NETWORKED EMBEDDED COMPUTING GROUP, MICROSOFT RESEARCH

Subjective Sensing

Mission Statement and A Research Agenda

Jie Liu, Michel Goraczko, Aman Kansal, Dimitrios Lymberopoulos, Suman Nath, and Bodhi Priyantha

7/25/2010

The essence of subjective sensing is to understand human individuals as subjective entities rather than objects in motion.

1. Introduction

Through the last decades, we are witnessing rapid changes in the world of computing. New paradigms are emerging, and they are making profound impact on how people view, use, and purchase information technology products.

One of the most visible evidences in this computing revolution is the proliferation of *mobile devices*, such as smart phones, that have continuous data connectivity, reasonable processing power and inexpensive storage space. Increasing number of sensors, such as for location, image, motion, and gesture sensing, are built into these devices. They are carried by people throughout the day. For millions of people, they have become the entry point to the cyber world and a portal for all digital contents, public or private. Another emerging trend is *cloud computing,* where data and computation are outsourced to dedicated service providers. An average developer or a small business can purchase computing resources on a need-to-use basis. Cloud computing lowered the barrier for getting access to unprecedented computing resources at a low cost (due to sharing of infrastructure and system management taskforce). A third trend, which is more subtle than consumer and cloud computing, is *community computing*. Digital devices are getting people connected via social network (e.g. facebook), shared interests (e.g. Amazon, Netflix), and shared experiences (e.g. Yelp). Contents and experiences are more than ever easier to be created and shared. The amount of data accessible by a service provider gives opportunity to solve some of the hardest problems on understanding people's preference and intent. Recommendation systems, collective intelligence, crowd sourcing are examples of mining contents and soliciting input from community of users to enable new services.

Under these emerging computing paradigms, the role of software has changed from productivity-enabler to content enabler. New business models have emerged to combine services with advertising. For many people, software products are purchased (explicitly through payment or implicitly through advertisers) not for performing computation (e.g. processing document, analyzing data, or creating content) fast, rather for organizing and efficiently accessing the right content at right time and right place. Another difference between purchasing packaged software and purchasing services is that in the latter, the user may not explicitly express, or sometimes know, what he/she wants. If we take this view that software is developed to provide services, then the better we understand costumer needs, the better we can serve them, and the more they will come to use our software. The more users we can attract, the more we can learn from them, and the better services we can provide in the future. This positive feedback loop is the key for success in the new world of computing.

Sensing is the way to understand people. Users leave traces through their online and physical activities, which effectively reflect their location, context, interaction, interests, and intents. All devices that a person encounters in the daily life – phones, vehicles, PCs, and TVs, can be leveraged as platforms for understanding users. Among these devices, mobile phones are particularly interesting (and often times overlooked). A high-end mobile phone today has more than 8 sensors, such as microphone(s), camera, touch screen, GPS, accelerometer, gyro, digital compass, and proximity sensors. People physically carry these devices, and at the same time perform a large amount of on-line activities through these devices.

Collecting, analyzing, and exploiting this information will enable us to tailor services to specific users or user groups. In the rest of this document, we focus our discussion on mobile phone based sensing.

Detecting user context and activities using mobile devices has been researched for a long time. Take location sensing as an example. There are many ways to get location information on a mobile device, such as using GPS, cell tower IDs, cell tower triangulation, WiFi or FM radio signatures, etc. These technologies represent different accuracy and power consumption trade-offs[1]. More recently, researchers have looked into using other sensing modalities such as accelerometer to determine when to turn on GPS, or selecting the right location sensing modality depending on required accuracy [5]. Similarly, extensive work exists on user state estimation, in term of standing still, walking, driving, or taking a bus [14][3][10].

However, there is still a big gap between acquiring *objective* sensor data to use such information for delivering services that meet users' *subjective* needs. Currently, our capabilities of sensing and making sense of the data collected from and around people are still limited. Although mobile phones that are carried by users throughout the day, they are activated only *sporadically* due to user attention span and energy concerns. So data gathering is also sporadic. In addition, if we are not careful, services built around personalization are under the risk of privacy infringement, thus are rejected by users.

This document is an attempt to describe the vision of *subjective sensing*, a notion for continuously sensing and learning from user activity to understand users' interests, intent, and interactions, and to support fine-tuned personalized services. We describe a research project called *Munich* (Mobile Users in a Non-intrusive Computing Hierarchy) and layout an agenda for tackling some of the key research challenges in subjective sensing, particularly in terms of system architecture, resource and data management, and fundamental service building blocks. The rest of the document is organized as follows. In section 2, we define the concept of subjective sensing. Then, in section 3, we list a few scenarios where personal, subjective data can be used to enhance information flow and user experience. In section 4, we drill down some key research challenges and give some possible directions to tackle them. In section 5, we describe the high level goal and architecture of Project Munich, and how we plan to build a system over mobile devices and backend servers to implement the vision of subjective sensing.

2. Mission Statement

Human behaviors and intents are complex but not random. They are governed by social and biological rhythms and can be learned to an extent by monitoring each individual and correlate among populations.

2.1 Activity Sensing

We start with a survey on what physical and cyber activities can be sensed from a mobile device. This is the foundation for subjective sensing systems.

---

[1] In fact, in ACM MobiSys 2010, there is an entire session devoted to this topic (http://www.sigmobile.org/mobisys/2010/)

2.1.1 Physical Status and Activity Recognition

The foundation of subjective sensing is the objective sensing capability on the mobile devices. We briefly describe what can be sensed by the phone and what sensing modalities can be used. Given enough samples and sensing duration, the following contextual sensing are done in practice or reported in research papers. This is by no means an exhaustive list, but examples of what can be done (in theory):

Context	Sensing Modality	Sensing Duration	Limitations	Related Work
Outdoor Location (infrastructure based)	GPS, Cell Tower ID, Cell Tower triangulation	Instantaneous	Power consumption is high in general	many
Outdoor Position (signature based)	WiFi or FM signatures	Instantaneous	Need profiling	Skyhook
Logical Location (e.g. store granularity)	A combination of GPS and contextual multisensor fusion (background sound, indoor light density, and carpet color scheme).	Instantaneous	Need profiling	SurroundSense [1]
Indoor Location	Step counting + direction + floormap; WiFi (FM) signatures; WiFi signal strengths; ultrasound	Instantaneous /continuous	Low accuracy; Signatures-based methods usually need profiling	Many: e.g. Radar[1] Cricket[11]
Transportation mode (e.g. still, walking, running, driving, taking a bus, etc.)	Accelerometer (+compass); GSM signal strengths; GPS traces	minutes	Need long sensing periods; Need training	[6][6]
Travel directions and speed	Accelerometer + compass; location traces + road grid info	minutes	Need long sensing period	[6][6]
Indoor/outdoor	GSM signal strength changes; GPS signal strength changes; light condition changes;	continuous	Need samples from before and after	
Facility type (e.g. home, office, bar, shops, etc.)	background sound; multi-sensor signatures;	instantaneous	Need profiling	
Audience/group (people nearby)	Sound + voice recognition; image	Continuous (for minutes)		

	+ face recognition			
Fine grained activity (watching TV, cooking, exercising, etc.)	Accelerometer + audio + infrastructure assistance	Continuous (for minutes)	Usually need body area network and multi-sensor data fusion	
Gesture (pointing, drawing in the air, throw files across devices, etc.)	Accelerometer, gyro, camera	Continuous (for seconds)	Likely need personalize training process	

However, other than location sensing, very few of these are used in practice, mainly due to the dependency on labeled training data and/or long period of data collection. Many features to be detected are subjective to uncertainties in the signals themselves. For example, individuals have different heights, arm lengths, and pace size. They carry the cell phone in different places – waist, front pocket, back pocket, knee-high pocket, or purse, which greatly changes the nature of the motion signals.

2.1.2 Cyber Activities

In compare to physical activities, cyber activities are more clearly defined, such as web site visited, search queries issued, email content, social network friends, event schedules in calendar, and media file watched or listened. Logging and using this information across applications is still a challenge, since most application cannot or is not willing to share data across application boundaries. An alternative is to logging such data inside the operating system, if we can gain trust from the users and address their security and privacy concerns. We will elaborate on this in section 2.3.

2.2 Subjective Sensing

Subjective sensing is a continuous process of sensing and learning from individual and community activities with ultimate outputs reflecting personal experiences and intent. The word subjective highlights the notion that the sensing results are modified or affected by personal views, experience, or background.

The mission of subjective sensing is to sense human as an individual with his/her own background and preference. It is to answer the question of "why" rather than simply "what" (although even answering "what" is quite challenge in many cases). For example, knowing whether a person is running is relatively simple by sensing acceleration and analyzing step periods. However, how to combine time, sound, context, and history data to differentiate whether the person is jogging for fitness or running away from danger is much harder.

In order to derive subjective data from objective physical and cyber activities, we must exploit the patterns in sensor data. We believe there are three key directions.

- **History**: Human activities are not random. It is guided by psychological and biological periodicities that can be used as strong priors to activity detection and classification. For example, if we observed in the past the user always read stock price and market news during his bus commute in the morning, and we detect that he is getting on a bus on weekday morning, we can pre-fetch those news to improve user online reading experience. Similarly, if the user repeatedly go to Chinese restaurants for lunch, when he submit a restaurant query in a new location, we can reasonably expect that he will more likely to consider Chinese restaurants and related offers.

- **Community**: Individualization does not imply that every human being is different in all aspects. Demographic, cultural background, and experiences lead people to make similar decisions. Classifying users, subject to location, time, and feature of interests, let us learn from other people's reaction within a group.
- **Probing**: Feedback extends what sensing alone cannot achieve by taking user reaction and ground truth into account. In many subjective sensing scenarios, ground truth is hard to obtain in a non-intrusive way. Designing *targeted*, *persuasive* experiments that a user is willing to annotate, share, and verify their ground truth information is an essential step towards improving learning accuracy.

So, the key challenge in subjective sensing is to enable and leverage analysis based on history, community, and feedback under the acceptance of common users.

At this point, a couple of words are due to clarify what Subjective Sensing is *not* about. Subjective Sensing is by no means to solve the generic AI problem, or to fill the "semantic gap" from data to information. Although we intend to leverage all latest results in machine learning and social sciences, the ultimate goal is not to understand human's mind. We have no ambition to *explain* or even *label* subjective behaviors; rather we want to acknowledge individual differences and to focus on their implications. In other words, we treat human mind as a black box and model the observable behaviors through monitoring and probing.

2.3 Privacy Concerns

Collecting, deriving, and exploiting personal subjective data inevitably causes privacy concerns. We believe this is manageable through both technology and business model innovations. The solution is more than simple protection, but lies in *trust*, *trade*, and *transparency*.

First of all, the user needs to trust the data collection. For that, we want to leverage the mobile devices and separate them from the cloud. People usually trust their devices and view then as personal. It is OK for the device to collect and store the data, as long as it does not share it (blindly) with any other devices or backend services without users knowing it.

Once data are collected, users may be willing to *trade* it to get advanced services. For example, current mobile users are willing to expose their location information to get traffic or turn-by-turn navigation services. They may be willing to store exercise data in the cloud if they can later track their performance or wellbeing. If the services are designed for the users' benefits, such as improving user convenience and efficiency, or saving them money, then the users are likely to share more personal data. Most banks and credit card companies give clients "trust" and "trade" level privacy preservation, however, their business models are not transparent. For example, they can sell client information to third parties, and once that happens, the clients completely lose control of the information.

So, the most challenging aspect of privacy preservation is transparency. It implies that the user should have the means to trace and control how personal information is used by various services. For example, if the service provider can *show* (rigorously or as a policy) that location information collected from the user is only used by the turn-by-turn navigation program and is not stored anywhere, then privacy should not be an issue. On the other hand, if it is used for targeted advertisement or commerce offers, then the user have the right to know. Ideally, the user should be able to delete some or all entries in the service provider's database and no trace can be used further. Unfortunately, this basic property is not respected by most online service providers.

Making transparency work across business boundaries is even more challenging. When a user task involved multiple service providers and most of them are not visible to the users, how to maintain transparency in user private information requires deep research.

3. Subjective Sensing Scenarios

We use a couple of examples to motivate how knowing subjective data about a user and her context can help providing better services. Instead of making science fictions, we specifically choose examples that are within the reach of what the current technologies can offer. The following scenarios have increasing complexities in terms of implementation and infrastructure requirements.

3.1 Adaptive Location Service

Location is a common service on mobile devices. Recreated from **Error! Reference source not found.**, Table 1 shows the accuracy and energy consumption tradeoff among some localization mechanisms on a typical mobile phone. We can see that their energy requirements can be orders of magnitude different.

Table 1. Accuracy and Energy Tradeoff for Various Location Technologies

Location Sensing Modality	Typical Accuracy	Average Energy Cost (mJ) per reading
GPS	10m	8000 (cold start)
Cell Tower ID	Cell size	20
WiFi Signatures	100m	800
Bluetooth Signatures	10m (not always available)	5000
Camera Science Matching	10~100m (not always available)	70

So, depending, subjectively, what the user's current mode (walking or driving) and what she wants to do, e.g. turn-by-turn navigation, closest gas station, local business search, or sightseeing, the system can choose the best location technologies that give enough accuracy, yet preserved battery life[5].

3.2 "Kill-Time" Mode

How to detect whether the use is in a "hurry" mode or a "kill-time" mode? The implication is profound, since a user will react to ads and offers quite differently depending on the modes.

One possible direction for tackling the problem is through a combination of 1) monitoring how frequent and how long the user uses the phone, 2) classification of online activity, e.g. websites, email accounts, tweets etc. 3) monitoring background noise, and 4) monitoring the transportation mode.

3.3 Context-Aware Reminders

Reminder services like those built into Outlook are useful to get people on schedule. However, current reminder services are solely based on time. That is, Outlook can remind the user for an appointment precisely at the time that the user set, e.g. 15 minutes before the appointment. For most users, this default leading time is always used regardless where the appointment is.

Assume that we can continuously monitoring user locations, a more advanced scenario for mobile devices is context-based reminders. One may use location, transportation options, and background noise to fine tune when to remind the user. For example:

- **Delay-aware reminders**: If the next meeting is remote, a more intelligent reminder system should take many factors into account, for example, the user's transportation options (e.g. does she have a car today), traffic conditions, how important is the next meeting (can the person be late, for how long), etc. This subjective data can be collected by knowing who else are going to the meeting, how the user reacts to similar meeting in the past, and current traffic data obtained from the web.
- **Audience-aware reminders**: Mobile phone audio signals are cues to detect who the owner is talking to over the phone line or in person. By combining continuous audio sensing and voice recognition in the cloud, the phone can remind the user when she meet someone (even unexpectedly). For example, one can create reminders in Outlook, such as "when I meet Alice next time, remind me to talk about the following things..." The cloud voice recognition is trained to recognize phrases like "Hi, Alice" and other names in the address book. The phone continuously senses in the background and when it detects greetings by recognizing "Hi" (locally), it records the next several seconds of sound input and sends it to the cloud. The cloud service recognizes the name ("Alice"), and sends corresponding reminder back to the phone. The phone now vibrates and the items about meeting Alice are brought to the foreground.

These scenarios can be further refined by subjective context, such as whether the user is in a hurry; what mood she is in; and how important the reminder is (e.g. how important is the next meeting or how often she meet the particular audience.)

3.4 Local Search and Advertising

Local search services provide accurate results when the user gives a very precise query, such as: "closest Star Bucks coffee shop". However, often times, local search queries are vague (e.g., "restaurant" or "movie") because the user may be open to suggestions or does not know enough about the area. In these cases, the relevancy of local search results becomes subjective. We can leverage the context of the query (e.g., time, weather, and traffic) and the personal preferences of the user in the past to delivery personalized results.

As a first step, we have analyzed the contextual data that affects how people react to search results. We only relied on information already being collected in the local search logs. Similar personalization can be done more advertisement (or better yet offers). If the phone can collect and expose necessary user preference and intention information, then commercial offers can be every effectively reach the target population.

3.5 Participatory Reality Search

Consider a crowd sourcing system for real time, reality data, such as street pictures, restaurant waiting time, event crowdedness, road conditions, or sun shining spots in Seattle. Imagine that when Alice issues a query such as "restaurant" into a search interface, in addition to restaurant names and address information returned on the page (ranked by her favorite cousins, etc.) , she can click a button to request real-time information from someone in that restaurant. A piece of video for the restaurant returned to her after a few minutes (with people's faces blurred). Now Alice can tell from the video the atmosphere in the restaurant, the arrangement of the tables, and even what kind of music is playing. This rich information will help her make a decision on whether to make a reservation.

The real-time reality information in this scenario comes from millions of people carrying their mobile phones and having the leisure time and incentive to participate other people's queries. For example, a mobile user can register with a central service. A subjective sensing system thus monitors the user's location, urgency, and reputation. The central service, when receiving Alice's request, tasks a few people in the restaurant who are likely to reply this kind of request, have the reputation of sending trustworthy videos, and are in the "kill time" mode. The returned information is further archived and can be used to serve other people with similar requests.

4. Research Challenges

Implementing subjective sensing on current mobile phones is facing a number of challenges. In addition to challenges on machine learning and data mining algorithms, there are several new research topics from a computing systems point of view, such as energy constraints, sensor selection, cloud-side data organization, privacy and access control, and learning from user feedback.

4.1 Sensing and Resource Management

Energy constraints: Mobile devices are battery powered. While the computing and communication capacity of these devices improve exponentially over time, the capacity of batteries are lacking behind the Moore's Law. Users' expectation for phone to last at least a day set hard constraints on what they will choose to turn on. For example, running GPS constantly to collect user location traces will deplete the battery in a few hours. So unless the phone is plugged into a car, the chance that the user turns on continues location sensing (using GPS) is very low. In addition, although some sensors, like accelerometers, do not consume much power, reading them periodically requires running the entire operating system stack on the application processor (AP), which requires orders of magnitude more energy.

Sensor selection: Energy (and communication bandwidth) constrains prevent us from turning on all the sensors and upload all the data. Data samples among these sensors can be redundant. Selecting the right set of sensors under the particular circumstances is the key to make sensing solutions practical.

Application management: Not all application needs to get access to all sensors and subjective data all the time. They may not willing share personal information derived within the application with other applications. We need to control access of physical and virtual sensors as well as the boundary between applications to prevent information leaking and resource draining.

4.2 Information Management

Signal processing: Deriving useful information from raw sensor data is nontrivial. In addition to deciding which sensor to use, the sensing fidelity varies depending on the use scenarios. These sensors are subject to random noise, user disturbance, and environmental uncertainly. Signal filtering and sensor fusion are essential to clean the data and make them useful.

Uncertainty management: Even after data cleaning, what sensor data and online traces represent are observable part of a complicated underlying human behavior and intent. Uncertainty will be inherent in this problem domain. How to represent user states, how to learn them and refine them from external signals require deep study.

Privacy concerns: Collecting and deriving personal information inevitably trigger privacy concerns. We must strike a balance between providing useful service and revealing private user information to service providers, who may use the information beyond user intent.

Data placement: Closely related to privacy concerns and resource management is the placement of subjective data. Data placed on device are more private and trustworthy. However, in order to provide services by learning across users, certain information must be shared with the cloud services. Data placement also dictates communication cost between device and the cloud.

4.3 Application Design

Service personalization: The ultimate tests of how useful human sensing is are through application scenarios. We envision a new class of applications that are tailored to individuals. These services need to anticipate how an individual user may react, to persuade them to take desired actions, and to learn from their responses. We may incorporate results from psychology and sociology with computing technologies.

5. Project MUNICH

MUNICH (Mobile Users in a Non-Intrusive Computing Hierarchy) is a research project at Microsoft Research addressing some of the systems research challenges of subjective sensing. In this project, we explore client and cloud system architecture, mobile continuous sensing, resource management, data access control, privacy preservation, and example applications. This section outlines the core components of the design.

System Architecture

Figure 1 shows an architectural design for the Munich platform. The core of the system is a set of Mobile User Classes (MUC), which captures the storage, update, and access of subjective data on a mobile device. These classes contribute to and benefit from cloud-side services to track and learn individual preference and priorities. MUC provides a set of API for high level applications to request information (under privacy preserving access control policies), register for long term monitoring, and contribute their application specific knowledge about the user.

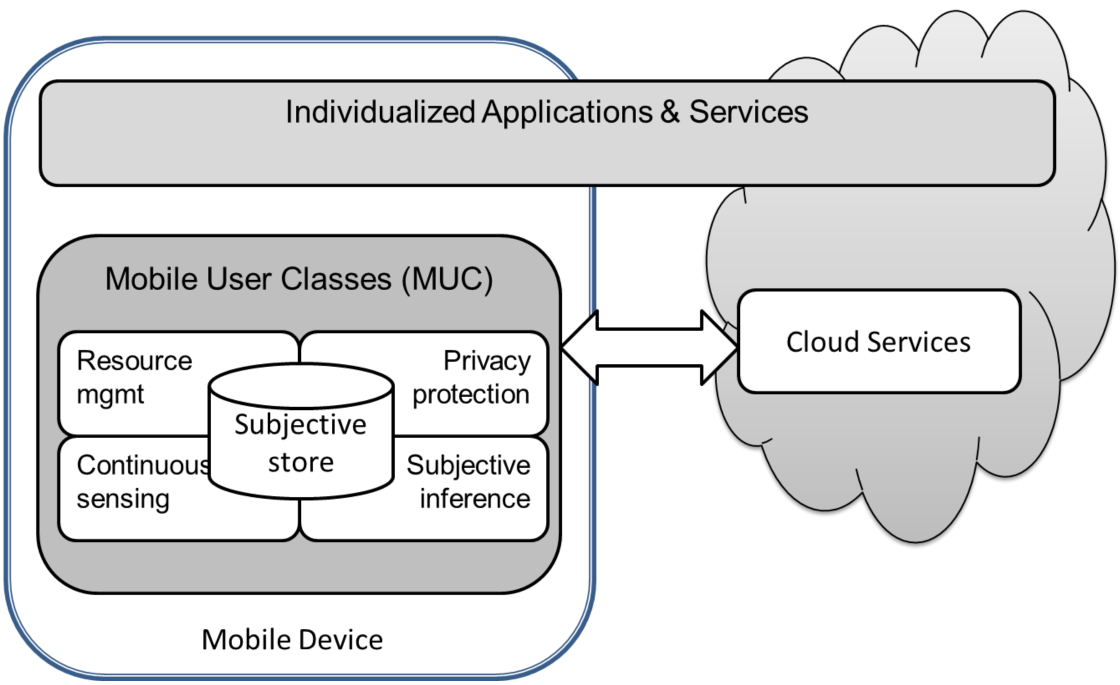

Figure 1. Munich Architecture.

Subjective Store

The core of MUC is the data abstraction of subjective data. Since we must keep histories of users' physical and online activities, and may use the device as the first and foremost trusted storage for such information, the storage management must be addressed. We are facing a tradeoff in storage. Raw data are useful when we do not know a priori what will be useful in the future. However, keeping all data in its raw form is not efficient to use and process.

Uncertainty is fundamental in subjective data. We not only need to represent them in an application programmer friendly way, but also need to make sure that confidence and beliefs are not over/miss counted when we run multiple sensing and inference algorithms. In essence, tracking *information lineage* and answering questions like "how do you arrive in such belief" makes subjective store very different from existing data structures.

Continuous Sensing

In order to update subjective store with useful user activity, the mobile device needs to overcome the current use pattern that the phone is only activated sporadically by the user. The main obstacle that prevents the phone from continuously sensing its context is the power constraints. In the current phone architecture, the sensors are driven by the main processor. Although the sensor themselves uses very little energy to operate, the main processor consumes significant energy to even read simple sensor data periodically. At MSR, we designed the LittleRock platform to facilitate continuous low power sensing on mobile phone [12]. Figure 2 shows the architecture design for LittleRock, where the sensing modules are attached to a low power microcontroller rather than the main application processor. Figure 3 shows a prototype board with a few sensors. As a result, the microcontroller can be always on, driving the sensors, and keeps the main processor in a standby mode.

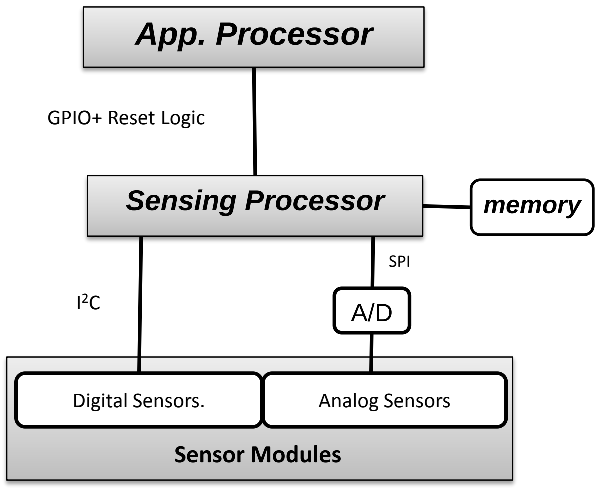

Figure 2. The LittleRock Architecture for Mobile Phones to Enable Continuous Low Power Sensing.

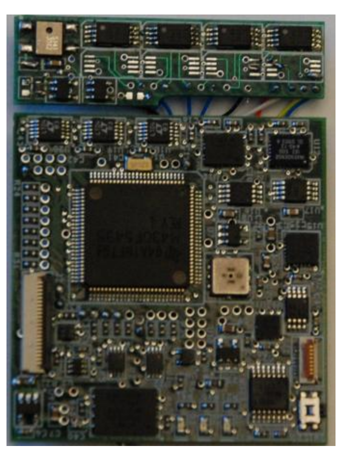

Figure 3. A Prototype LittleRock Board with Sound Capturing Extension. The board has an accelerometer, a digital compass, a gyro, a temperature sensor, and a pressure sensor.

For example, if we implement a pedometer application for step counting using on-board accelerometers, then the LittleRock design is two orders of magnitude more energy efficient than using the current architecture. In other words, instead of having the phone battery drained in about 4 hours, the pedometer can run on LittleRock for more than 10 days.

The LittleRock hardware design itself does not solve all continuous sensing problems. For example, GPS is one of the most useful, yet most power consuming sensors on the phone, regardless of the usage of the application processor. Furthermore, it is unnecessary and unpractical to continuously logging all sensor data blindly. A software stack for tasking and managing LittleRock platform and sensors is under development.

Subjective Inference

The subjective inference engine encapsulates a set of learning and data mining algorithms, with the help from cloud services, to derive more information from raw sensor data. In addition to the algorithms themselves, when and how they are triggered is an interesting. Depending on how urgent a piece of information is; what connectivity that phone has; and how much resource (e.g. battery energy) is left in the phone, certain parts of the algorithm can be offloaded to the backend cloud services.

Another reason to leverage cloud services is the access to models learned or built from mining communities of users. The cloud needs to assist the phone with the latest findings from community users to make local inference up to date.

Resource Management

Subjective sensing imposes significant challenges to resource management on the phones. There will be long running background tasks (e.g. continuous sensing), long running queries and processing (e.g. subjective inference), and communication with the cloud. This level of intelligence changes the current mobile phone use pattern, which is primarily to react to user commands.

Core resources in this context are energy, sensing, and communication. In addition to be more sophisticated about the triggering, scheduling, and offloading, the transparency of resource usage is important for users to directly prioritize and manage the phone behavior. Monitoring and deriving per software process energy consumption is a fundamental service for resource management.

Privacy Control

We foresee two aspects of privacy control in Munich: access control and data perturbation. Access control is conceptually straightforward. A user should have intuitive and comprehensive control on what information is used by what services and application. She should have the capability to track and invoke protection mechanisms on how the information is used.

Data perturbation is a way to share private data by adding noise to them such that the information is private under the differential privacy definition [4], yet still be useful to community analysis such as aggregation and grouping [14].

Cloud Services

The cloud services that are built in Munich are different from traditional services. We foresee two kinds of cloud services to help derive subjective data: private services and public services. Private services are trusted computing and storage resources in the cloud or on user PC. They assist mobile devices to archive and learn user activities and preferences. The public cloud, on the other hand, relies on community data to provide cross users training data to refine private models.

In Munich, we plan to build new application and improve existing cloud applications using subjective data obtained from subjective sensing. Some example scenarios are listed in section 3. These applications will provide real constraints and evaluation mechanism to show how useful subjective sensing is and how users accept them.

6. Conclusion

In this document, we describe the opportunity and challenge to derive personalized sensor data through subjective sensing. We believe that this opens a new dimension in leveraging personal experiences and priorities to deliver better services to user and create bindings between users and service providers. The technical challenges are significant. The Munich project is a step towards building a practical subjective sensing system and show how it may enable new services.

7. References

[1] Martin Azizyan , Ionut Constandache, and Romit Roy Choudhury, “SurroundSense: mobile phone localization via ambience fingerprinting.” In *Proceedings of the 15th Annual international Conference on Mobile Computing and Networking* (Beijing, China, September 20 - 25, 2009). MobiCom '09. ACM, New York, NY, 261-272.

[2] Paramvir Bahl, Venkata N. Padmanabhan, “RADAR: An In-Building RF-Based User Location and Tracking System,” IEEE INFOCOM 2000

[3] Tanzeem Choudhury, et. al. “An Embedded Activity Recognition System,” Pervasive Computing, April-June, 2008, pp 32-41.

[4] Cynthia Dwork, “Differential privacy,” in ICALP 2006

[5] Kaisen Lin, Aman Kansal, Dimitrios Lymberopoulos, Feng Zhao, “Energy-Accuracy Aware Localization for Mobile Devices,” ACM MobiSys 2010

[6] Jonathan Lester, Tanzeem Choudhury, and Gaetano Borriello, “A Practical Approach to Recognizing Physical Activities,” Pervasive, 2006

[7] Ben Nham, Kanya Siangliulue, and Serena Yeung, “Predicting Mode of Transport from iPhone Accelerometer Data,” Stanford University class project. (http://www.stanford.edu/class/cs229/proj2008/NhamSiangliulueYeung-PredictingModeOfTransportFromIphoneAccelerometerData.pdf)

[8] Jeongyeup Paek, Joongheon Kim, Ramesh Govindan, “Energy-Efficient Rate-Adaptive GPS-based Positioning for Smartphones,” ACM MobiSys 2010

[9] Jun-geun Park, Ben Charrow, Dorothy Curtis, Jonathan Battat, Einat Minkov, Jamey Hicks, Seth Teller, Jonathan Ledlie, "Growing an Organic Indoor Location System," ACM MobiSys 2010.
[10] Donald J. Patterson, Lin Liao, Dieter Fox, and Henry Kautz, "Inferring High-Level Behavior from Low-Level Sensors," UbiComp 2003, LNCS 2864, pp. 73–89, 2003.
[11] Nissanka B. Priyantha, Anit Chakraborty, Hari Balakrishnan, "The Cricket location-support system," MOBICOM 2000
[12] Bodhi Priyantha, Dimitrios Lymberopoulos, and Jie Liu, "LittleRock: Enabling Energy Efficient Continuous Sensing on Mobile Phones," Microsoft Research Tech Report, MSR-TR-2010-14, 2010.
[13] Energy-Aware Localization Using Mobile Phones
[14] Vibhor Rastogi and Suman Nath, "Differentially Private Aggregation of Distributed Time-Series with Transformation and Encryption," SIGMOD 2010
[15] Timothy Sohn, Alex Varshavsky, Anthony LaMarca, Mike Y. Chen, Tanzeem Choudhury, Ian Smith, Sunny Consolvo, Jeffrey Hightower, William G. Griswold, and Eyal de Lara, "Mobility Detection Using Everyday GSM Traces," UbiComp 2006.
[16] Yi Wang, Jialiu Lin, Murali Annavaram, Quinn Jacobson, Jason I. Hong, Bhaskar Krishnamachari, Norman M. Sadeh, "A Framework of Energy Efficient Mobile Sensing for Automatic User State Recognition," ACM MobiSys 2009.
[17] Zhenyun Zhuang; Kyu-Han Kim; Jatinder Pal Singh, "Improving Energy Efficiency of Location Sensing on Smartphones," ACM MobiSys 2010