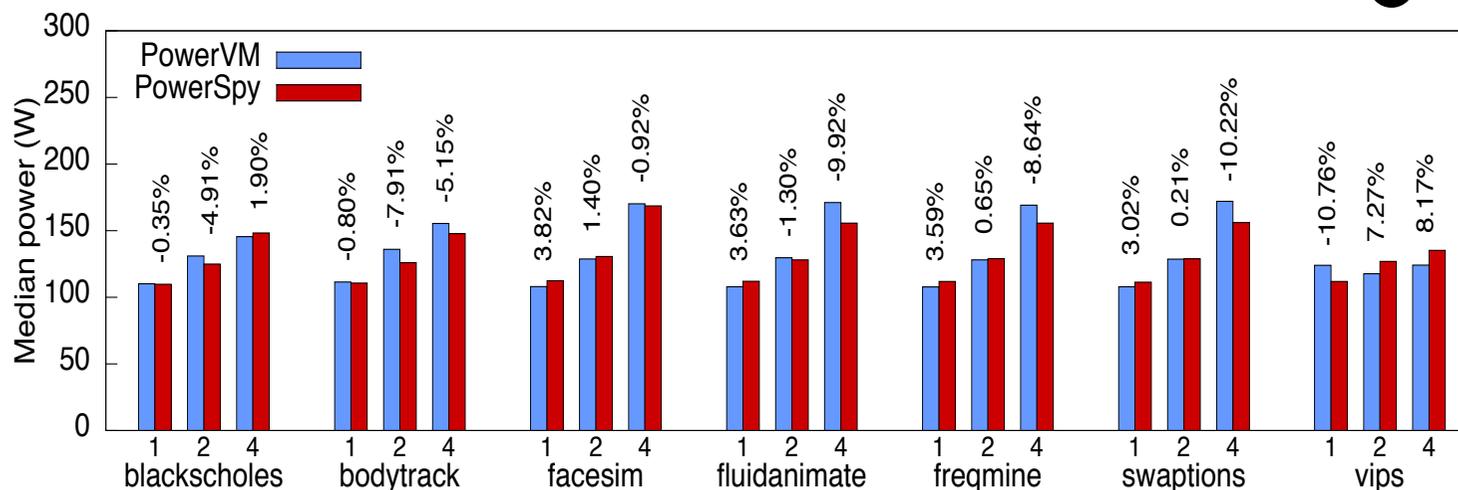
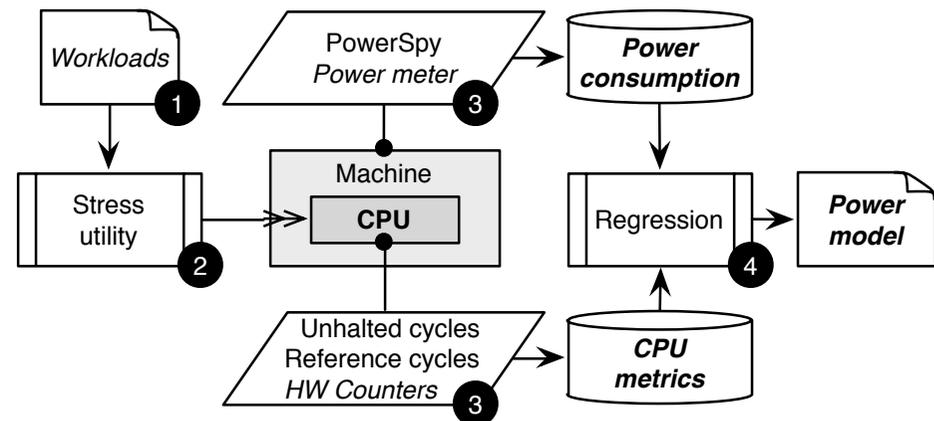
– *Efficient context monitoring middleware*

- Enabling in-depth monitoring with PowerAPI
- Controlling in-breadth monitoring with APISENSE®

Perspective for PowerAPI (1/2)

- PhD thesis of Maxime Colmant (ADEME / Lille 1)
 - *Reducing the power consumption of applications running on multi-core architectures*

- Advanced CPU features
- Virtualisation layers
- Asymmetric CPUs
- GPUs



Perspective on APISENSE® (2/2)

- PhD thesis of María Gomez (Inria)
 - *Debugging mobile applications in the wild based on the quality of experience perceived by the crowd*
 - Collecting QoE (traces + feedbacks)
 - Capturing crashes (tests)
 - (Auto) Fixing bugs
 - Deploying patches

