Ad Hoc Cloud Computing using Mobile Devices

Gonzalo Huerta-Canepa and Dongman Lee
KAIST

MCS Workshop @ MobiSys
2010

Agenda

- Smart Phones are not just phones
- Desire versus reality
- Why using mobile devices to create an ad hoc cloud?
- Considerations and Architecture
- Current implementation and evaluations
- Summary and Future Work

Smart Phones are not just phones

- Mobile phones/terminals are becoming more powerful and used for personal information processing
- They have become a central point of ubiquitous information processing
- Mobile phones are considered new portable sensors and expected to be a main driver of smart applications

Smart Phones are not just phones

- Eric Schmit, CEO of Google said
  - “you'll see things you can't even imagine now. As examples of what's being done with Android mobile phones, application that can take pictures of barcodes, identify the corresponding product, and compare prices online. An application that can take a picture of a menu in a foreign language and translate it.”

InformationWeek, 10/28/09

(http://www.informationweek.com/story/showArticle.jhtml?articleID=220900806)

Today’s Smart Phone is not enough

- People cannot do their intelligent information management only with Smart phone due to its limitations
  - CPU (~10MIPs) → no 1TB knowledge processing
  - GPU (~30MT/s) → no real-time 3D information visualization
  - Networks (~10Mbps, single hop) → not enough for multimedia transmission
  - Storage (~150GB) → no anywhere/anytime personal data storage
  - Display (~ 4”/800x400) → no UHD
  - Interface (keypad & touch) → no natural human-friendly interface
  - Battery-dependency → limited lifetime

Mobile Super Computing

- Large amount of computing & networking resources are allocated to a user's smart phone

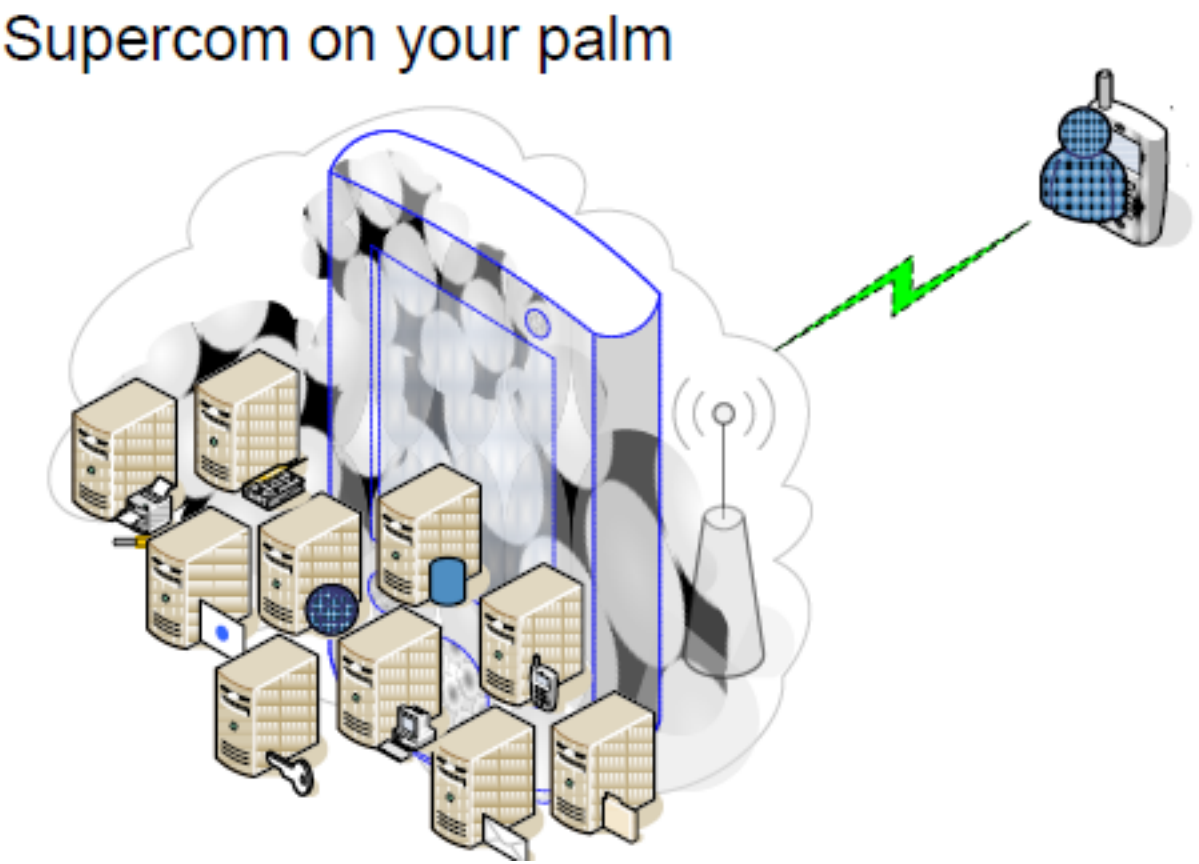

Mobile Super Computing is not always ‘the’ solution

- Cost of 3G
  - In South Korea the subscription plans for the iphone3GS (32GB) are near 70 USD per month, and if the user wants to download 1 GB of data he

Therefore, an initial solution will be to use resources from nearby devices connecting via a wireless radio

- WIFI connectivity is not always present
  - Less than 20% in US cities (Balasubramanian et al, 2010)

Ubiquitousness of Mobile Devices

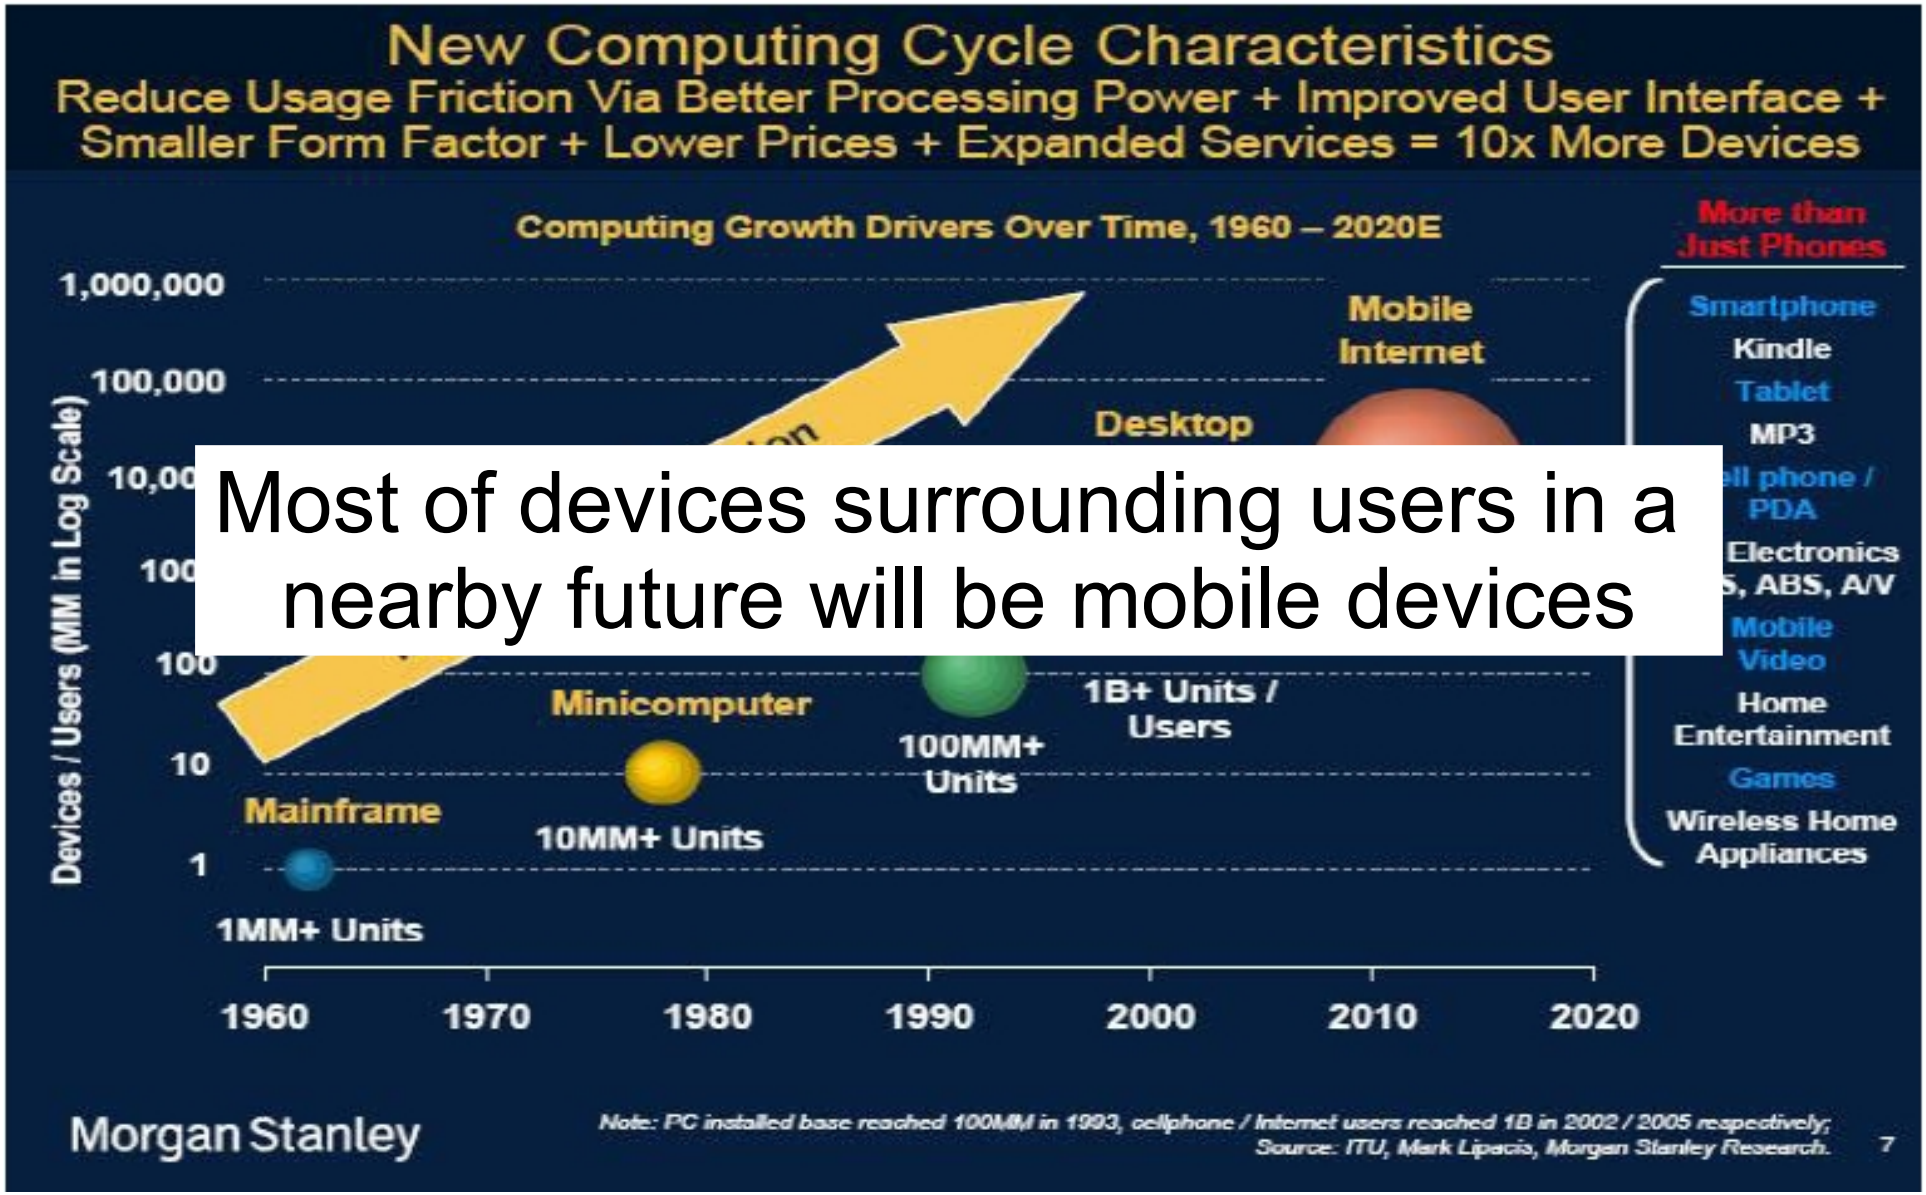

Most of devices surrounding users in a nearby future will be mobile devices

Mobile phones are a source of idle resources...

- Falaki et. Al, 2010
  - Less than 25% per hour of usage

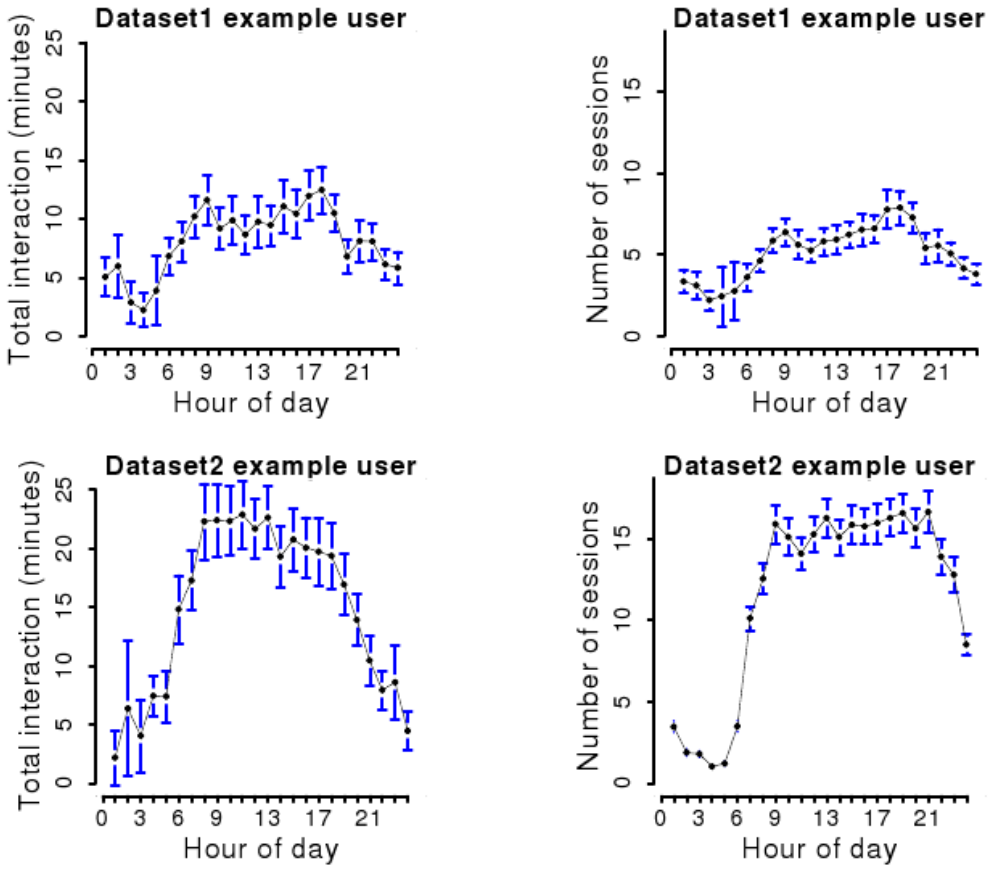

(a) Interaction time (b) Number of sessions

… and will be able to handle stronger computations

- Qualcomm's dual-core ARM CPUs clocked at up to 1.2 GHz presented at Computex 2010

  (http://www.dailytech.com/ARM+Readies+Dual+Core+CPU+Onslaught+With+12+GHz+Snapdragon/article18575.htm)

- Samsung is planning to release a quad-core ARM processor for mobile phones in 2012

  (http://www.mydigitallife.info/2010/05/03/samsung-plans-quad-core-arm-cortex-a9-processor-due-by-20122013/)

Scenario

- Task aware collaboration

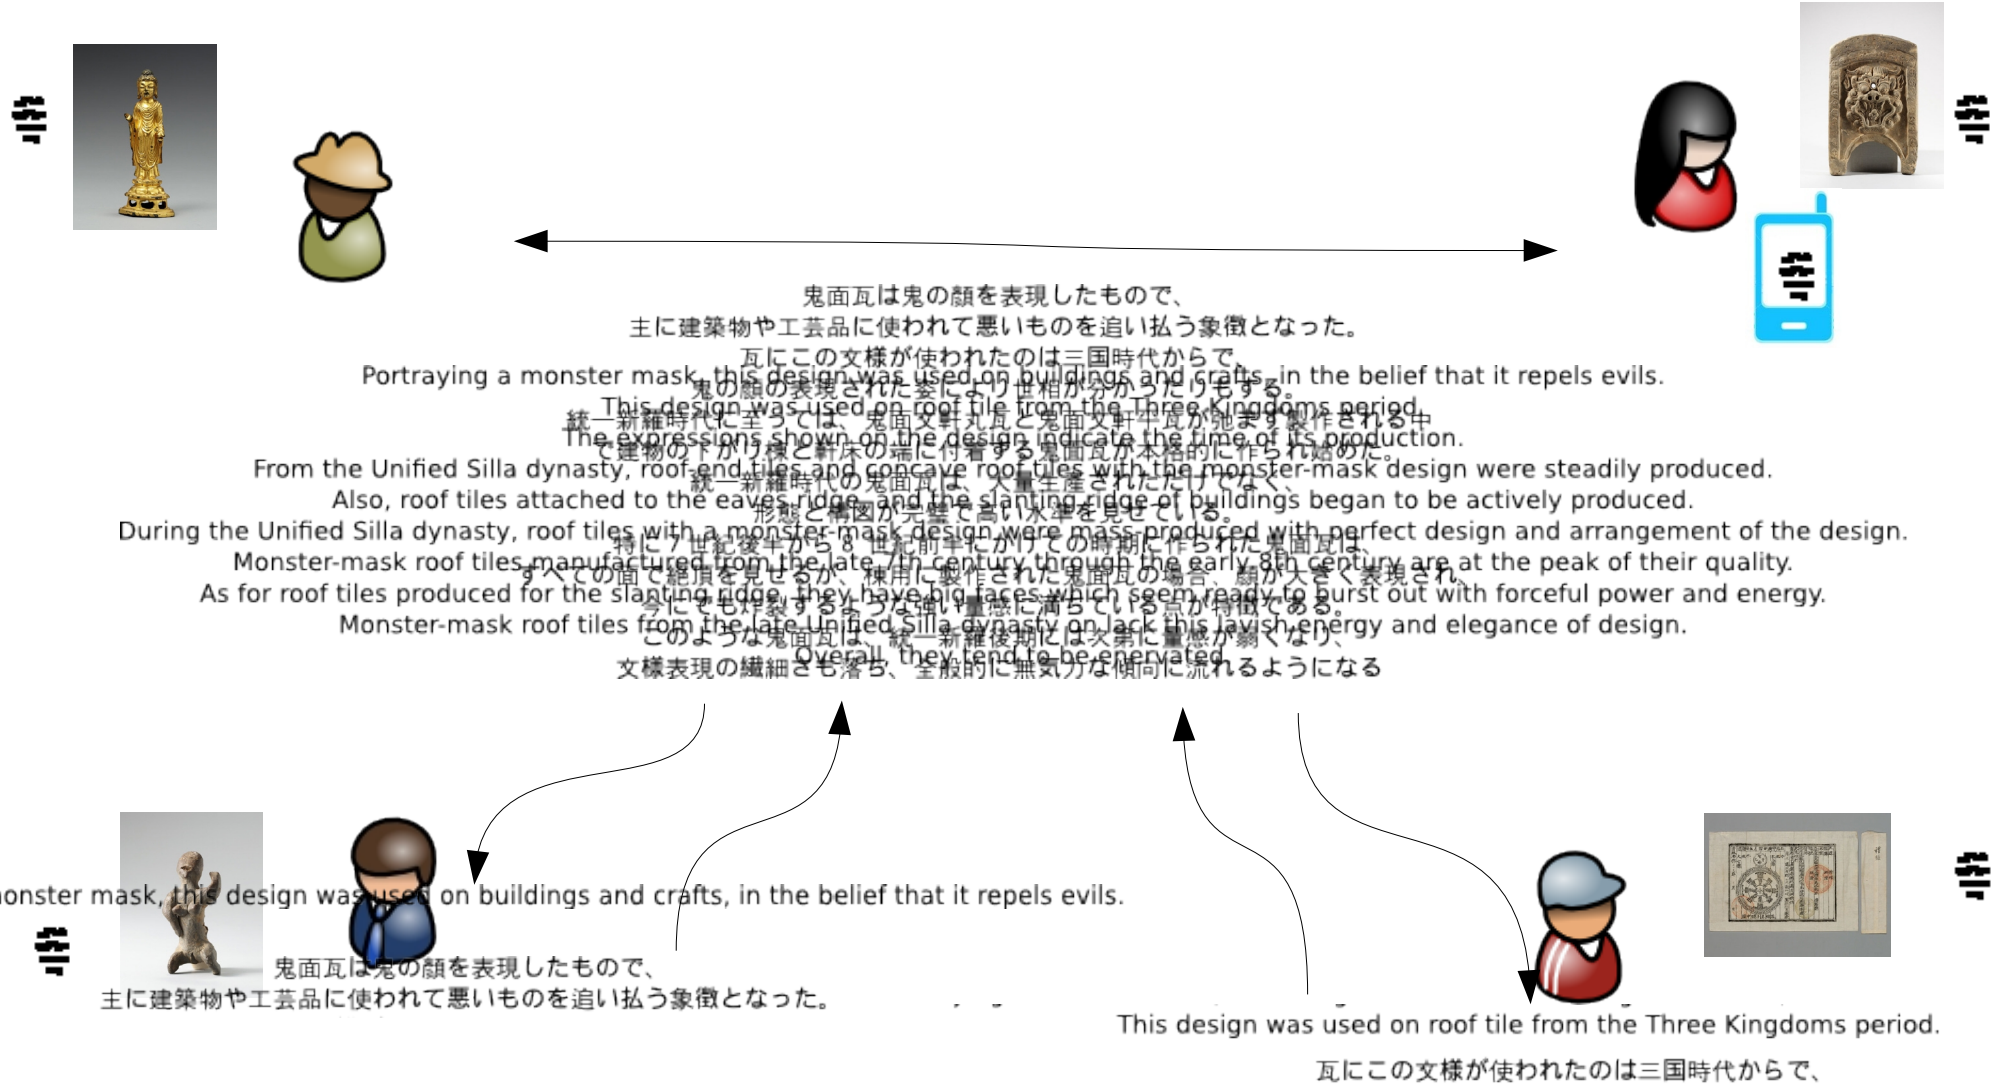

It is not always useful

- Cost of accessing or limited connectivity to networks
- Number of surrounding devices is high enough
- Tasks can run on a single device or can be divided into smaller tasks
- Users are willing to share resources (usually to save theirs)

→ It is a complement to current mobile cloud solutions not a complete alternative

Design Considerations

- Size of tasks
- Future mobile activities of users
- Activity/Task awareness
- Application behavior
- Mobility

Architecture

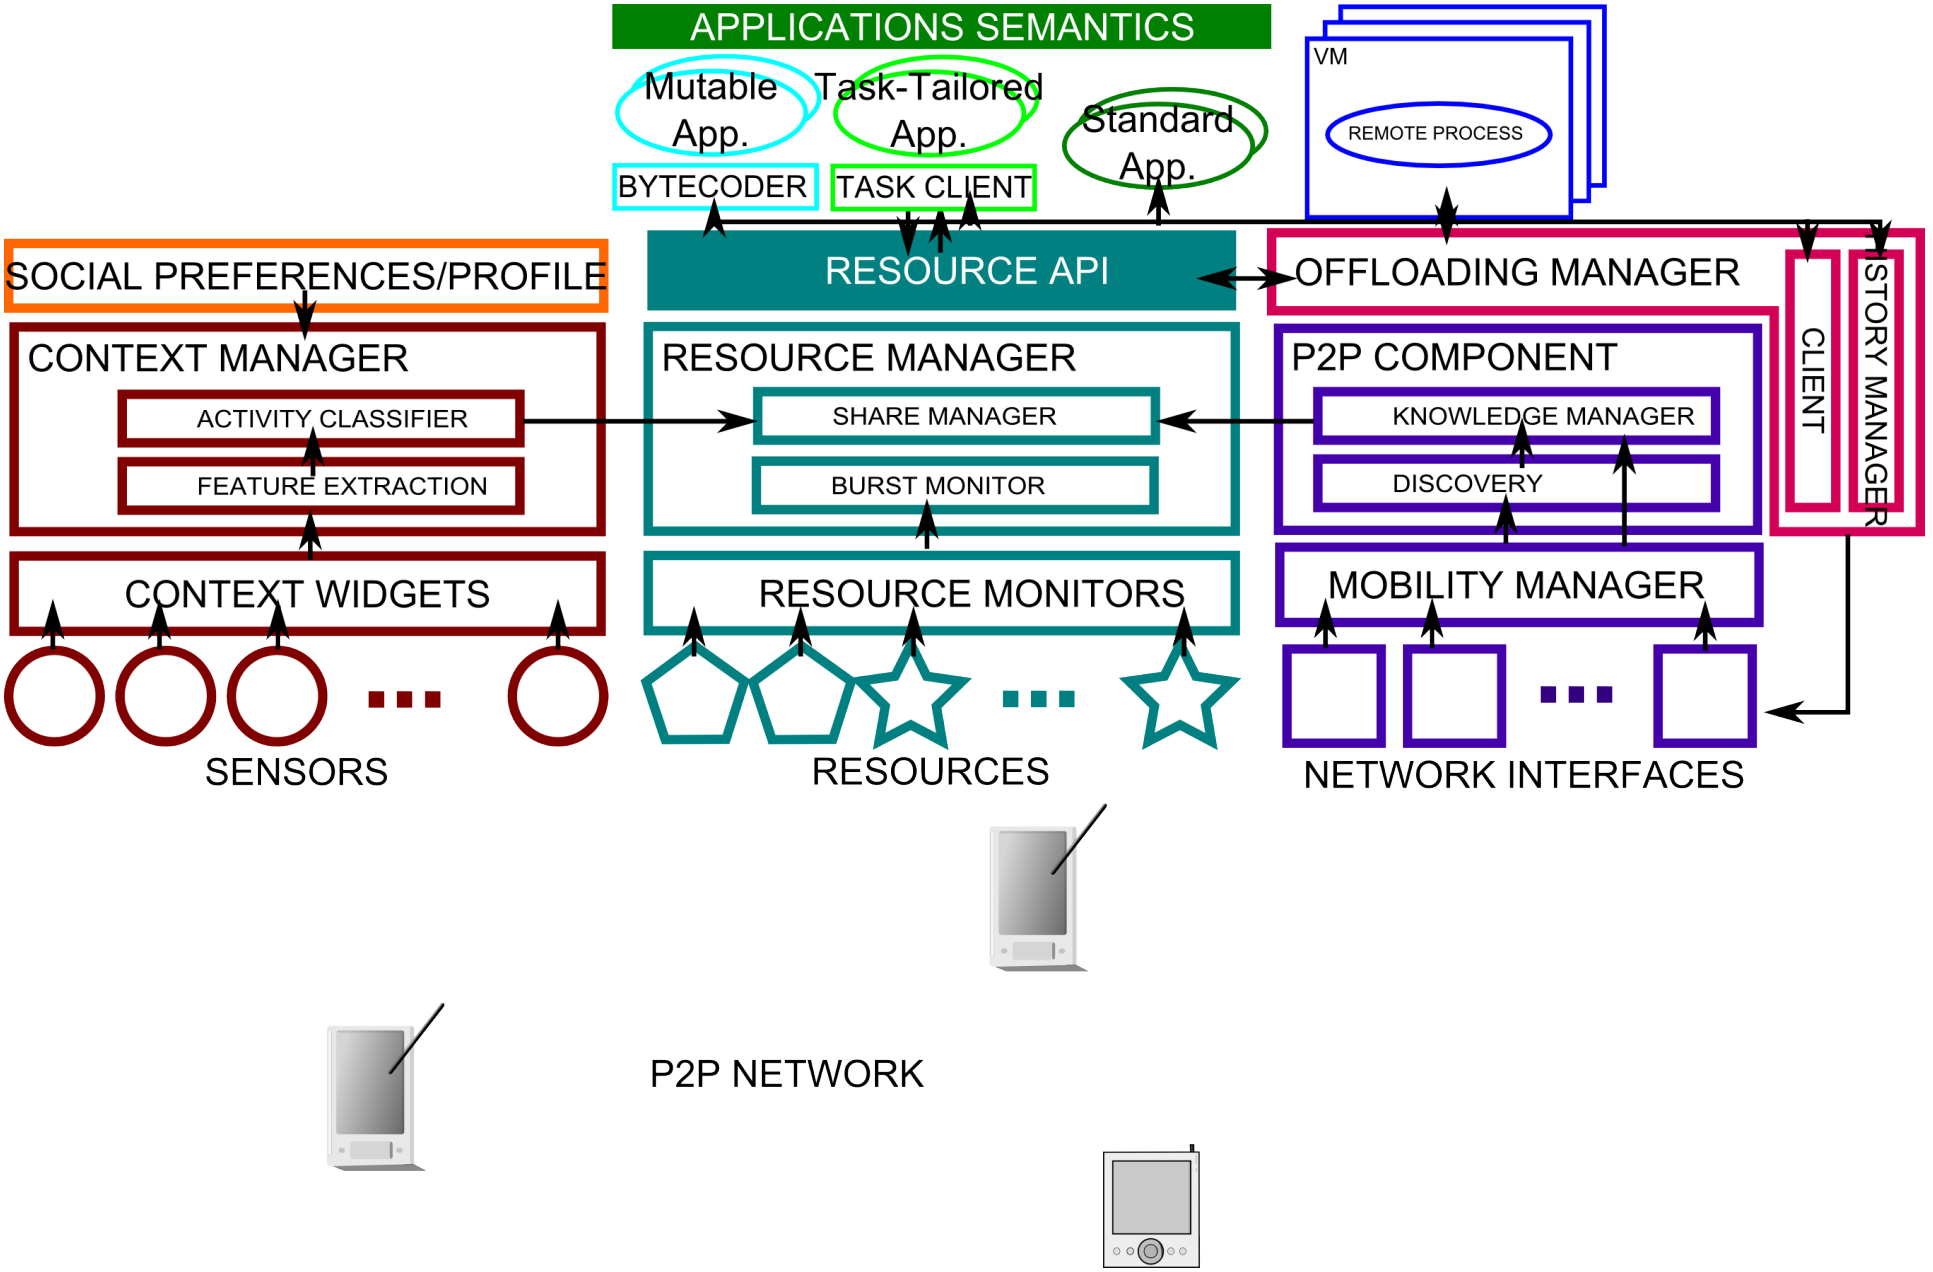

Current Implementation

- Mobile cloud computing client and an ad hoc cloud computing peer
  - Both are developed based on Hadoop
- Java as programming language
  - JamVM used on iPhones to run the framework
  - Bytecoding done with javassist
  - Communication between mobile phones based on Yaja!, an XMPP middleware (Added: RPC and Serverless messaging)

Preliminary Evaluations

- Based on translation scenario
- Cloud computing cluster based using Hadoop
  - Five servers
- Three iPhone clients
  - But only one selected as potential surrogate
- WiFi used as communication medium
- We simulate mobility
  - Creation of ad hoc cloud happens only when mobile nodes are at stable places

Preliminary Evaluations

- Ad Hoc cloud computing versus local execution
  - Performance-wise is not harmed

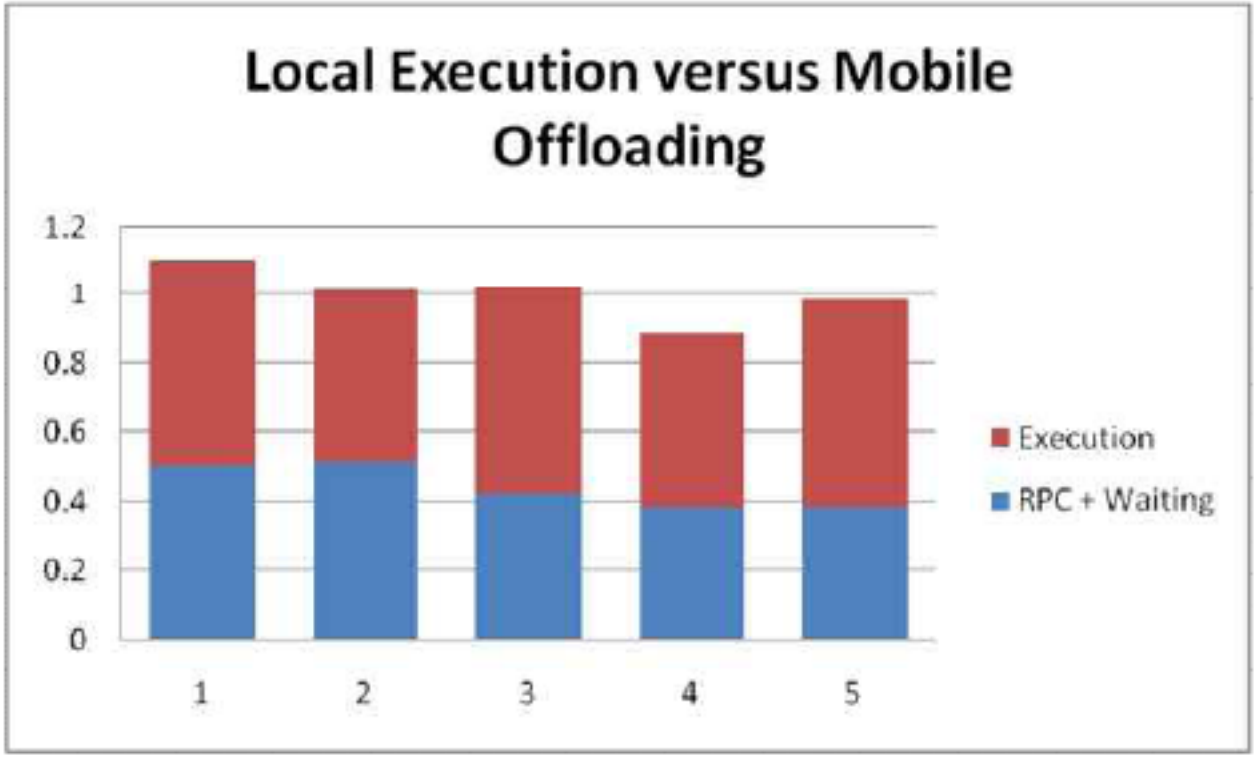

Figure Performance of Mobile Offloading compared to local execution. Results are normalized to local execution (value 1).

Preliminary Evaluations

- Hadoop suffers from low performance with small files
  - We modify this behavior configuring with an infinite number of reuse.
- The gain was of 2% for the average performance and of 3% for the best one.
  - Must tailor hadoop to mobile devices!!!
- There is one more problem with small files that we found during our test: During start up, each DataNode scans its file system and provides the NameNode with the information which files it is storing. The more files there are, the longer this takes in scanning and networking.

Summary

- Mobile phones are becoming the ubiquitous companion for users
  - But they are still resource constrained and battery dependent
  - The usage of remote surrogates is a way to overcome this deficiency
- Limited communication and cost if a barrier for mobile service computing
- A complementary solution is to take advantage of nearby mobile devices
- Main usage: collaborative or ‘same time’ tasks
- Requires to incorporate context and task awareness

Future Work

- The usage of mobility traces will be considered to create stable communities and not only places
- Considerations related to the usage of context awareness for fault tolerance will be introduced
  - Event-based alerts
- Implement a partitioning scheme for offloading between mobile phones

Questions?