UrbanSense08

Proceedings of

International Workshop on Urban, Community, and Social Applications of Networked Sensing Systems

November 4, 2008

Raleigh, North Carolina, USA

Table of Contents

Sensing and Localization Panel

Chair: Andrew Campbell (Dartmouth)

Sensing and Measurement Panel

Chair: Landon Cox (Duke)

Preface

Sensing is going mobile and people-centric. Sensors for activity recognition and GPS for location are now being shipped in millions of top end mobile phones. This complements other sensors already on mobile phones such as high-quality cameras and microphones. At the same time we are seeing sensors installed in urban environments in support of more classic environmental sensing applications, such as, real-time feeds for air-quality, pollutants, weather conditions, and congestion conditions around the city. Collaborative data gathering of sensed data for people by people, facilitated by sensing systems comprised of everyday mobile devices and their interaction with static sensor webs, present a new frontier at the intersection between pervasive computing and sensor networking.

Welcome to the Workshop on Urban, Community, and Social Applications of Networked Sensing Systems (UrbanSense). This workshop is the third in a series of meetings held at SenSys over last few years.

The first workshop in 2006 dealt with the concept of the world sensor web and the second at SenSys 2007 focused on sensing on everyday mobile phones in support of participatory research.

This workshop promotes exchange among sensing system researchers involved in areas, such as, mobile sensing, people-centric and participatory sensing, urban sensing, public health, community development, and cultural expression. It focuses on how mobile phones and other everyday devices can be employed as network-connected, location-aware, human-in-the-loop sensors that enable data collection, geo-tagged documentation, mapping, modeling, and other case-making capabilities.

This year's meeting focuses on four topic areas: (1) Participatory sensing; (2) Emerging applications; (3) Sensing and localization; and (4) Sensing and measurement. The program comprises 11 papers presented in the form of four panels to encourage discussion of the ideas, challenges and technical issues presented by the speakers.

The program also includes an invited panel organized and chaired by Frank Bentley (Motorola) on "Applications of sensor-enabled mobile devices - how your phone can be location-aware, keep you fit, and save your life". Panelists include Rahul Nair (Yahoo!), Pedja Klasnja (Intel Research/UW) and Tim Bergin (Motorola). We would like to thank Frank for putting this exciting panel together.

We would also like to thank the technical program committee for their hard work in putting the program together:

Frank Bentley, Motorola
Assaf Biderman, MIT
Péter Boda, Nokia Research
Gaetano Borriello, University of Washington
Andrew Campbell, Dartmouth College
Hae Don Chon, Samsung
Landon Cox, Duke University
Deborah Estrin, UCLA
Lama Nachman, Intel Research
Tapan Parikh, UC Berkeley
Matt Welsh, Harvard University
Sean White, Columbia University
Feng Zhao, Microsoft Research

We are grateful to Emiliano Miluzzo (Dartmouth College) who managed the workshop webpage where you will find the papers and slides. Emiliano also put the proceedings together.

We look forward to an informative and productive meeting.

Andrew Campbell
Deborah Estrin

Raleigh, North Carolina, USA
November 4, 2008

Evaluating Participation and Performance in Participatory Sensing

Sasank Reddy, Katie Shilton, Jeff Burke, Deborah Estrin, Mark Hansen, Mani Srivastava
Center for Embedded Networked Sensing, University of California, Los Angeles (UCLA)
Email: {sasank, kshilton, jburke, mbs}@ucla.edu, destrin@cs.ucla.edu, cocteau@stat.ucla.edu

Abstract

Because participatory sensing – targeted campaigns where people harness mobile phones as tools for data collection – involves large and distributed groups of people, participatory sensing systems benefit from tools to measure and evaluate the contributions of individual participants. This paper develops a set of metrics to help participatory sensing organizers determine individual participants' fit with any given sensing project, and describes experiments evaluating the resulting reputation system.

I. Introduction

The rapid adoption of mobile phones over the last decade and an increasing ability to capture, classify, and transmit a wide variety of data (image, audio, and location) have enabled a new sensing paradigm – participatory urban sensing – where humans carrying mobile phones act as, and contribute to, sensing systems [1], [2], [3]. In this paper, we discuss an important factor in participatory sensing systems: measurement and evaluation of participation and performance during sensing projects.

In participatory sensing, mobile phone-based data gathering is coordinated across a potentially large number of participants over wide spans of space and time. We draw from three pilot projects to illustrate participatory sensing and describe the unique challenges to measurement and evaluation provided by "campaigns": distributed and targeted efforts to collect data. Project Budburst [4], Personal Environmental Impact Report (PEIR) [5], and Walkability all situate "humans in the loop", but have critical differences in their goals and challenges (Table I).

Campaign	Goal	Data Collection	Evaluation Challenges
Project Budburst	Gather data about the flowering of native plants to study climate change	Mobile phones upload time-stamped, geo-tagged plant photos	Ensuring data quality to meet scientific standards, including photo resolution and reliable capture of events of interest
PEIR	Enable individuals to collect data about environmental impact and exposure	GPS traces create estimates of carbon emissions/pollution exposure	Continual and long-term participation yields most meaningful results
Walkability	Gather data about neighborhood sidewalk hazards	Mobile phones upload geo-tagged photos of cracks, gaps, and impediments	Systematic coverage and verification to make a strong case to city councils

TABLE I
GOAL, DATA COLLECTION, AND EVALUATION CHALLENGES ASSOCIATED WITH CAMPAIGNS.

These pilots have raised complex problems of reputation and reliability. "Human in the loop" sensing relies upon community expertise, drawing upon mass contributions in much the same way as web-based systems such as Wikipedia and Slashdot [6]. But these projects also demand data quality that meets criteria set by scientists or legislators. Participatory sensing reputation metrics must therefore incorporate expertise, data quality, credibility, and certainty, while encouraging participation and development of expertise among amateur volunteers.

To address this challenge, we have developed two indexes for measuring participation and expertise among campaign participants: cross-campaign and campaign-specific measurements. Cross-campaign metrics are résumé-like measures of previous experience and commitment. These include number of previous campaigns undertaken and the success of a participant in previous campaigns. Experience metrics can play a key role alongside other factors such as sensing modality, coverage, and cost in enabling campaign monitoring services to select participants who can achieve the highest effectiveness for a particular data collection initiative. Tracking participation in this way is not new; it has been employed widely by Internet businesses and services [7]. Systems that provide a marketplace for commissioned work, such as Amazon Mechanical Turk and GURU.com, keep detailed statistics tracking the performance of requesters and workers [8], [9]. Systems for question answering, such as Amazon Askville and Yahoo Answers, use credits to track a participant's performance [10], [11], and auctioning systems such as E-Bay have transaction ratings to help evaluate whether a particular participant is trustworthy [7].

Campaign-specific metrics provide project organizers with something existing systems do not offer: evaluation of participation during the deployment of a campaign over weeks or months. Instead of the résumé-like measures used by existing reputation systems, we compare campaign-specific metrics to a project review. As a campaign progresses, a "watchdog" module can observe quality and utility of a participant's contribution relative to campaign needs. The module can then send alerts and recommendations to participants, dynamically change incentives for existing participants, and recruit new participants to the campaign.

II. Cross-Campaign Participation: The Résumé

Our work builds on ideas of monitoring participant behavior and provides metrics to evaluate participation and performance in sensing campaigns. To suggest metrics useful to filter or rank potential participants based on past performance, we consider the requirements volunteers must meet before joining a campaign and how the metrics that make up these requirements are delineated. Cross-campaign metrics can relate to either campaign participation or campaign performance.

A. Campaign Participation

Participation requirements allow a campaign organizer to recruit participants that have a certain level of experience or have been active recently. Participation metrics include: a) number of campaigns volunteered for, b) number of campaigns accepted for, c) number of campaigns participated in, and d) number of campaigns abandoned. Individual metrics can be associated with other information about a campaign, such as size, lifetime, and type of sensing required. Examples of requirements could include: participants who have been accepted for image-sensing campaigns in the last 6 months, have participated in a certain number of location-sensing campaigns, or have less than 10% abandonment rate.

B. Campaign Performance

To create metrics that represent a participant's performance, campaign organizers need a language to express the campaign contributions they require. A successful contribution can be defined according to a number of qualities, including: what sensor type should be used and what modalities employed; the spatial or temporal context in which the sample is taken; and timeliness, relevance and quality of the sampled data. Timeliness represents the latency between when a phenomenon is sampled (or occurs) and when it is available to a data processing module. Relevancy indicates how well the sample describes the phenomenon of interest ("did the participant photograph a flower?") Quality describes the ability of the system to determine a particular feature in a sample ("can the system detect a sidewalk hazard in this photograph?") Quality includes the probability of detection, probability of a false positive, or probability of a false negative. Campaign performance may also include metrics that describe the responsiveness of a participant, or similarly, the amount of responsibility taken by a participant. Campaign organizers might consider how frequently a participant checks in with the system, whether a participant uploads regularly, or whether the participant takes privacy precautions with their data such as blurring third party images in photographs. Because the meaning and importance of each of these variables can vary based upon the needs of the campaign, it will be important for campaign organizers to have control over setting definitions and levels of cross-campaign metrics.

Using the above metrics, campaign organizers can define a useful campaign entry and set benchmarks that indicate participant success. Performance benchmarks can be absolute, set limits for performance categories, or relative, based on performance of other members. Also, the scales can be translated into user-friendly forms for querying such as the 5-star system popular on many Internet platforms [12]. As future work, we will consider enabling participants to set their own benchmarks for success.

III. Campaign-Specific Participation: The Project Review

Analysis of pilot campaign participation has shown that it is important to monitor participant contributions while a campaign is running, as well. Participants may realize during the course of a campaign that data collection does not fit their schedule or interests, or organizers may discover that participants are not keeping up with their data collection duties. Campaign-specific monitoring enables a monitoring service to adapt the participant list, coordinate reliable contributors to collect or verify information, or adapt feedback and incentive mechanisms.

To generate campaign-specific participation and performance measures, campaign organizers could choose several mechanisms. Campaigns could incorporate a "calibration" phase paired with reoccurring "check-ups" where experts or campaign organizers obtain ground truth information for a particular area of interest. Participants would then be coordinated to monitor the same area. The observations made by participants could then be compared to the ground truth to obtain a reliability measure. For campaigns where obtaining ground truth information is not practical or possible, indirect "collective" observations made by participants can generate a reliability score. In this case, geographic coverage overlaps between participants would be found, and reliability scores can be calculated by measuring the similarity of observations in these areas. Evaluating participant reliability by comparing overlapping results is similar to the social science practice of calculating interrater reliability: consistency among responses when assigning values to subjective data [13].

We propose a mathematical model to represent campaign-specific performance, update it based on measures of participant reliability, and translate it into a participant trust metric. Existing reputation systems used in applications include: cumulative, where a user's reputation ratings are summed; average, where the reputation score is computed by averaging all scores; blurred, where a weighted sum is used to down weight old ratings; and adaptive, where the current reputation score affects to what degree new observations make a difference [14].

The above mentioned reputation systems only capture stochastic uncertainty (due to randomness of the system). We instead want a reputation system that captures both stochastic and epistemic uncertainty (due to lack of knowledge about the randomness of the system) so we adopt the Beta distribution to model the reputation of a participant. The distribution is indexed by two parameters, alpha and beta, which define the number of successful and unsuccessful transactions that have occurred in the past. A participant's reputation can be found by calculating the expectation of the Beta distribution, E(alpha,beta) which is the stochastic uncertainty. The confidence factor is the posterior probability given the actual expectation value lies within an acceptable level of error [15]. Prior to any interaction, alpha and beta are set to 1, which results in a uniform distribution where all values are considered equally likely. We refer the readers to [16] for a full description of the Beta distribution.

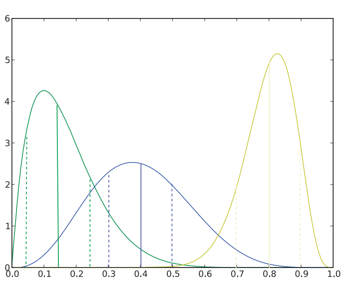

$$f(p|\alpha,\beta) = \frac{\Gamma(\alpha+\beta)}{\Gamma(\alpha)\Gamma(\beta)} p^{\alpha-1}(1-p)^{\beta-1}$$

$$E(\alpha,\beta) = \frac{\alpha}{\alpha+\beta}$$

$\alpha =$ successful entries $+ 1$

$\beta =$ unsuccessful entries $+ 1$

Beta plot generated for three participants (alpha, beta): (2,10), (4,6), (20,5) with mean of (0.16), (0.40), (0.80) and confidence value of (0.67), (0.48), and (0.80) when 0.10 is used for the acceptable level of error.

Fig. 1. Density of Beta functions for various types of participants.

Figure 1 shows the density of Beta function for different participants. The acceptable error level was defined as 0.1. By having more evidence for a certain hypothesis, as is the case with the participant with alpha of 20 and beta of 5, the level of confidence is large, 0.80. A participant with only a few entries, alpha of 4 and beta of 6, the confidence factor is significantly lower, 0.48. The Beta formulation for reputation affords us other features that are useful in monitoring participants as well. For instance, an aging factor can be introduced to account for quality variations over time by discounting past contributions as a campaign executes, and higher weights can be introduced for contributions that are made in high priority contexts (discussed in Section IV) [12]. Also, contributions do not have to be binary: ratings between 0 and 1 can be made by appealing to the Dirichlet process [17].

IV. Preliminary Evaluation

Since we are in the preliminary stages of several campaigns, we focus on evaluating campaign-specific participation metrics using information gathered from pilot studies of Walkability, PEIR, and Project Budburst campaigns.

A. Walkability and Image Quality

To pilot the Walkability campaign, we asked 6 participants to walk the Westwood neighborhood of Los Angeles and take geo-tagged images of sidewalk segments with visible damage such as cracks. The campaign ran for 2 weeks. Figure 2 shows an analysis of whether participants' contributions were adequate to assess the state of the sidewalk. Images deemed too blurry or dark by a human were considered inadequate. A sense of the epistemic uncertainty helps analyze whether a given participant would be useful as a campaign continues. For example, participant #1 has a high mean (likely to contribute adequate information) but our confidence in his ability is low since the number of contributions he made is small. The confidence factor, however, for the other participants is high because have contributed much more data. We can consequently be much more certain about the abilities of the other participants. Note that we used an error level of .1 when calculating the epistemic uncertainty.

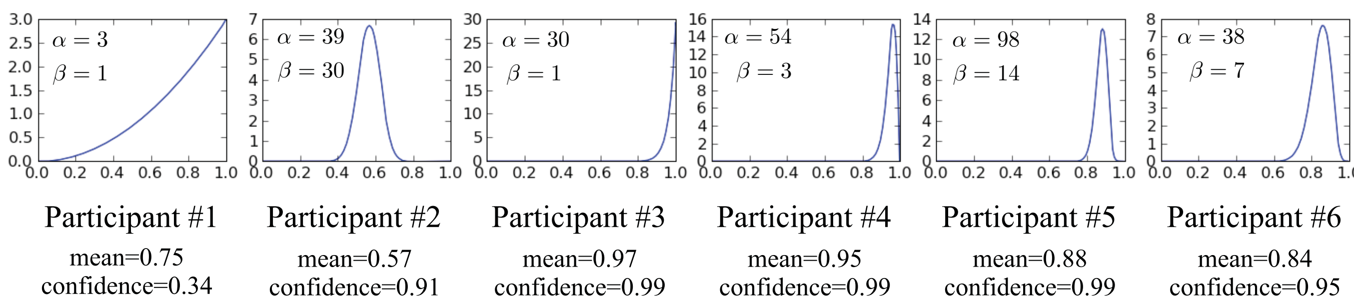

Fig. 2. Walkability pilot analysis showing the likelihood of contributing adequate contributions.

B. PEIR and Long-Term Participation

During the PEIR technical pilot campaign, participants were encouraged to contribute location traces as frequently as possible to test the performance and accuracy of the system. Figure 3 shows the analysis of participation over 61 days for the 26 participants in the pilot. PEIR participants can be clustered into three types of users: consistent, bursty, and sporadic. Consistent contributors contributed data in a dedicated manner throughout the campaign. Bursty participants showed concentrated bursts of contribution, perhaps due to reminders sent to solicit participation. Sporadic participants were very inconsistent or even one-time contributors. By breaking up the campaign into intervals and evaluating participants using the Beta distribution, we could cluster participants according to participation and send feedback based on this information.

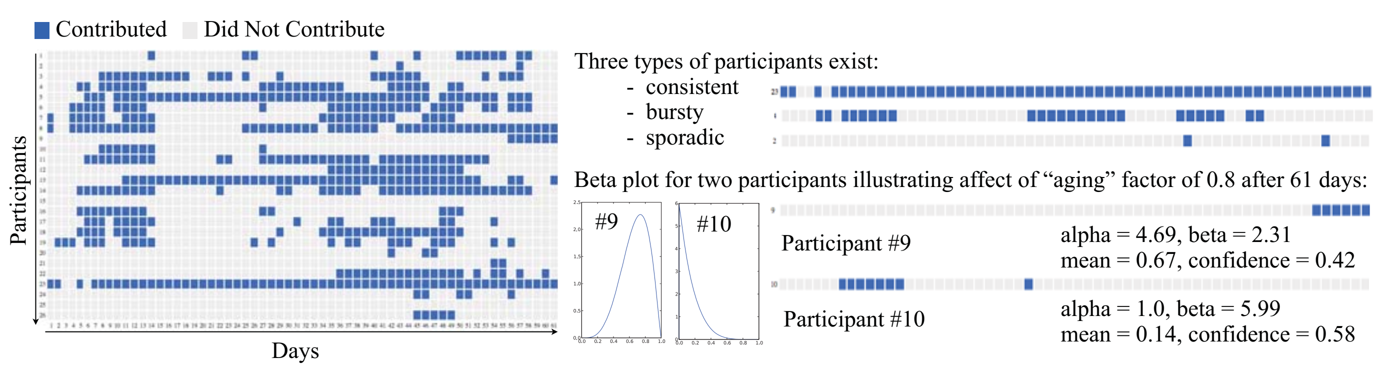

Fig. 3. PEIR analysis showing long-term participant contributions and aging factor.

A feature of the Beta distribution beneficial for long-running campaigns is the aging factor, which we demonstrate by analyzing PEIR participants #9 and #10. Participants #9 and #10 contributed roughly similar amounts of data, but did so during different periods of the campaign. Participant #9 contributed during the tail end of the campaign while participant #10 was heavily involved during the beginning. By having a weight of .8 for the aging factor, we see that participant #9 would be more likely than #10 to contribute data (thus giving #9 a higher mean) in the immediate future. Incorporating an aging factor into the reputation mechanism indicates reliability over time, important since campaigns with a long temporal duration may experience bursty participation.

C. Project Budburst and Reliability

For the Project Budburst pilot we recruited 11 participants. Because we did not have any prior measure of their reputation as data collectors, we initiated a short calibration exercise. The campaign organizer documented flowering plants using geo-tagged images along three specific routes. Participants then traversed the same routes looking for flowering plants. Routes #1 and #2 were short and consisted of 15 and 7 flowering instances, respectively. Route #3 was longer and had 23 instances of flowering. By comparing participant contributions to the calibration phase, we can quickly identify highly reliable participants as well as ones who may required feedback to adjust their data collection practices. For instance, by looking at the results from routes #1 and #2, we could predict which users would be most effective in route #3. The participants who achieved mean of .88 and mean of .92 (with confidence level greater than .90 when the acceptable error is .1) using the Beta distribution for the first two routes, contributed samples that most closely matched the ground truth for route #3. On the other hand, the participants who performed poorly, mean of .58 and .71, on the first two routes followed up with samples that were least consistent with the ground truth. This illustrates that using the Beta distribution with a calibration phase could effectively provide an initial measure of a participant reliability useful for adaptive recruitment or feedback.

V. Discussion and Future Work

Both the challenge and the promise of participatory sensing emerge from involving people in the sensing process. The ultimate "smart sensors," people can make decisions that increase data accuracy, but they also vary according to participation and performance. This paper presents a set of participation and reputation metrics along with a model to help organizers of sensing campaigns determine the reputation and fit of potential participants and whether adjustments are needed during campaign execution. This work is just a first exploration; further study will take place as we incorporate an adaptive participant recruitment system. We will explore how best to communicate reputation ratings, provide actionable feedback to improve a participant's reputation, and determine whether the metrics can become incentives. Another area for future work is exploration of how attributes used in recruitment, such as social network membership or external credentials, may affect reputation measures.

References

[1] J. Burke, et. al., "Participatory Sensing," *ACM Sensys World-Sensor-Web*, 2006.
[2] S. Eisenman, et. al., "MetroSense Project: People-Centric Sensing at Scale," *ACM Sensys World Sensor Web Workshop*, 2006.
[3] E. Paulos, R. Honicky, and E. Goodman, "Sensing Atmosphere," *Workshop on Sensing on Everyday Mobile Phones in Support of Participatory Research*, 2007.
[4] National Phenology Network, "Project budburst," 2008. [Online]. Available: http://www.windows.ucar.edu/
[5] E. Agapie, et. al., "Seeing Our Signals: Combining location traces and web-based models for personal discovery," *HotMobile*, 2008.
[6] S. David, "Toward Participatory Expertise," *Structures of Participation in Digital Culture*, 2007.
[7] P. Resnick, et. al., "Reputation systems," *Communications of the ACM*, vol. 43, no. 12, pp. 45–48, 2000.
[8] Amazon, "Amazon mechanical turk," 2008. [Online]. Available: http://www.mturk.com
[9] GURU.com, "Guru.com - freelancers at online service marketplace." 2008. [Online]. Available: http://www.guru.com
[10] Amazon, "Amazon askville," 2008. [Online]. Available: http://www.askville.com
[11] Yahoo, "Yahoo answers," 2008. [Online]. Available: http://answers.yahoo.com
[12] A. Jøsang, R. Ismail, and C. Boyd, "A survey of trust and reputation systems for online service provision," *Decision Support Systems*, vol. 43, no. 2, pp. 618–644, 2007.
[13] E. Babbie, "The practice of social research," 2007.
[14] A. Schlosser, M. Voss, and L. Bruckner, "Comparing and Evaluating Metrics for Reputation Systems by Simulation," *Paolucci*, 2004.
[15] W. Teacy, J. Patel, N. Jennings, and M. Luck, "Coping with inaccurate reputation sources: experimental analysis of a probabilistic trust model," *Conference on Autonomous Agents and Multiagent Systems*, pp. 997–1004, 2005.
[16] A. Jøsang and R. Ismail, "The beta reputation system," *15th Bled Electronic Commerce Conference*, pp. 324–337, 2002.
[17] S. Ganeriwal, L. Balzano, and M. Srivastava, "Reputation-based framework for high integrity sensor networks," *IEEE TOSN*, 2008.

Participatory Sensing in Commerce: Using Mobile Camera Phones to Track Market Price Dispersion

Nirupama Bulusu[1], Chun Tung Chou[2], Salil Kanhere[2], Yifei Dong[2], Shitiz Sehgal[2], David Sullivan[2] and Lupco Blazeski[2]

[1]Computer Science Department, Portland State University, Portland, OR 97207, USA
Email:{nbulusu}@cs.pdx.edu
[2]School of Computer Science and Engineering, University of New South Wales, Sydney, NSW 2052, Australia
Email:{ctchou, salilk, ydong, shitiz.sehgal, dsul945, lbla241}@cse.unsw.edu.au

Abstract

In economics, price dispersion refers to the price difference of a homogeneous good across different vendors. According to [1] "The empirical evidence suggests that price dispersion in both online and offline markets is sizeable, pervasive, and persistent." Not surprisingly, there exist several popular web commerce sites such as Froogle that enable users to track consumer pricing information in online markets. In this paper, we present and explore our vision that participatory sensing can be employed in this new application domain to track price dispersion in homogeneous consumer goods even in offline markets. We discuss two proof-of-concept participatory mobile camera-phone sensing systems that we have built: (1) automating fuel price collection, and (2) semi-automated scanning of receipts.

I. Introduction

Price dispersion of *homogeneous* goods is a fact of life [1]. We emphasize homogeneity because if two goods are not homogeneous, such as televisions of different brands, then there is a quality difference which makes them hard to compare quantitatively. We have encountered myriad real life examples of price dispersion. For example, the following homogeneous goods were sold at different stores at fairly different prices at the same time in June 2008. We observed a $10 price difference for multivitamins (a $30 product) between Costco and RiteAid stores, and nearly a $200 price difference for HDTVs (a $2000 product) between Circuit City and Best Buy. Online, the quoted air fare for the same flight was $600 higher at Expedia than Lufthansa $2600 at the same instant of time.

Price dispersion is attributed to several causes. A seminal article by Varian [2] suggests that price dispersion might be a deliberate marketing ploy by retailers to entice consumers into exploring their choices. Nevertheless, a major cause is the consumer search cost incurred in collecting pricing information from competing retailers, including the opportunity cost in time in acquiring this information [Baye06]. Price dispersion remains widely prevalent on the Internet (15-17%) [3], although studies have speculated that the low Internet search cost, where alternate retailers are often just a mouse click away, will eliminate price dispersion [4]. Not surprisingly, numerous web commerce sites such as Shopzilla[1] and Amazon[2] try to remedy this situation in online markets by providing a clearinghouse of price information for a homogeneous good for different e-retailers.

There are compelling reasons for creating such a clearinghouse of up-to-date product pricing information, even for offline markets of brick and mortar stores. It could create arbitrage opportunities, wherein an enterprising person can leverage the price difference for profit. The availability of real-time price dispersion information can empower consumers to more effectively negotiate prices [5]. In online markets, studies cited by [1] show that savvy consumers who use on-line price comparison sites save up to 16% in consumer electronics purchases.

Numerous consumer communities are already tracking price dispersion manually. A group of Hong Kong

[1] http://shopzilla.com/
[2] http://www.amazon.com/

housewives divide themselves into teams to manually copy prices of selected staple grocery items in major supermarkets and local grocery stores, and upload the prices to a website, prompting a major Chinese newspaper to advertise weekly grocery prices across different stores on its website[3]. In several countries, petrol price information is collected manually, by volunteers or employees of websites such as gaspricewatch[4] (USA) and motormouth[5] (Australia). Manual price information collection is cumbersome, error-prone and not up-to-date.

Our vision is to apply participatory sensing to share consumer pricing information and reduce the search costs of tracking price dispersion in offline markets. We are motivated by the success of the Wikipedia, Youtube and BitTorrent applications that are driven by altruistic user participation. In this paper, we explore two participatory camera phone sensing systems: (1) automating fuel price collection, and (2) semi-automated scanning of receipts.

II. Related Work

Participatory Sensing enables collection and dissemination of environmental sensory data by ordinary citizens, through devices such as mobile phones, without requiring any pre-installed infrastructure [6]. Researchers have recognized its potential and applied it in many domains, including but not limited to, health (DietSense) [7], intelligent transportation (TrafficSense) [8] and air-quality monitoring [9]; however, to the best of our knowledge, participatory sensing has not been applied in commerce. As in our proof-of-concept systems, DietSense and TrafficSense use camera phones. Researchers are also developing geo-mapped clearinghouses such as SensorMap[6] to simplify sensor data sharing. Our goal is to extend this idea to pricing information collected by image sensors.

The use of mobile phones to enable micro-transactions in commerce has burgeoned over the past few years, particularly in the developing world. It is estimated that Indian farmers get only about 20-25% of the final purchase price of their agricultural produce (about 40-50% for farmers in USA), while most of the rest goes to middlemen. The recently introduced Reuters Market Light services provides farmers with up-to-date information on crop prices and related agricultural news via SMS messages to their mobile phones [10]. Key distinctions between this work and our vision are that we focus on empowering the *consumer* community, and focus not only on modes of disseminating pricing information to users, but also modes of collecting information from consumers. Parikh has used camera phones to scan loan applications for supporting rural microfinance in the CAM system [11]. As in CAM, we use the camera phone to scan receipts, albeit to support participatory data collection.

III. Challenges

Significant challenges remain to be addressed. Data gathered from camera phones is not in a consistent format, making it hard to aggregate pricing information across different retailers. In contrast, aggregation on systems such as Shopzilla is much easier, as they operate over Web-XML data with well defined schemas. Moreover, the sheer numbers of goods and consumers make it difficult to collect information in a single database. Because of this, information clearinghouses such as gaspricewatch tend to focus on a single good. Recent database research, such as COLR-Tree is exploring scalable indexing for SensorMap [12]. The computer vision aspects of extracting price information are also non trivial. Another concern is the optimal positioning of camera phones for image capture, making automation difficult. In our proof-of-concept systems, we exploit the availability of GPS (Global Positioning System) and GIS (Geographical Information System) software to simplify image processing.

Other challenges seem inherent to all participatory sensing systems, whose success hinges on achieving high user participation. How do we promote collection and sharing of pricing information? Two types of incentives are possible here. The first incentive is to lower the technical and monetary barrier for participation. The designer must make the system easy to use, ensuring that it takes minimal effort and very low monetary cost to share data.

[3] http://price.mingpao.com/
[4] http://www.gaspricewatch.com
[5] http://motormouth.com.au/default_nf.aspx
[6] http://atom.research.microsoft.com/sensormap

The system should be automated as far as possible to reduce reluctance to participation. Free text messaging or WiFi capability in some phones could eliminate the monetary cost of sharing data. We have investigated these in PetrolWatch, the fuel price collection application. It is possible to provide the consumer an information reward proportional to her contribution, as has been explored in systems like BitTorrent.

The final concerns are security, privacy and data reliability. How does a user upload pricing information without exposing her shopping behavior? Here, solutions must safeguard not only the user's location privacy but also her shopping pattern privacy. There is a monetary value for knowing what products a consumer buys. If anonymity is not essential, users can be provided the choice to contribute data anonymously or not. It is also critical to ensure the integrity and reliability of the contributed pricing data, ensuring that no bogus data is contributed. This is difficult, because data integrity is at odds with privacy. We expect to build upon the solutions being developed by researchers in this community to address these challenges in the long term.

IV. PROOF OF CONCEPT SYSTEMS

We have built two systems, *PetrolWatch* [13] and *MobiShop* (demonstrated at [14]) that process and deliver product pricing information from street-side shops or gas stations to potential buyers, on their mobile camera phones, and have similar client-server architectures (see Fig.1). They can also serve as an effective indirect advertising medium for gas stations or shops. They operate in two modes: (i) price collection and (ii) user query.

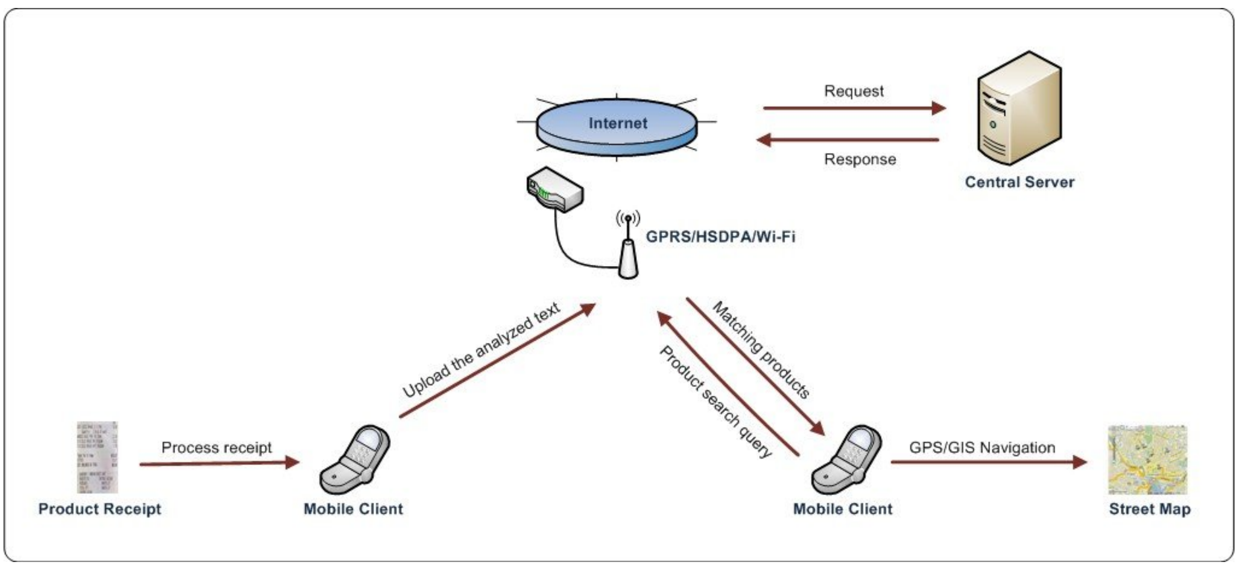

Figure 1: Mobishop System Architecture

A. *PetrolWatch: Automated Fuel Price Collection*

The goal of PetrolWatch (see Fig. 3) is to automate collection of fuel prices, by triggering the mobile phones of contributing users to photograph the roadside fuel price boards when they approach service stations. A central server implements computer vision algorithms for processing images and extracting fuel prices. To deal with a non-structured environment, and to reduce computer vision complexity, it relies on the GIS database and GPS location to know the service station brand and uses the fact that each brand uses a specific color for its fuel price board. The meta-data (location coordinates, service station brand and time) are extracted and stored separately.

The images and fuel brand information are passed on to the image processing engine. The first step detects the existence of a fuel price board. For each service station brand, we employ a tailored color thresholding that can capture regions within the images, having a color scheme similar to the fuel brand price board. In certain situations, surrounding objects in the image may have colors resembling the board, e.g.: the blue sky may be similar to the Mobil fuel price board. In this case, we use post-processing techniques to narrow the search. We use the price board dimensions to exclude some of the candidate regions selected by color thresholding. This is further refined by comparing the color histogram of all candidate regions with that of a sample fuel price board image. The detection concludes by identifying the precise board location in the image. The image is cropped to contain only the board, and normalized to standard size and resolution. We convert the color image to binary and use connected component labeling to extract the individual numeral characters. A Feedforward Backpropagation Neural Network algorithm is used to classify the price numeral characters. The extracted prices are stored in a database, linked to a GIS road network database populated with service station locations, consistent with the GIS database installed on the phones. The server updates fuel prices of the appropriate station in the database if the current image has a newer timestamp. The past station fuel price history is also recorded to analyze pricing trends.

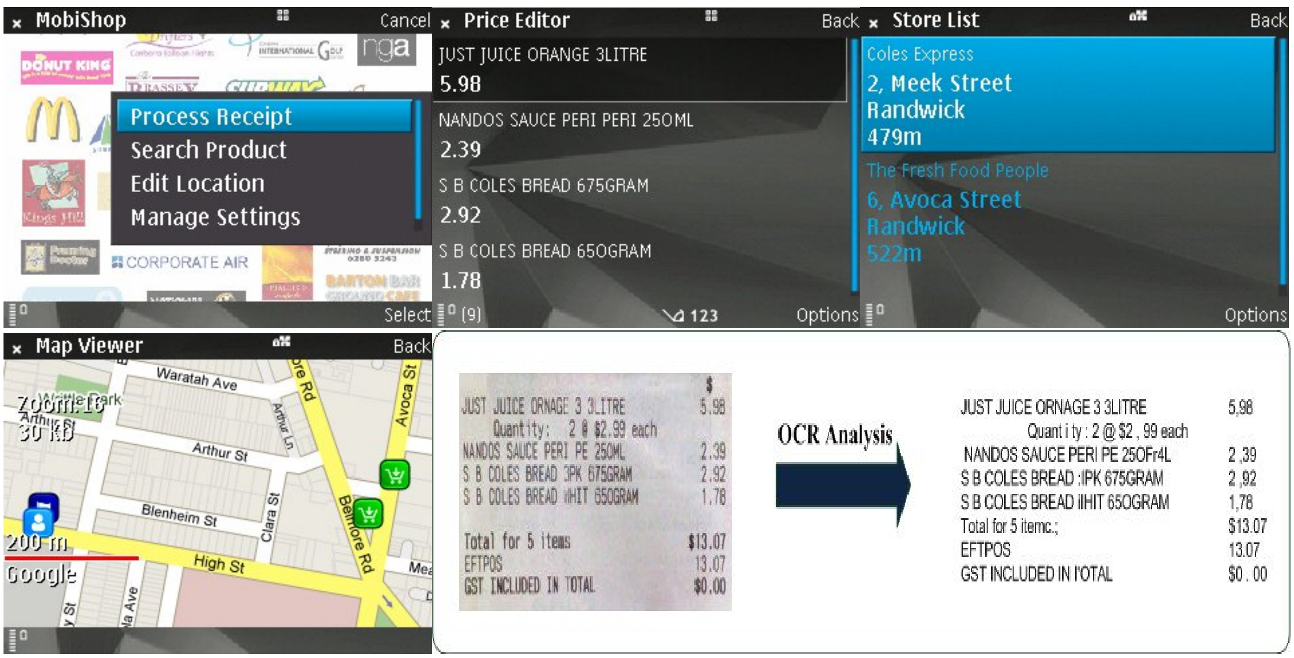

Figure 2: Mobishop screen shots and OCR

Figure 3: PetrolWatch screen shots.

B. *Mobishop: Semi-automated Scanning of Receipts*

To contribute pricing information in Mobishop (see Fig. 2), the user photographs the store receipt with his camera phone (more efficient than photographing product tags), which lists the products and their prices. MobiShop implements Optical Character Recognition (OCR) on the mobile device to extract the pricing information from the image. The user is given an option to edit the extracted text to fix any mistakes, and also to allow her to delete personal information such as credit card details. The products and prices are uploaded with the GPS coordinates of the user and the time of purchase to the central server using a TCP connection over built-in GSM/GPRS/3G/HSDPA or 802.11 interfaces. The server collates user inputs and maintains an updated repository of product prices at different stores. This database is interfaced to a GIS street map populated with store locations.

The MobiShop client has been primarily implemented in Java ME to ensure portability across devices on a Nokia N95 8GB phone. We used the native Symbian OS 9.2 OCR engine, which can accurately detect about 60% of item prices on the receipt. A simple GUI is provided for user input. User location is determined by querying the GPS receiver. The client interfaces with an external GIS library, J2MEMap[7] so store locations can be highlighted on a street map for navigation. The server program is written in Java and is executed as a daemon on an always-on workstation. In future work, we intend to improve the OCR accuracy, and enable self-registration of stores.

V. CONCLUSION

We explored participatory camera phone sensing for tracking price dispersion in offline markets in two systems – *PetrolWatch* and *Mobishop*. They address the challenge of collecting offline non-structured information. This system model could be extended to help users keep track of their shopping habits, in receiving frequent-buyer promotions, and track price dispersion elsewhere, like rates for various city parking structures.

ACKNOWLEDGMENT

This research is funded by the Australian Research Council Discovery Grant DP0770523 and National Science Foundation CAREER Grant 0747442.

REFERENCES

[1] M. R. Baye, J. Morgan and P. Scholten, "Information, Search, and Price Dispersion," *Handbook of Economics and Information Systems* (T. Hendershott, ed.), Elsevier Press, Amsterdam, 2006.

[2] H. R. Varian, "A model of sales", *American Economic Review*, vol. 70, pp. 651-659, 1980.

[3] M. Baye, J. Morgan, P. Scholten, "The value of information in an online consumer electronics market", *Journal of Public Policy and Marketings*, vol. 3, pp. 481–507, 2003.

[4] T. Gupta and A. Qasem, "Reduction of Price Dispersion through Semantic E-commerce: A position paper," in *International Workshop on the Semantic Web*, Hawaii, USA, 2002.

[5] B. Tedeschi, "It's a Buyer's Market. Haggle", *Money Magazine*, August 2008.

[6] J. A. Burke, D. Estrin, M. Hansen, A. Parker, N. Ramanathan, S. Reddy, M. B. Srivastava, "Participatory sensing", In *Proceedings of the Workshop on World Sensor Web*, Denver, USA, 2006.

[7] S. Reddy, A. Parker, J. Hyman, J. Burke, D. Estrin, and M. Hansen, "Image Browsing, Processing, and Clustering for Participatory Sensing: Lessons From a DietSense Prototype", in *Embedded Networked Sensors*, Ireland, June 2007.

[8] P. Mohan, V. Padmanabhan, and R. Ramjee, "TrafficSense: Rich Monitoring of Road and Traffic Conditions using Mobile Smartphones", *Microsoft Research Technical Report, MSR-TR-2008-59*, Apr 2008.

[9] E. Paulos, R. Honicky, E. Goodman, "Sensing atmosphere", In: *Workshop on Sensing on Everyday Mobile Phones in Support of Participatory Research*, Sydney, Australia, Nov 2007.

[10] D. Grammaticas, "Text messages empower poor farmers", BBC News, 6 May 2008. http://news.bbc.co.uk/2/hi/south_asia/7385542.stm

[11] T. S. Parikh, "Using Mobile Phones for Secure, Distributed Document Processing in the Developing World", *IEEE Pervasive Computing Magazine*, April 2005.

[12] Y. Ahmad and S. Nath, "COLR-Tree: Communication Efficient Spatio-Temporal Index for a Sensor Data Web Portal", In *Proceedings of ICDE*, Cancun, Mexico, 2008.

[13] Y.F. Dong, S. Kanhere, C.T. Chou, N. Bulusu, "Automatic Collection of Fuel Prices from a Network of Mobile Cameras", In *Proceedings of the Distributed Computing in Sensor Systems*, Santorini, Greece, June 2008.

[14] S. Sehgal, S. S. Kanhere, C. T. Chou, "MobiShop: Using Mobile Phones for Sharing Consumer Pricing Information", In *Demo Session of the Intl. Conference on Distributed Computing in Sensor Systems*, Santorini, Greece, June 2008.

[7] http://j2memap.landspurg.net/

Participatory Sensing for Urban Communities

D. Airantzis[1], A. Angus[2], G. Lane[2], K. Martin[2], G. Roussos[1], and J. Taylor[1]

[1]School of Computer Science and Information Systems,
Birkbeck College, University of London, Malter Str, London WC1E 7HX, UK
Email:{dimitrios, gr, jenson}@dcs.bbk.ac.uk
[2]Proboscis, 101 Turnmill Street, London EC1M 5QP, UK
Email:{alice, giles, karen}@ proboscis.org.uk

Abstract

Social Tapestries views participatory sensing as the principal supporting technology to enable grass-roots groups and communities to track and act on information about their local environment. In this article, we report on our experiences with a low-cost open source hardware and software platform, which we specifically developed for this task. We describe how we employed this platform to support community workshops and art, and highlight the lessons learnt through our involvement with urban communities in London, UK. We conclude by identifying the main ingredients for the development of a successful strategy for the use of this and other similar platforms in supporting environmental sustainability through sustainable communities.

I. Introduction

Public authoring is the mapping and sharing of local knowledge using pervasive computing technology to create and support relationships beyond established social and cultural boundaries and the development of new practices around place, identity and community. Social Tapestries (ST) is a programme of research into the potential costs and benefits of public authoring to communities and individuals [1, 5]. Within ST, several projects have focused specifically on how public authoring can support grassroots participatory sensing activities with a view to allow local urban communities to take action towards environmental sustainability. Here, we present three such projects, namely Feral Robotic Public Authoring, Snout and Everyday Archaeology, which explore complementary ways to achieve this aim. All projects share a commitment to the principal premise of ST that is, to contribute to an alternative experience commons where people are presented with the opportunity to be agents, actors and authors. They also make use the same open-source hardware and software participatory-sensing platform, which we specifically developed to support these activities. We describe the development of the platform in parallel to our discussion of our developing understanding of public authoring in this context.

II. Public Participation and Social Tapestries

Public authoring takes the view that true daily life is richer and more complex than consumption of services, products and media, relying as much on social networks, personal experiences and chance interactions and connections. Pervasive computing applications should attempt to reflect this richness and complexity. At the core of such diverse everyday activities lies social knowledge, a term used in ST to refer to the passing communications that are the glue of society and communities: the everyday and essential sharing of information, stories, knowledge and memories with friends, family, neighbours and strangers.

The practice of public authoring can offer opportunities to individuals and groups to intervene in situations that have previously been tightly controlled. For example, in the ST Eyes on the Street project, residents of the Havelock Estate in Ealing, London, are engaged in public authoring with a view to employ local knowledge to support the operation of a tenant organisation that aims to take over the management of their estate. In this case, we have found that public authoring may provide insights not otherwise available by creating a record of living that far exceeds what is possible through centralized estate-management services. Such activities should not necessarily be seen as threats to established authoritative sources of knowledge but rather as people's desire to participate.

This desire for participation in public life stemming from the grassroots is particularly attractive to ST. As a result, following on Eyes on the Street and a subsequent project with St Marks Housing Cooperative in West London, it became evident that an investigation was required into how participation by local communities in data collection through sensing in their immediate environment, can enable them to become actively involved in promoting sustainable environmental practices. Note that he concept of participatory sensing has been used elsewhere as a model for sensor networks [2]. In our work, we use the term giving emphasis on its social, rather that its technological meaning.

Many of the explorations within Social Tapestries are supported by the Urban Tapestries software platform [1], which has been specifically designed to enable public authoring. In UT, users as authors go about their everyday activities as they usually would, but whenever they wish to add new content they do so using their mobile phone. This task is facilitated by the UT client that allows them to annotate a place with media including text, sound, images or video. Authors can link such pockets of content together into threads with a specific theme. Threads and pockets published on UT weave together into an information tapestry overlaid on the urban structure. Users as consumers can search for, browse and access content published by other participants, also using the UT client application on their mobile phone. For our investigation of participatory sensing for environmental sustainability we extended UT with new information harvesting, management and visualization capabilities, and developed a new client platform specifically for sensing.

III. PROTOTYPE ONE: ROBOTIC FERAL PUBLIC AUTHORING

Natalie Jeremijenko in her Feral Robotic Dogs project proposed ways to reconfigure toy robot-pets that became popular in the early 2000s, with a variety of low-cost chemical sensors that can sense, record and in some cases trace environmental pollution [3]. The aim of doing so is to create an opportunity for public discourse by providing the tools to construct open-ended interpretations of the evidence at hand. Such experiments open up new possibilities for exploring local environments to detect the presence of many kinds of emissions and map them using the UT toolset. A large variety of low-cost sensors are readily available including carbon monoxide and dioxide, solvent vapours, electro-magnetic emissions (for example, those coming from mobile phone masts, electricity generators and so forth), and light and noise pollution. Adding the sensor readings to UT makes evident the relationships between the physical environment and communal places. It enables people to feel they can learn about their environment and have the evidence to do something about it. By linking robot building and mapping workshops into traditional community events for example, fetes and local festivals, a wide range of people can become involved in gathering and sharing knowledge about their environment.

The Feral Robotic Public Authoring (RFPA) project takes exactly this point of view. To achieve this goal, we augmented the initial Feral Dog design with wireless networking, location and advanced environmental sensing capability, and linking it to the internet and the UT platform. Our priority is to develop a low-cost platform that can be built easily out of widely available commodity components and with very limited technological resources and skills. Out intention is that all software and designs would become available on the web for everyone to freely re-use, build or modify.

Our design implemented Gumstix, an open-source small and inexpensive computer running the Linux operating system. We extended the core design with cooling and power management sub-systems, which were critical for the consistent operation of the device, an external GPS receiver connected over Bluetooth, and customised its Analog-to-Digital Converter for use with a radio controlled car. And of course, we build in sensing capability using a variety of inexpensive sensors typically used in fire and carbon monoxide alarms, and home and car ventilation systems [6]. The full designs including bill of materials, assembly instructions for the enclosure and mounting, Gerber files for production of the printed circuit board, sensor calibration procedure and the software repository are available online via http://socialtapestries.net/feralrobots/

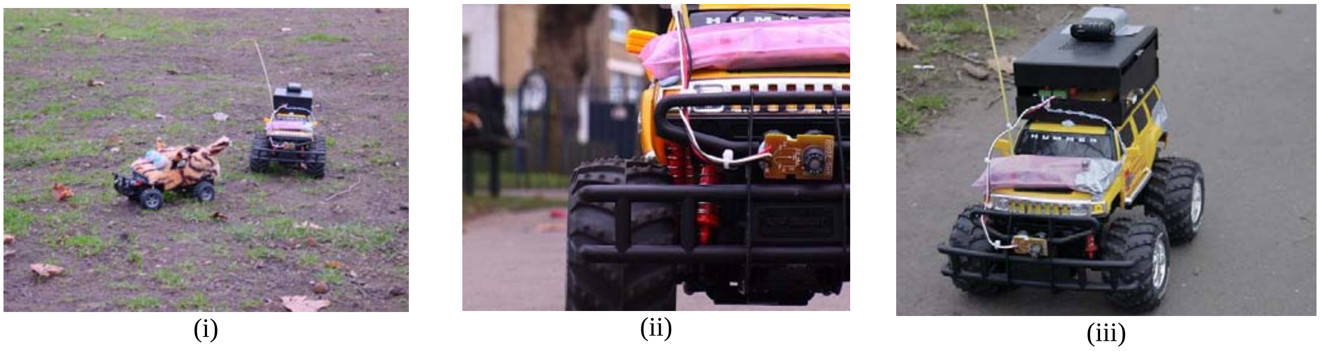

Fig. 1 (i.) The original Feral Robot (left) followed by the networked and location aware version developed for RFPA . (ii.) A close-up of the sensor modules at the front of the RFPA remotely controlled car. (iii.) The complete RFPA design with sensors, GPS, processing and networking unit assembled and operational during the London Fields trial.

IV. COMMUNITY MAPPING WORKSHOPS WITH RFPA

London Fields is a popular park in Hackney, East London, and an important resource for local communities in a built up area. The park is used by local people for a variety of activities: as a space to play and socialise in, for championship cricket and football games, dog walking, and as a popular walking and cycle route. In its relatively long history, London Fields and the area around it have adapted to accommodate the differing needs of the surrounding population.

Air quality in London is monitored on an hourly basis by the London Air Quality Network (LAQN), through an extensive network of observation stations in fixed locations across its Boroughs. LAQN is an important resource but considering that Hackney itself only has one station for the entire borough there is clearly ample opportunity to examine air quality at a more localized resolution. Yet, the collection of this kind of data by non-experts is not necessarily useful, and some would argue that such activities would lack scientific rigour and would thus not be comprehensive or authoritative. Data collected through the RFPA devices can provide a snapshot of pollution in a specific place at a specific time and is not designed to replace or replicate LAQN. Instead, it aims to trigger an open dialogue about how pollutant sensing technology placed at a grassroots level can function and its potential applications for community action and interaction.

Community pollution mapping workshops were organised in collaboration with SPACE Media Arts, a local arts and education charity, which allowed us to access their local community networks. We found is that grassroots pollution mapping is not necessarily about producing accurate scientific data. Instead, it is a tool to highlight concerns, to map knowledge, to enable involvement in the data collection process thus reinforcing perceptions of the area, and provide the focus for communities to come together. As one workshop participant remarked: “we have come to accept air pollution because we are culturally habituated in it -- that’s got to change and if this doesn’t happen at a grassroots level with tools that we can handle ourselves, governments will not shift because they are in with the big corporations.”

Nevertheless, not all workshop participants took the same view, and others expressed the opinion that ordinary people do not have any control over their local environment. For example, vehicle emissions are the major cause of air pollutants in London and in many cases they are due to pass-through traffic, about which local people have little power to intervene. This point of view can lead to passivity and resigned acceptance of the situation as expressed by one participant who said that “the more I think about it, the less I want to have any access to any data about air pollution in my locality, or information about this park. I don’t have a garden, I have a kid, and I’ll always use it.”

In addition to views of the here and now, community mapping workshops prompted participants to reminisce about the history of the Fields, highlighting past activities in the area which could have left an environmental footprint. This type of local knowledge is invaluable and can help locate pollution hotspots that would otherwise require an extensive survey. An expert coming from the outside would not have access to this knowledge without considerable resources for research.

A second series of RFPA workshops was run in collaboration with the Jenny Hammond School in Waltham Forest, also in East London. In July 2006, a week-long workshop with 30 students aged between nine and 10 years old involved several activities including extensive use of the RFPA platform to gather evidence about the world around them. These activities were linked to specific modules within the Key Stage 2 national curriculum, in particular transport, architecture and climate, and allowed the students to gradually develop associations and connections between these areas and how they all fit within the environmental sustainability agenda.

V. Developing the platform

RFPA allowed us to experiment with participatory sensing and gain experience in what works both from a technical and a community perspective:

Embedded interaction. The RFPA cars, although appealing (especially to younger and male users) provided limited opportunities for interaction beside the remote control. While running, they collect data silently and relay it to the UT server for further processing without any perceptible indication that this has occurred. Their operators are only able to tell if the different system components function correctly due to several LED indicators, but are unable to get feedback about the current detected levels of air pollution for example.

Media scavenging. UT was designed as a stand-alone software system well before the emergence of Web 2.0 and due to our resource limitations in supporting the software development process over an extended period of time, it lacked features that users identified as important. As a result, we adopted an approach based on scavenging functionality that could be mashed up with UT to provide the missing features, such as social networking as implemented by Ning.com.

Everyday archaeology. Since we formulated the participatory sensing approach, we have seen its emphasis shift from pollution mapping to what we now describe as everyday archaeology. In doing this, we shifted our focus away from the specifics of data collection and focus on the process of excavating information about the local environment and its relationship to communities.

VI. Prototype Two: Snout

Our next investigation of participatory sensing was through the development of a community art project in partnership with the International Institute of Visual Arts (inIVA). We consider community art to be an appropriate approach for community development because it reflects the main principles of participatory sensing: it is rooted in a shared sense of place, tradition and spirit; it is as much about the process of involving people in the making of the work as the finished object itself; and it is situated in public, accessible and resonant places, geared to a specific audience and a specific time. It seems that this point of view can offer new opportunities for the development of a community around environmental sustainability by providing both the practical and the conceptual framework required. Indeed, community art and grassroots activism are about knowledge and building social capital in the form of the grassroots networks that enable people to move information and ideas to a broader audience and make change happen.

The specific performance developed for Snout explored relationships between the body, community and the environment. The concept of the performance was developed around the Carnival with two typical costumes build and instrumented with participatory sensing capabilities. The reason for this choice was that carnival is a time of suspension of the normal activities of everyday life – a time when social hierarchies are inverted and when everyone is equal. This view is highly compatible with the Snout objective to invite participation. The characters selected were Mr Punch and the Plague Doctor (Figure 2).

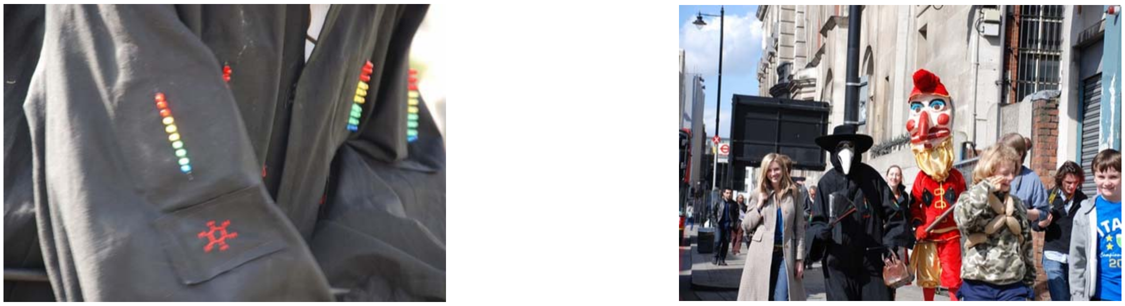

Fig. 2 Snapshots from the Snout performance.

VII. Lessons learnt through Snout

Snout enabled us to take further our ideas about participatory sensing, especially in identifying effective ways to facilitate the development of grassroots communities.

Inspiration rather than prescription. We initially considered our open-source platform as the main ingredient of any community project around participatory sensing. However, it turned out that the complexity of the hardware construction aspect in particular caused significant difficulties to the general user, despite the fact that is considerably more accessible than any other alternative available.

Multiple ways to access information. A single web interface as the only means to interact with the captured data appeared to be far too limiting and unable to address the needs and concerns of all the users involved. Instead, it was necessary to provide alternatives that could address the specifics of the situation in which access to the data is required as well as the skills of the particular user.

Access also via low-tech materials. Within the context of multiple modalities of access to the captured data, it is especially important to note approaches that do not employing information or communication technology. For example, fabrication of physical artefacts can be particularly effective in interpreting and communicating the data. We specifically experimented with Story Cube and Diffusion electronic notebooks.

VIII. Conclusions

We believe that this work has demonstrated that it is possible, using cheap electronics and publicly accessible mapping software, to create an engaging form of environmental sensing at a micro-local level. Although our prototypes require a level of electronics and engineering skill above that of most people, they are well within the realm of the hobbyist and we believe it is possible to further reduce its complexity as new platforms and products become more readily available and cheaper. We also hope that we have shown how artists and engineers can collaborate to bridge the gulf between pragmatic technical solutions to social problems and the cultural interventions that artists bring to their communities.

References

[1] A. Angus *et al*, 2008, Urban Social Tapestries, *IEEE Perv. Comp.*, special issue on user-generated content, Sep-Dec.
[2] J. Burke, *et al*, 2006, Participatory Sensing, *World Sensor Web Workshop* at Sensys 2006.
[3] D. Cuff, M., Hansen and J. Kang, 2006, Urban sensing: out of the woods, *Com. ACM*, 51(3):24–33.
[4] A. Galloway, 2004, Intimations of Everyday Life: Ubiquitous Computing and the City, *Cult. Stud.* 18(2-3), pp. 383-407.
[5] G. Lane, 2003, Urban Tapestries, *Per. Ubiq. Computing*, 7(3-4), pp. 69-175.
[6] G. Lane *et.al*, 2006, Community-based Public Authoring with Chemical Sensor Networks, *IEE Intell. Env.*, 2:23 - 29.
[7] P. Paulos and T. Jenkins, 2006, Objects of Wonderment: Hullabaloo, Demo at Ubicomp 2006, Orange County, CA.

Measuring the Pulse of the City through Shared Bicycle Programs

Jon Froehlich[1], Joachim Neumann[2], and Nuria Oliver[2]

[1]Department of Computer Science and Engineering, University of Washington, Seattle, WA 98195, USA
Email: jfroehli@cs.washington.edu
[2]Telefonica Research, Barcelona, Spain
Email:{joachim, nuria}@tid.es

Abstract

City-wide urban infrastructures are increasingly reliant on networked technology to improve and expand their services. As a side effect of this digitalization, large amounts of data can be sensed and analyzed to uncover patterns of human behavior. In this paper, we sense and analyze data from a new type of urban infrastructure called shared bicycling. We provide a spatio-temporal analysis of six weeks of usage data from Barcelona's shared bicycling system called Bicing. We show how these digital traces can be used to uncover daily routines, cultural influences and the role of time and space in city dynamics.

I. INTRODUCTION

We are nearing a time when even the most mundane objects and services will be digitized. As a result of this computational pervasiveness, our interactions in the physical world are increasingly leaving behind *digital footprints*. Recent work has shown the value in sensing these footprints to uncover new insights into human behavior [2], urban dynamics [3], and tourist movements [1]. In this paper, we explore the underlying "pulse of the city" of Barcelona through the lens of a new type of urban infrastructure: a 3rd generation shared bicycling program called *Bicing*. We emphasize not just what the data reveals about the patterns of human movement but also how these patterns reflect culture and the overall spatial context of the city. Our objective is twofold: (1) to highlight the potential of using shared bicycling as a new data source to gain insights into city dynamics and human behavior and (2) to introduce preliminary analysis techniques that we are developing to analyze the ever-expanding amounts of spatio-temporal data produced by urban infrastructures.

II. BICING: BARCELONA'S SHARED BICYCLING PROGRAM

Community shared bicycle programs offer an environmentally friendly, healthy and inexpensive alternative to automobile transportation. Recent technological advances have lead to a new generation of systems, which utilize technology such as RFID, mobile data services, and automated micro-payments to increase operational efficiency and reduce bicycle theft and vandalism. Barcelona's shared bicycle program, Bicing, was launched in March of 2007 (Figure 1). It is one of a series of extremely successful 3rd generation shared bicycle programs recently deployed in Europe. In late Summer 2008, Bicing grew to 373 stations with an average of 25.5 bicycle parking slots, 6,000 bicycles, and over 150,000 subscribers.

To check-out a bicycle, users swipe their RFID membership card at a Bicing station's kiosk, which then displays, via an LCD screen, the bicycle on the rack that has been unlocked. This information is uploaded to a server that keeps track of who has checked out what bicycle and updates the Bicing website with real-time information about the number of available bicycles and vacant slots at each station. A check-out provides 30 minutes of free ride time, every half-hour beyond that costs €0.30 for up to two hours. Bicycles can be returned to any station, where they are placed in an auto-locking rack. Warnings, monetary penalties (€3/hr), and eventually suspension of membership are possible if a user consistently returns a bicycle beyond the two hour limit. Bicing is open from 5AM to midnight on Sunday through Thursday and 24 hours during the weekend. To maintain even distribution of the bicycles (load balancing), a small number of trucks equipped with trailers move bicycles around the city.

Figure 1. (a) A nearly full Bicing station; (b) A station kiosk; (c) A close-up of a locked bicycle; (d) A map of Barcelona showing the location of the 373 Bicing stations. The five highlighted stations are discussed below.

A. Our Bicing Dataset

The Bicing website reports the status of all bicycle stations via a Google Maps visualization[1]. We scrape this webpage every five minutes and extract three data elements per station: the station's *geo-location*, the number of *available bicycles* and the number of *vacant parking slots*. We do not obtain personally identifiable information. Our data logger automatically accounts for new Bicing stations as they appear online. About 1% of the raw data extracted from the Bicing website suffers from one of the following problems: (1) the numbers are unreasonably large; (2) the numbers jump by more than five bicycles and then return to their previous values at the next observation; (3) sometimes all stations simultaneously report zero available bicycles and/or zero vacant parking slots. We clean the data simply by replacing the erroneous value with the most recent valid value. Our cleaned dataset includes six weeks[2] of station observations starting on May 23rd and ending on July 3rd, 2008 for a total of over 4.3 million data points. Note that 2 of the 373 stations consistently reported invalid data and were thus disregarded from our analysis.

In the following section, we analyze the temporal and spatial patterns of the Bicing dataset, in order to explore the underlying human behaviors and movement dynamics in the city of Barcelona.

[1] http://www.bicing.com/localizaciones/localizaciones.php
[2] Although our web scraper has been logging continuously since May 23rd, the Bicing website was down from July 3rd - 30th.

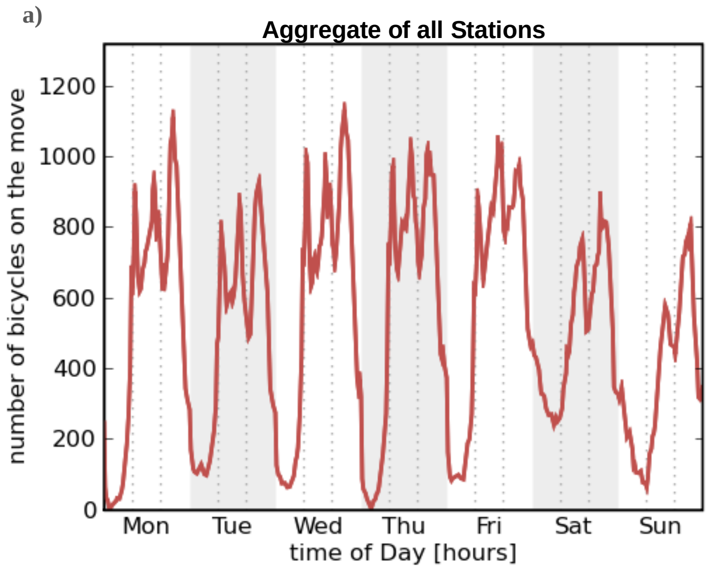

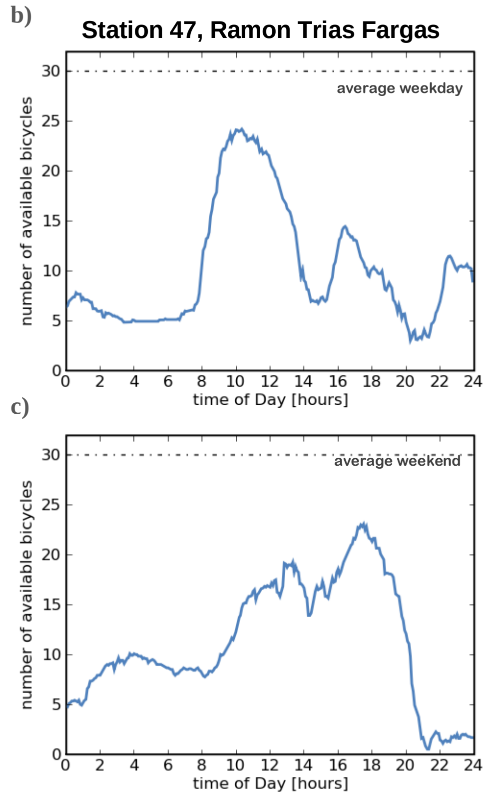

Figure 2. (a) The total number of bicycles checked-out from all stations over a week, averaged across the observation period. (b) The average number of available bicycles at Station 47, Ramon Trias Fargas, averaged across all weekdays (Mon-Fri) in our dataset (c) Same as previous but over the weekend (Sat-Sun). The dashed line in (b) and (c) indicates the number of parking slots at Station 47.

III. ANALYZING THE "PULSE OF THE CITY"

A. Temporal Patterns: Sensing Culture and Daily Routines

The temporal patterns of a city are a reflection of the daily routines of its citizens. Figure 2a shows the average number of bicycles "on the move" during a given point along the week. Perhaps the most salient feature of this graph is the repeating three-pronged spike, which corresponds to the morning, lunch, and evening commutes. As one might expect, the morning commute is absent in the two weekend days, resulting in a two-pronged spike. In addition, observe that the "lunch spike" occurs at 2PM, reflecting that Spaniards tend to eat a late lunch. In addition, the two most popular Bicing periods—*i.e.*, the periods with the largest number of bicycles on the move—are Monday and Wednesday night at around 10PM and that, on average, people tend to use Bicing more during the work week than they do during the weekend.[3]

To further highlight the contrast between weekday and weekend activity, Figure 2b and c portray the number of available bicycles at station 47, Ramon Trias Fargas, which is situated next to the University of Pompeu Fabra (see Figure 1d). Early on weekday mornings there are relatively few bicycles at the station. Then, at around 8AM, students, staff and faculty begin arriving on campus and the number of available bicycles increases quickly as people begin dropping them off. A local minimum occurs at 2PM as people leave for lunch and a second dip occurs around 7-8PM as people seem to be leaving for the night. In contrast to the weekday activity, on the weekends there is no sign of the 8AM commute. Instead, bicycles slowly trickle in throughout the day. Interestingly, however, both the weekday local minima and the weekend local minima seem to temporally align (at around 2PM and 8PM, respectively). We are not certain if this is a reflection of the lunch and dinner routines of these Bicing users or perhaps an artifact of the load-balancing implemented in Bicing via trucks. We are currently exploring methods to automatically detect the presence of the trucks in our data. However, our intuition is that they do not significantly bias our analysis.

[3] A keen reader may observe that the number of bicycles on the move does not drop to zero on weekday nights when the system is closed. We believe this an artifact of Bicing's operational logistics (bicycles are taken off racks during this time for maintenance) and not a reflection of human behavior.

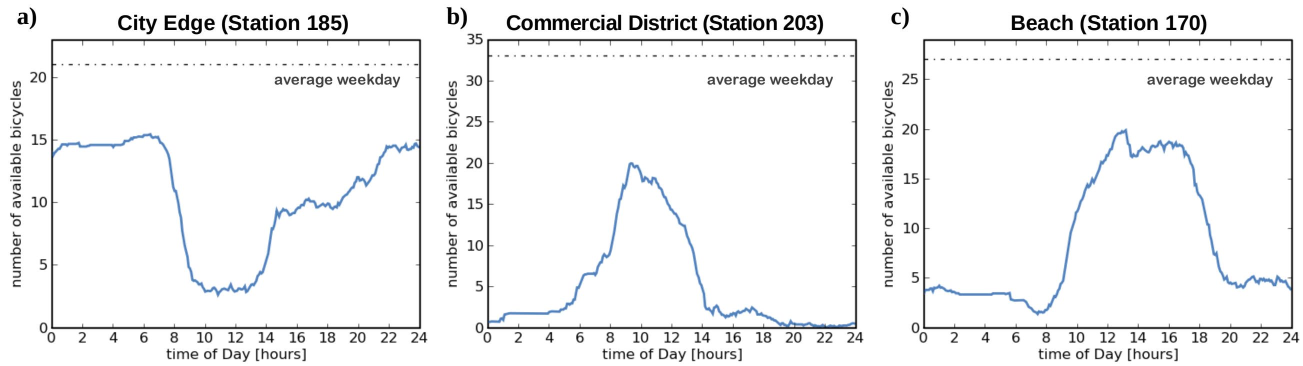

Figure 4. Average number of available bicycles at stations (a) 185 Olzinelles, (b) 203 Diagonal, and (c) 170 Litoral.

B. Spatial Patterns: Sensing the Flow of the City

The spatial layout of a city has an obvious influence on the movement patterns and social behaviors found therein. Barcelona has a mixture of residential, commercial, and recreational areas connected via narrow streets, one-way avenues and a multitude of public transportation options. In Figure 4, we explore the interrelation between a station's location and its underlying temporal usage pattern. What is interesting here is not just that the temporal patterns differ from place to place but what these differences seem to reveal about the *type* of place. For example, station 185 (Figure 4a) is on the city edge and station 203 (Figure 4b) is in a commercial district along Avinguda Diagonal, a major arterial road (these stations are marked in Figure 1d). Their weekday patterns are near opposites: on the city edge, bicycles are checked out starting around 7AM just as bicycles begin arriving in the commercial district. These usage patterns are an indication of the surrounding locale, e.g., the city edge is more residential: people take bicycles in the morning and return them in the evening as they commute to and from work.

Although not as pronounced, the subtle differences between station 203 and station 170 (Figure 4c) reflect a fundamental difference in the reasons why people travel. Whereas stations 185 and 203 are "commuter stations," station 170 is located at the beach. Thus, the 7AM rush of activity does not occur and the rise in incoming bicycles is less pronounced than in the commuter stations, as people casually arrive without the pressures of an explicit work schedule. Although the weekend figures are not shown above, we also observed that while the city edge and commercial stations have considerably different weekday patterns from weekend patterns, the beach station's usage pattern remains relatively constant throughout the whole week.

The "one hump" pattern depicted in both Figure 4b and Figure 4c is one of approximately eight common temporal patterns in our dataset. It is the simplest representation of an "incoming flow" station, often positioned in an area that attracts people during the day. In the next section we report a preliminary analysis that begins to tease out how different types of usage patterns relate across stations.

C. Understanding the Role of Time and Space

In previous sections, we observed the influence of daily routines, culture and location on a station's usage pattern. The usage of Bicing involves a multitude of underlying motivating factors such as commuting, shopping and going to eat. Clustering allows us to measure the relative importance of these effects. By placing the clustering results on a map, we can begin to see the interrelationship between activity and space. This allows us to explore questions like "Do co-located stations share the same usage pattern?" and "How are these usage patterns distributed in the city?"

For clustering, we used the Expectation-Maximization (EM) algorithm found in the Weka toolkit [4]. To compute our clustering features, we split each weekday into seven time bins: early morning (7-9AM), mid-morning (9AM-1PM), lunch (1-3PM), afternoon (3PM-5PM), early evening (5PM-8PM), late evening (8PM-12AM) and night (12AM-5AM). In addition, we considered the entire 24 hour period and only the "open" period (from 5AM to 12AM). For each time bin, we calculated the average number of available bicycles, the difference between the number of bicycles at the beginning and end of the bin's edge, and a measure of the station's activity defined as the percentage of change in the features with respect to the features observed 5 minutes earlier (i.e. the

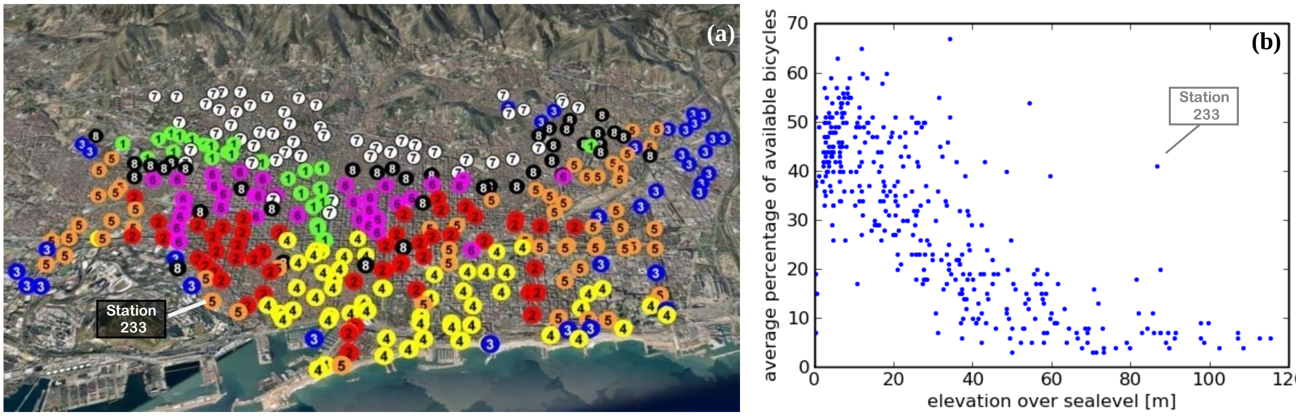

Figure 5. (a) The clustering results; (b) A scatter plot of station elevation vs. average number of available bicycles.

previous observation). Note that EM requires an expected number of clusters as its input. We implemented a visual analysis tool to inspect the temporal patterns of each station within their geo-spatial context. With this tool, we determined eight as the ideal number of clusters.

Figure 5a displays the clusters that contain stations with similar temporal patterns. Close stations tend to behave similarly. However, this is not *always* the case. Cluster 3 (the blue cluster) is clearly the “city edge” cluster; note the number of blue nodes that surround the east, west and southern edges of the Bicing system. Cluster 1 (green) is a “commercial district” cluster, tracing the outline of the northern part of Avinguda Diagonal. Cluster 7 (white) is scattered in the upper part of the map covering nearly the full extent of the city. This cluster represents the highest elevation in the Bicing system and is likely a consequence of a preference to bicycle downhill as Barcelona rests on a slight, ever-increasing slope. Figure 5b reinforces this dynamic by showing how stations with high elevation very rarely have available bicycles. Station 233, marked above, is an interesting exception. Although it is located 86 meters above sea level, its usage pattern is distinct from Cluster 7. This shows that although elevation is a dominant factor affecting usage, other factors may overcome this dominance. In this case, Station 233 is close to both a major park called Montjuïc and residential housing.

IV. Conclusion and Future Work

In this paper, we have introduced the notion of using shared bicycling’s digital footprints to gain an understanding of human behavior and city dynamics. Although our dataset does not contain details about individual movement, our results show how Bicing footprints expose the underlying daily routines and patterns of Barcelona’s citizens. We are currently working on incorporating data from other sources of urban infrastructure including cellular networks and automobile parking sensors to investigate how they might augment our analysis and provide different insights into human behavior. We are also interested in building internet- and mobile-based applications that suggest nearby stations based on predictive models. In addition, we have begun logging eighteen additional shared bicycling programs including Paris, Auckland and Washington D.C. We are planning a large-scale analysis that compares behavioral patterns across cities. Some of our results have obvious implications for the design and operation of the shared bicycling system itself. For Bicing, we have met with the Barcelona city government to discuss our findings and continue to correspond with them about our progress.

Acknowledgements

We thank Fabien Giradin for early discussions about this work and Mauro Cherubini, Sunil Garg, Volker Hohmann, Travis Kriplean, Neal Lathia, and Karen Züpferwomann for their comments on early drafts. We also thank Dave Chui for his Bicing pictures.

References

[1] Girardin, F., Calabrese, F., Dal Fiore, F., Ratti, C., Blat, J. (2008). Digital footprinting: uncovering the presence and movements of tourists from user-generated content. *IEEE Pervasive Computing.*

[2] González, M. C., Hidalgo, C. A., and Barabási, A. L. (2008). Understanding individual human mobility patterns. Nature 453, 779-782.

[3] Ratti, C., Pulselli, R. M., Williams, S., and Frenchman, D. (2006). Mobile landscapes: Using location data from cell-phones for urban analysis. Environment and Planning B: Planning and Design, 33(5). 727 – 748.

[4] Witten, I. H. and Frank, E. (2005). Data Mining: Practical machine learning tools and techniques. 2nd Edition, Morgan Kaufmann, San Francisco.

Multimodal Sensing for Pediatric Obesity Applications

M. Annavaram†,N. Medvidovic†, U. Mitra†, S. Narayanan† G. Sukhatme†,
Z. Meng‡, S. Qiu‡, R. Kumar†, G. Thatte†, D. Spruijt-Metz§
† Viterbi School of Engineering, University of Southern California, Los Angeles, CA 90089, USA
Email: {annavara,ubli,thatte}@usc.edu, {neno,gaurav}@cs.usc.edu, shri@sipi.usc.edu
‡ Tsinghua University, Beijing China
Email: {third,forth}@institution.edu
§ Keck School of Medicine, University of Southern California, Los Angeles, CA 90089, USA
Email: metz@usc.edu

Abstract

In this paper, a wireless body area network comprised of heterogeneous sensors is developed for wearable health monitoring applications. The ultimate application space is in the context of pediatric obesity. The specific task examined herein is activity detection based on heart rate monitor and accelerometer data. Based on statistical analysis of experimental data for different key states (lying down, sitting, standing, walking and running), a multimodal detection strategy is proposed. The resulting detector can achieve 85-95% accuracy in state detection. It is observed that the accelerometer is more informative for the active states, while the heart rate monitor is more informative for the passive states.

I. Introduction

Wearable health monitoring systems coupled with wireless communications are the bedrock of an emerging class of sensor networks: wireless body area networks (WBAN). The objectives of such WBANs are manifold from diet monitoring [14], activity detection [3], [4], and health crisis support[6]. These new networks demand significant technological advances from sensor development to novel software engineering, signal processing, wireless communications and networking. Importantly, WBANs must be designed with application-specific design and end-use requirements in mind. These advancements are necessary to cope with the unique challenges introduced by deployment **on** people, such as: unpredictable mobility, heterogeneous sensor nodes, new wireless channels, very low power requirements, non-invasive sensing and the need for sensors with small footprints. Furthermore, drawing robust inference from sensor streams requires information from multiple, often disparate, sources. In the current work, we provide preliminary results from the construction of a WBAN which we will use to drive the development of assessments and interventions for pediatric obesity applications.

Pediatric obesity has emerged as a major national and international health crisis. National collected data from 2003-2006 show 11.3% of adolescents aged 12 - 19 years by some measures could be designated as obese; a further 16% would be classified as overweight and 32% considered at risk for being overweight [13]. While physical activity (PA) is tightly related to lower obesity rates in children [11], [7], there are additional factors leading to obesity. The increasing environmental stress may promote both general obesity (through lifestyle behaviors such as decreased physical activity) and visceral obesity (through hypothalamic-pituitary-adrenal axis activation and increased cortisol secretion)[5]. Current monitoring systems validated for research in children typically monitor physical activity only (such as the much-used Actigraph accelerometer). However, in order to truly understand and reverse childhood obesity, we need a multimodal system that will track stress levels, PA levels, blood glucose levels and other vital signs simultaneously, as well as anchor these levels to context such as time of day and geographical location. Our preliminary KNOWME network is a first step towards such a system.

A key aspect of our work is the unified design and evaluation of multimodal sensing and interpretation, for automatically recognizing, predicting and reasoning about human physical activity and socio-cognitive behavior states. On the one hand, this meets the needs of traditional observational research practices in the obesity and metabolic health domain (based on, and validated through, careful expert human coding of data) while on the other, this enables new analysis capabilities that have not been possible before such as providing information on user emotional state in conjunction with physical activity and energy expenditure.

Many aspects of human behavior are inherently multimodal or require multimodal processing. For example, measuring and understanding energy expenditure and its etiology requires processing not only activity from accelerometers but other data such as pulse rate, ECG, oxygen intake, as well as contextual information such as emotions that are marked by humans through their voice, body posture and through physiological signals skin conductance measures (electro dermal response). Hence, to model human behavior and task-specific activity, both in terms of what people do, how they do it, and *why they do it*, it is critical to understand and capture the interplay between such multimodal streams. Multi-modal coverage of our approach enables cross-channel comparison and verification (allowing us, for example, to capture relationships between increased heart rate, increased emotional activity, and changes in physical activity). Our approach to this problem is grounded in statistical signal processing.

In the current work, we summarize preliminary results on activity assessment. We consider a mix of low mobility (lying down, sitting, standing) and higher mobility (walking, running) states. Features of our problem and approach do appear in the prior literature. Much work on activity detection appears to center on accelerometer data alone (*e.g.*[8], [3], [10]) with some systems employing many accelerometer packages. On the other hand, multi-sensor WBANs have been implemented and deployed (see *e.g.* [12], [9], [6]); however in those works, the emphasis was on the higher layer communication network processing and hardware design – signals from each sensor were transmitted directly to a central decision making unit. Our focus is on a modest number of heterogeneous sensors and the utilization of multi-modal signal processing methods; we wish to design decision making and data interpretation methods that will reside within the WBAN and allow for interaction with the WBAN wearer. For our pediatric obesity application, activity detection is an indirect measure of energy expenditure quantification as discussed above. In [4], multi-modal classification **is** considered. There are some key differences to the approach taken herein. First, while different sensors are employed, they are similar in the types of measurements taken (*e.g.* accelerometers, gyroscopes and tilt measurements), herein we use sensors which measure fundamentally different quantities that are correlated, but the statistical relationships are unclear *a priori*. The goal of [4] is to determine a sampling scheme (with respect to frequency of sampling and sleeping/waking cycles) for multiple sensors to minimize power consumption. The authors show that their new methods achieve reduced power relative to classical joint schemes. Our goal is on classifier performance with heterogeneous sensors – future versions of our methods could incorporate power minimization strategies of [4]. An important question to address is how the correlation between measurements affects power minimization. We conjecture that the sensors employed in [4] have more highly correlated observations with regards to the states of interest than our sensors and thus greater power minimization is possible through the use of their methods.

As our WBAN must be used for a diverse set of decision making processes, all sensors may not be uniformly useful for each task. We, in fact, see this with the activity detection problem considered herein.

II. KNOWME Network Architecture

The basic foundation of the KNOWME network is our three tier network architecture as depicted in Figure 1. The first tier's goal is data collection based on the heterogeneous sensors that are coupled to a mobile phone which acts as a "base station," equipped with data transmission and processing capabilities. The second tier is a web server that receives data and can perform additional processing; the web server transmits the data to the final tier: a back-end database server that stores the information. In the sequel, we shall discuss the specific sensors employed.

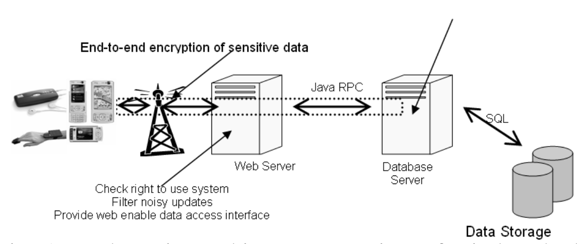

Fig. 1. Three-tier architecture overview of wireless body area network sensor system.

Currently, the primary focus of this research is to perform multi-modal sensing and interpretation of data to serve some of the end-user needs. As such, significant effort has been spent in integrating heterogeneous sensors to a mobile phone. One challenge in integrating heterogeneous sensors is that these sensors have different APIs, packaging, and data collection methods. In addition to integrating multiple sensors, synchronization of the data received from multiple sensors in the phone is critical for statistical correlation of sensor data and to perform the multi-modal data processing. Sensor information is continuously recorded on the local storage on the mobile phone. Our mobile device platform has a 8GB in-built flash memory that can be used for storing sensor information. Sensor data rates vary from 300bps for the accelerometers to 100 bps for the heart rate monitor. Using these data rates, we estimate that our 8GB local storage can store 1000 days worth of data. As the Bluetooth wireless link is a bottleneck for our current data collection, we use time-division multiple-access to schedule the data from different sensors (equal time share).

The software development phase uses well-known unit testing to extensively test the mobile software suite. In order to minimize errors in configuring the software, our software has several built-in checks to advise the user if any of the sensor readings do not match expected sensor behavior. Since the mobile device has to transmit the data to the backend servers, we are currently developing an opportunistic data transfer mechanism that uses an open WiFi network where available to transfer data both efficiently and cheaply. In the absence of WiFi networks, the mobile software is configured to automatically use the cellular data network to transmit the data. Our initial deployment is mostly with graduate and undergraduate student test subjects with limited (on-going) pilot experiments with children in the Exercise Physiology Lab at the USC Keck School of Medicine.

A. Sensor Systems

The sensor layer is a collection of off-the-shelf devices that measure features which can provide insight about metabolic activity; most (with the exception of galvanic skin response) are also capable of wirelessly transmitting this data over a Bluetooth interface. The current study employs an Alive Technologies[1] electrocardiograph (ECG). The ECG is a single channel device with 8 bit resolution and a peak sampling rate of 300 samples/second. The pulse-oximeter, also from Alive, provides non-invasive monitoring of oxygen saturation (SpO2) and pulse rate. The oximeter is a Bluetooth slave device that supports the Bluetooth Serial Port Profile (SPP). We also have BodyMedia WMS sensors [2] to measure Galvanic Skin Response (GSR) [1] and motion estimation using accelerometers. We use feature rich Nokia N95 as the mobile phone platform. N95 supports Bluetooth 2.0 + EDR for quick pairing with external Bluetooth sensors, and has 3G and WiFi radios for high bandwidth data transfer. In addition to the high bandwidth radio capabilities, the N95 mobile phone platform has a highly accurate built-in assisted GPS unit that uses a combination of GPS satellites, cellular tower and WiFi scanning to obtain a GPS position lock in less than 10 seconds. The stated location accuracy of GPS unit is 30 meters. We have observed accuracy at less than 3 meters in practice. The data collected from multiple sensors is geo-tagged using the location data collected from the in-built GPS. Furthermore, our system is also capable of audio and video tagging to assist users to supplement the automatically collected sensor data (as in [14]). Some WBAN components are depicted in Figure 2.

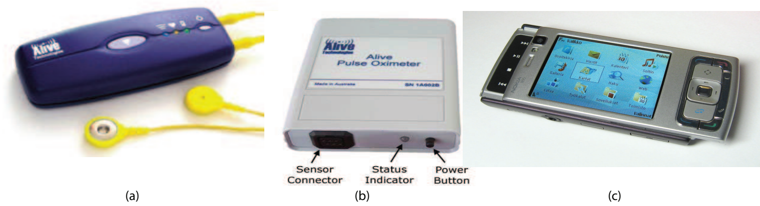

Fig. 2. (a) ECG monitor, (b) pulse oximeter, (c) Nokia Smartphone (GPS and accelerometer).

III. Activity Modeling

Data collected from our experimental system setup can be used in multiple contexts, for instance by the users to regularly monitor their physical well being as well as by medical practitioners in assessing the physical health of their patients. Here, we describe one such application of using the data to automatically derive the activity of a person with data collected from multiple sensors. Statistical modeling of various test subject states was undertaken based on the data collected from the WBAN. We examined five different states: lying down, sitting, standing, walking and running. Again, to reiterate, activity detection has been previously considered with an emphasis on the use of many accelerometers, yielding a cumbersome network to wear. We conjecture that multimodal data analysis will enable the achievement equal or even better accuracy and robustness in activity detection with fewer sensors.

[1] The data of the WMS GSRs are not currently included due to issues with time synchronization.

In this research, multiple distributions were considered to fit the data which for each sensor was predominantly unimodal in nature. After extensive experimentation, the use of the pulse oximeter sensor was abandoned due to limited change in readings for any of the states of interest for our activity detection problem. Thus, we focused on ECG and accelerometer data. The distributions under consideration were: T location-scale, Gaussian, log-normal, logistic, log-logistic, one-side Gaussian and Laplacian. Where possible, Gaussian distributions were selected to facilitate the determination of joint densities. The ECG data were preprocessed as follows: peak detection was performed and the inter-peak time collected. The inter-peak time was modeled as a Gaussian random variable. An average of the empirical variance for each of the axes over a pre-specified window of time for the accelerometer data was employed. The walking and running state data were modeled as Gaussian; however, the lower-activity level data (lying down, sitting and standing) was modeled as a Laplacian to achieve a better fit. Figure 3 (L) and (R) shows the ECG and accelerometer data for the running and sitting modes, respectively. We see that both states are relatively well distinguished from each other with significant differences in the accelerometer data.

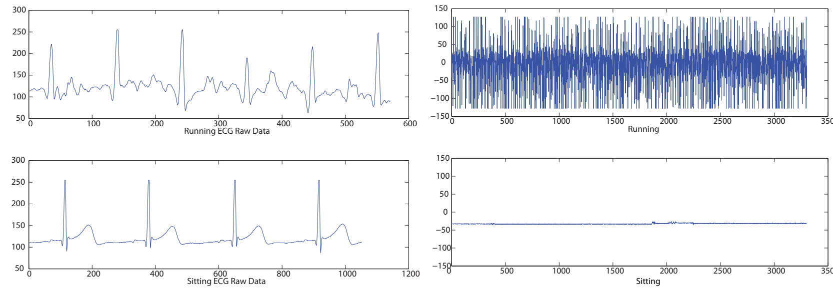

Fig. 3. (L) ECG and (R) accelerometer data from the heart-rate monitor for sitting and running.

Not surprisingly, ECG and accelerometer data had different discriminatory properties for the various states, underscoring the benefits of multi-modal sensing and signal processing. In Figure 4, we see the statistical fits for the accelerometer data for high activity states and the statistical fits for the ECG data for low activity fits. To develop bivariate models (joint densities) for the ECG and accelerometer data, additional processing (resampling) was required to determine the correlation between the ECG statistic and the accelerometer statistic in the high-activity levels.

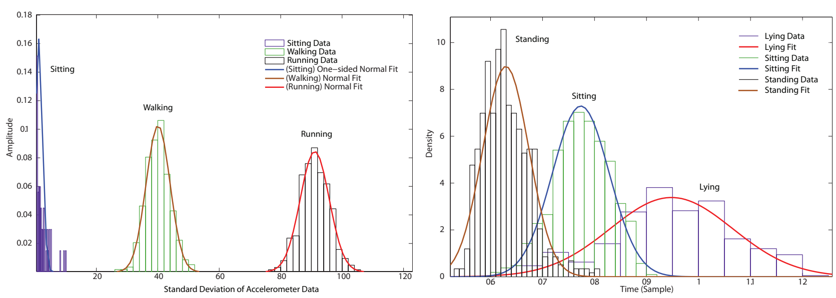

Fig. 4. (L) Statistical fitting for higher activity states (accelerometer data): sitting, walking, and running. (R) Statistical fitting for lower activity states (ECG data): lying down, sitting, and standing.

In the low-activity level cases, the ECG and accelerometer statistics were assumed to be independent. The resulting bivariate densities for each of the five hypotheses are shown in Figure 5(L) and (R). For clarity, the low activity states are shown separate from the higher activity states. Bivariate testing yielded state detection rates on the order of 85% to 95% – achieving detection rates with two heterogeneous sensors comparable to the rates found in [3], where nine single mode (accelerometer) sensors were employed.

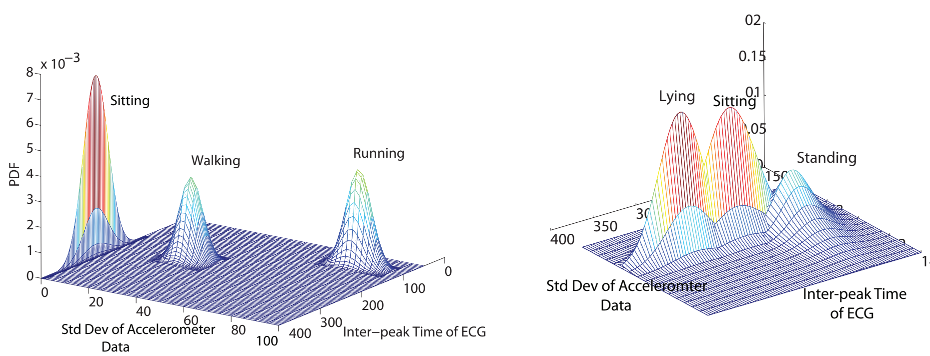

Fig. 5. Bivariate distributions for (L) running, walking and sitting and for (R) lying down, sitting and standing.

IV. Observations and Ongoing Work

Our preliminary system successfully collects data and transmits it to the cellular phone. We conjecture from our experiments that a few heterogeneous sensors may offer better discrimination and robustness than many homogeneous sensors. Our preliminary data for activity detection in comparison to [3] appears to bear this out this conjecture. There are however important engineering challenges associated with WBANs, especially for activity detection. For our particular set up, we are limited by the mobile phone platform which can only accommodate a maximum of eight different sensors. If all sensors sample at their maximum sampling rate, the expected throughput would exceed the capabilities of the Bluetooth link leading to dropped packets. The battery power of the cellular phone is another bottleneck for the system. Finally, for activity detection, high activity/mobility can impair a sensor's ability to sense. This fact can be viewed two ways: it is detrimental in that we lose sensor accuracy, on the other hand, new features are introduced into the signal which are still indicative of high activity. Our preliminary results suggest that sensor selection and prioritization will be important to ensure that packets are not lost; furthermore energy aware sensor management will be critical.

We have recently conducted a pilot study with two pre-adolescent girls following an observation protocol typical for pediatric obesity studies. We are currently analyzing this data, including designing multi-modal detection algorithms for deciding between the various states. We hope to share those findings at the workshop. Finally, introducing contextual cues for use of the WBAN in everyday life will be extremely important; to this end, the image processing and analysis methods of DietSense [14] will prove very useful. Finally, as noted earlier, power minimization is of high importance for WBANs and their attendant applications; we expect the methods of [4] will have promise when properly adapted to our context.

Acknoweldgements

The research is funded in part by Qualcomm. Additional support is from Nokia for the Nokia Smartphones used in our studies.

References

[1] Wireless Health Monitors from Alive Technology, www.alivetec.com, Retrieved on July 24, 2008.
[2] Sensewear BMS available at http://www.sensewear.com/solutions_bms.php
[3] S. Biswas and M. Quwaider. Body Posutre Identification using Hidden Markov Model with Wearable Sensor Networks. *Proc. of BodyNets Workshop 2008*, Tempe, AZ, March 2008.
[Sensys 2007:]
[4] A. Benbasat and J. Paradiso , A framework for the automated generation of power-efficient classifiers for embedded sensor nodes. *Proceedings of Sensys*, pp. 219-232, Novemberm 2007, Sydney, Australia.
[5] P. Bjorntorp, R. Rosmond. Neuroendocrine abnormalities in visceral obesity. *International Journal of Obesity & Related Metabolic Disorders: Journal of the International Association for the Study of Obesity* 2000;24 Suppl 2:S80-5.
[6] T. Gao, C. Pesto, L. Selavo, *et. al.*. Wireless Medical Sensor Networks in Emergency Response: Implementation and Pilot Results, *Proc. 2008 IEEE International Conference on Technologies for Homeland Security*, May, 2008.
[7] M. Goran, K. Reynolds, C. Lindquist. Role of physical activity in the prevention of obesity in children. *International Journal of Obesity & Related Metabolic Disorders: Journal of the International Association for the Study of Obesity*. 1999;23:S18-33.
[8] S. Jiang, Y. Cao, S. Iyengar, *et. al.*. CareNet: An Integrated Wireless Sensor Networking Environment for Remote Healthcare. *Proc. of BodyNets Workshop 2008*, Tempe, AZ, March 2008 (work in progress paper).
[9] E. Jovanov, A. Milenkovic, C. Otto and P. C. de Groen. A wireless body area network of intelligent motion sensors for computer assisted physical rehabilitation. *Journal of NeuroEngineering and Rehabilitation*, 2:6, March 2005.
[10] A. Kalpaxis. Wireless Temporal-Spatial Human Mobility Analysis Using Real-Time Three Dimensional Acceleration Data. *Proc. of International Multi-Conference on Computing in the Global Information Technology* , March 2007.
[11] S. Kimm, N. Glynn, E. Obarzanek, *et. al.* Relation between the changes in physical activity and body-mass index during adolescence: a multicentre longitudinal study. *The Lancet* 2005;366:301.
[12] D. Konstantas, A. Van Halteren, R. Bults, *et. al.* MobiHealth: Ambulant Patient Monitoring over Public Wireless Networks. *Proc. Mediterranean Conf. on Medical and Biological Engineering (MEDICON)*, August 2004, Ischia, Italy.
[13] C. Ogden, M. Carroll, and K. Flegal. High Body Mass Index for Age Among US Children and Adolescents, 2003-2006. JAMA 2008:299:2401.
[14] S. Reddy, A. Parker, J. Hyman, *et. al.* Image Browsing, Processing, and Clustering for Participatory Sensing: Lessons From a DietSense Prototype. *Proc. EmNets'07*, June 2007, Cork, Ireland.

Personalized Awareness and Safety with Mobile Phones as Sources and Sinks

Anna Yu, Athanasios Bamis, Dimitrios Lymberopoulos, Thiago Teixeira and Andreas Savvides
Embedded Networks and Applications Lab
Electrical Engineering Department
Yale University, New Haven, CT 06520
Email:{firstname.lastname}@yale.edu

Abstract

Today's mobile phones are equipped with an increasing set of sensors including GPS, accelerometers, cameras, and more that makes them ideal source nodes in urban sensing applications. The growing displays and internet connectivity also makes these phones excellent sinks of just-in-time information, including information from other sensors deployed in the infrastructure. Exploiting these features, our personalized safety and awareness system tries to enhance personal safety of users around the clock through a collection of services that process personal and aggregate community data to track, escort, flag, supervise behaviors and help users coordinate to enhance the collective safety of the group. Designed for individuals and groups that operate on campuses and beyond, it intends to make campuses safer, by going beyond the processing of individual location data and by providing services based on the application of intelligent behavior sensing algorithms and collaborative models to aggregate sensor data.

I. INTRODUCTION

Most campus security plans consist of scattered emergency phones, scheduled shuttles, and foot or vehicle escorts, but such plans are not always very effective with rising student population numbers and sudden spikes in localized security demand. Moreover, many campus security implementations are unreliable and unable to meet the needs of busy individuals. Confusing timetables, unclear pick-up locations, and limited hours of service discourage many people from actually using campus security services. To compound this problem, many people hesitate to call for a security escort out of embarrassment, or false beliefs that they are immune to danger. To address some of these challenges, our system takes a broader view to personal safety, leveraging GPS mobile phones and social networking to introduce dynamic safety practices. Our system provides user customizable activity monitoring that begins to form the basic virtual escort tracking for small trips on foot, longer term tracking during travels, and escalates up to model-driven activity monitoring and community based coordination for safety. Instead of focusing on security and privacy issues our research is directed towards the creation of semantic meaning from the sensor data, particularly reasoning with user locations in time and space, also using context information extracted form maps. Privacy issues are implicitly handled by exploiting the phones' local processing capabilities, provisioning for the use of security and privacy from other researchers [8] and by operating in community mode where users are willing to share some level of private information in aggregate form with other members of the community to enhance community-wide safety.

In this position paper we describe a personal safety and awareness framework that is currently being developed as part of the Behavior-Scope (BScope) project at Yale [1]. This is centered on the use of smartphones as sources and sinks of information and involves coordination among multiple phones as well as other sensors deployed in the environment. The mobile phones coordinate with a central server to provide a set of services to the users. For instance, when walking across campus, users can put their mobile phone client application in a virtual escort mode. This service provides a panic button option and tracks the travel progress of the user to ensure that the user safely reached the intended destination. During longer trips, a travel service sends automatic emails and text messages to family and friends providing updates about the trip. For more general personal safety, the phone also learns the daily routines of the user and notifies a set of registered recipients at different levels of behavior deviations. Finally, a set of aggregate location information and inputs from campus security are used to coordinate the movements of users and campus security personnel during late hours to ensure maximum safety coverage while moving around campus. Examples of such coordination include pairing up members of the same group to walk together at night, providing safe walking route advice (i.e route that currently has the most members or most security personnel around) and dynamically re-positioning security officers (on foot, bike or car) to the places demand surfaces. We also anticipate enhancing safety and awareness through coordination between friends and social networking. The same system can also bind into the previously developed BScope home monitoring infrastructure, to monitor user safety at home and also to notify caregivers of the status of a loved one. In all cases, the mobile phones are used are the primary user interface and sink device.

The rest of the paper outlines the system architecture in Section II; and the current application features in Section III. Section IV and provides a brief description of our results on a rule behavior engine and activity modeling engine we are developing. Section V concludes the paper with a set of research challenges considered by the BScope project.

II. SYSTEM ARCHITECTURE FEATURES

The broader personal safety model introduced by our architecture (shown in Figure 1) is driven by the interpretation of spatio-temporal properties of multimodal sensor measurements. This process is primarily driven by two interpretation engines: a probabilistic rule-based behavior engine and a dynamic activity modeling engine. The network traffic flows over a heterogeneous set of links including Wi-Fi, IEEE 802.15.4, GPRS, CDMA data links and Ethernet. This forms a dataflow graph in which data propagates in the direction of a central server and back to the sink devices. The central server controls all the data processing modules and can make decisions to diffuse them into the network as needed to improve response times, conserve bandwidth and to interpret data closer to the sinks to manage complexity and increase system robustness.

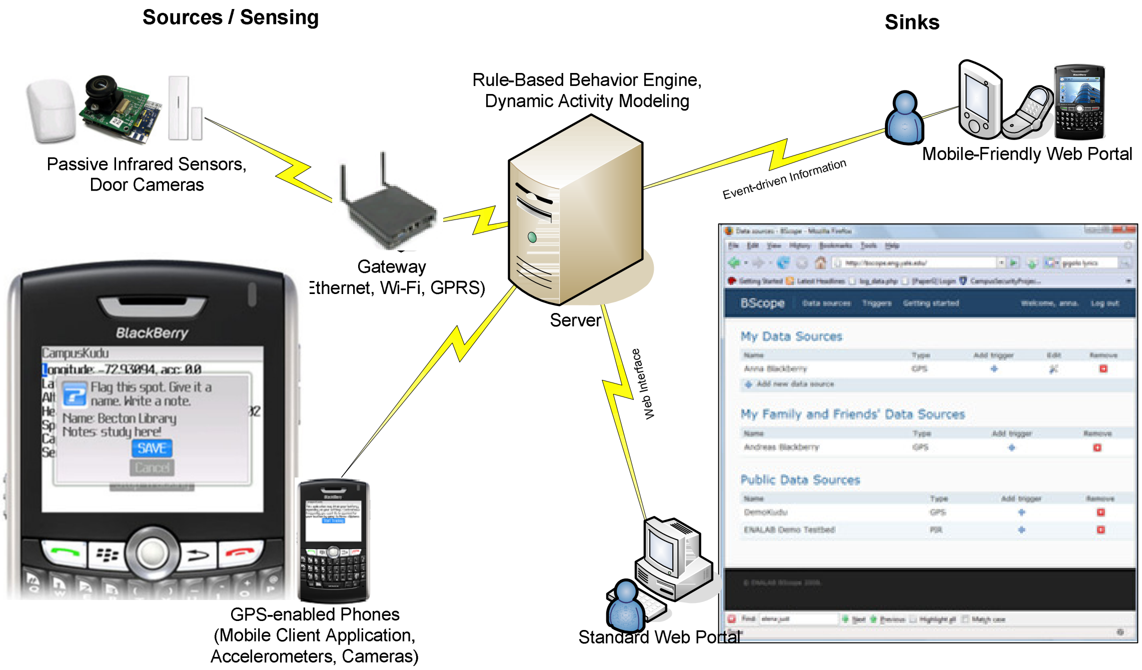

Figure 1: System Architecture

The system uses GPS phones (currently supporting BlackBerry GPS phones) for outdoor environments a collection of cameras, passive infrared and door sensors for indoor environments. Users can configure how the system should work for them and their level of participation in the system directly from their mobile phones. A more sophisticated web portal allows users to interact with the systems at different levels, allowing them to define automated notifications, develop new applications and specify how the incoming data should be processed and plotted. The users can specify how they want the system to work for them as their preferences essentially configure parameters in the rule-based behavior and the dynamic activity modeling engines described in section IV. The two engines reflect the design philosophy of the system when it comes to monitoring human behaviors and providing services. A growing part of the system is centered on learning behaviors and patterns to detect deviations from the norm. Nonetheless, our testbed experience (and of course human nature) demands checking for certain conditions and rules that are not directly extractable. Furthermore, statistical outliers are not always alarming and vice-versa. Because of this, a core aspect of our system relies on the collaboration and sharing of parameters between the two engines.

III. CURRENT MOBILE CLIENT FEATURES

Our initial deployment in an urban setting heavily relies on the client application running on mobile phones. The deployed mobile currently supports several features than enable it to be an active contributor to the overall system architecture. The mobile client automatically connects to our deployed web portal (www.sense4care.com), where users can also configure their preferences in more detail. The mobile client (currently supports GPS-enabled Blackberry phones) can also be downloaded from the web portal. Its key features are highlighted below.

A. Tracking

The mobile client contains a basic tracking feature that users can select to turn on to allow their location be sent to a central server. The user can choose to have this feature on 24-hours a day, or just during select commutes or times of day. Ideally, the application will be running continuously so as to collect as much information and allow as much personalization as possible. This location information is then accessible from any computer or mobile device that has access to the internet via either the standard web portal or the specially-designed mobile-friendly web portal. The benefit of this is that you don't need to have a GPS-enabled mobile phone in order to participate in this system. In order for individuals or groups to access a user's location information, the user must first authorize this access.

B. Virtual Escort

The "Virtual Escort" feature is an integral part of the new campus safety model. It allows users to have an escort when the user cannot find anyone else to walk with. This feature is useful for users who frequently need to walk outside late at night and are seeking a cost-effective and time-efficient solution to staying safe and gives the user access to a programmable PANIC button that can let security know there's trouble and exactly where the user is.

C. Triggers

The web interface also allows users to set up triggers that inform family and friends via SMS or e-mail alert when the user is leaving or entering a pre-defined space. It automates the process of checking a user's location by having an automated message sent out according to the pre-specified preferences. These triggers can be simple conditions about geographic locations or more elaborate behaviors as described in section IV.

D. "Auto-K.I.T"

If the user defines certain areas as being associated with specific activities, the system engine can write automatic digests of a user's day to send to friends and families. This would allow a user to automatically "keep-in-touch" with everyone, even when very busy and has no time to call or e-mail. Also, at the central security server, this is driven by a powerful behavior interpretation engine that would allow system management, and multiple levels of user-defined safety mechanisms.

E. Mobile Phone Power Management

Since the goal of this application is to make it easy to stay safe and secure anywhere, at anytime, the application needs to efficiently manage its power consumption and make its state known to the server at all times. The application informs the server of its status on power up and shutting down, loss of GPS signal, and feature usage. To conserve power, local processing, intelligent sampling and other sensors such as accelerometers need to be exploited. Our prototype experiences have shown that reading the GPS alone can take a noticeable toll on the phone's battery lifetime [4]. Such excessive power consumption could be reduced by utilizing accelerometer sensors and context inferred from the behavior monitoring applications to intelligently manage the GPS sampling and communication frequency. The BlackBerry smart phones used in our prototype deployment do not have accelerometers but other phones such as the Nokia N95 and iPhone already have them. We anticipate that more phone models will have them in the future.

IV. SENSOR DATA INTERPRETATION

A. *Probabilistic Rule-Based Behavior Engine*

This is based on BScope's hierarchical probabilistic grammar framework detailed in [2,3]. The framework follows a language-based approach that uses semantic-level sensor outputs as a sensing abstraction. To sense an activity, the environment needs to be instrumented by a set of sensors to extract a string of phonemes. The collection of phonemes is parsed by small libraries of probabilistic grammars we call *sensory grammars*. The outputs of these grammars are higher order phonemes that can be parsed by other grammars in the hierarchy. A time-abstraction layer allows the sensory grammar framework to reason with temporal quantities. Spatial quantities are implicitly considered by the system through sensor labeling.

In the personal safety system discussed here, the rule-based behavior engine can be applied in several different ways. Users can enter simple rules as regular expressions, and can then form higher-level rules using the outcomes of pre-programmed behavior grammars, or they can develop elaborate grammars from scratch. This hierarchical mode of operation can also be perceived as a higher order, sensor composition tool that can tie together simple sensors in time and space to form a more complex sensor. A detailed example of the composition of a cooking sensor can be found in [5]. The application of this system in elder monitoring and the network architecture used can be found in [6].

B. *Activity Model Extraction*

Although the Rule-Based Behavior engine can classify sensor data into pre-defined activities, the extraction of intuitive activity models directly from the data should be an integral component of an autonomous personal safety system. Ideally, the personal safety application running on a mobile phone should be able to extract the daily or weekly habits of the user and use them to measure deviations and generate alarms when behaviors deviate by a certain threshold. This is a challenging task since the system needs to identify the recurring activities over a time-window without any prior definition of the activities. The mobile phones can collect GPS locations but locations alone only provide a low-level dataset that does not directly imply specific activities. Furthermore, the models need to capture not only the spatial meaning of the data but also its temporal properties. For instance, a consistent everyday visit at a certain location has a different meaning than visits to the same location at unusual times. Our initial effort in this direction has developed a four step model extraction method that operates on the three main attributes of sensor measurements: location, time and duration. The four steps are summarized below:

1. Decritize locations using the map context, that is, name the data according to the named area on the map. This converts locations in the dataset to a series of labeled events.
2. For each event type, group together all event instances that have similar start times and durations over a time window and relabel the events with spatio-temporal labels according to their cluster classification. This part applies a new clustering algorithm we have developed to group together events that have similar start-times and durations over the course of a time-window (i.e a day or a week). This is done by applying a set of similarity metrics, without knowing the number of clusters in advance. This step essentially descritizes time into a finite set of labels and appends these labels as suffices to the labels of step 1.
3. Using the resulting sequence from step 2, extract activities by mining out the most frequently recurring subsequences. This can be done using a data mining algorithm such as the apriori algorithm described in [9].
4. The outcome of step 3 provides enough information to build a state machine that describes the activity. The spatio-temporal event labels of the mined subsequences become the states of the model. The transition probabilities between the states are computer by counting and normalizing the transitions among spatio-temporal events from step 3.

The above methodology has been tested on an indoor trace originally collected for elder monitoring. Inside homes people locations are sensed with cameras and passive infrared sensors but the mechanism of the model are very similar. Some of our initial results can be found in [7]. The modeling methodology is currently being extended and tested on GPS data collected from mobile phones.

V. FUTURE RESEARCH DIRECTIONS

The safety system described in this paper opens a new set of interesting research problems on how to collect, interpret and utilize sensor measurements across communities of mobile phone users. Several aspects of the safety coordination, route recommendations and security officer patrols could be treated as dynamic sensor coverage problems. Urban environments however, and groups make this a challenging problem. Sensors are mobile, the terrain is not a flat plane and map context should be exploited.

From an architecture perspective, our position is that a user reconfigurable data interpretation service that can distribute processing modules as needed across a network of heterogeneous wireless links would provide a powerful feature for developing new services. Finally, we advocate that mobile phone technology has reached a level where it can become the primary information sink for people. Mobile phones have become deeply pervasive to everyday lives and providing an event-driven sensing framework that can notify users about interesting matters in their lives as they happen will form the basis for a new generation of applications. The BScope project at Yale is working towards this direction by using mobile phones as part of a campus security program and as the main information sink for elder monitoring. In the elder monitoring scenario, stakeholders and caregivers get event-driven phone notifications on their mobile phones by configuring the two engines described in this paper. More up-to-date information about the project can be found on the BScope project website [1] and its affiliate website at [10].

ACKNOWLEDGMENT

This work was partially funded by the National Science Foundation under projects CNS 062802 and CNS 0751513. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation (NSF).

REFERENCES

[1] BehaviorScope Project Website. http://www.eng.yale.edu/enalab/behaviorscope.htm

[2] Lymberopoulos, D., et. al. Macroscopic Human Behavior Interpretation Using Distributed Imager and Other Sensors, Proceedings of IEEE, August 2008.

[3] D. Lymberopoulos T. Teixeira and A. Savvides, Detecting Patterns for Assisted Living Using Sensor Networks, to appear in the proceedings of SensorComm, Valencia, Spain, October 2007

[4] Yu, Anna S. “Personalized Security and Social Networking Services over GPS-Enabled Mobile Phones”, Final Senior Project Report, Electrical Engineering Department, Yale University, May 2008.

[6] A. Bamis and D. Lymberopoulos and T. Teixeira and A. Savvides Towards Precision Monitoring of Elders for Providing Assistive Services to appear in the proceedings of the First International Conference on Pervasive Technologies Related to Assistive Environments, Athens, Greece, July 2008

[7] D. Lymberopoulos and A. Bamis and A. Savvides Extracting StatioTemporal Human Activity Patterns in Assisted Living using a Home Sensor Network, Proceedings of the First International Conference on Pervasive Technologies Related to Assistive Environments, Athens, Greece, July 2008

[8] Murali Annavaram, Quinn Jacobson, John Shen, HangOut: A Privacy Preserving Social Networking Application, First International Workshop on Mobile Device and Urban Sensing (MODUS), St. Louis, April 2008

[9] R. Agrawal and R. Srikant. Mining sequential patterns. Proceedings of ICDE. Washington, DC, USA, 1995. IEEE Computer Society.

[10] Sense4Care Website. http://Sense4Care.com

Increasing the precision of mobile sensing systems through super-sampling

RJ Honicky[†], Eric A. Brewer[†], John F. Canny[†], Ronald C. Cohen[‡]
[†] Department of Computer Science, UC Berkeley
Email: {honicky,brewer,jfc}@cs.berkeley.edu
[‡] Department of Chemistry, UC Berkeley
Email: rccohen@berkeley.edu

Abstract

Sensors integrated into mobile phones have the advantage of mobility, co-location with people, pre-built network and power infrastructure, and potentially, ubiquity. These characteristics, however, also present significant challenges. Mobility means non-uniform sampling in space, and also constrains the size and weight of the sensors. In this paper, we focus on non-uniform sampling, and imprecision. We investigate the question, "Assuming well calibrated sensors, what precision can we expect from a network of sensors embedded in location aware cell phones?" We briefly describe some results that suggest that a Gaussian process based model is appropriate.

I. Introduction

With increased public focus on environmental conditions and increasing industrialization of developing countries, the need for environmental monitoring has increased significantly. Current air pollution monitoring systems typically consist of highly sensitive, bulky equipment placed in a few strategic locations. These systems, such as the California Air Resource Board (CARB) monitoring system mostly monitor ambient levels over large geographic areas [1]. Not only do systems like CARB have very coarse granularity, but they also only measure the human and environmental health impacts of pollution indirectly.

The Networked Suite of Mobile Atmospheric Real-Time Sensors (N-SMARTS) project [2] aims to radically improve the geographic coverage and granularity of environmental monitoring by integrating pollution (and other environmental) sensors into location-aware mobile phones. Our current sensor devices connect to the phone via Bluetooth, and will eventually fit into a modified battery pack, for tight ergonomic integration. Sensors integrated into mobile phones have the advantage of mobility, co-location with people, pre-built network and power infrastructure, and potentially, ubiquity.

These characteristics, however, also present significant challenges. Mobility means non-uniform sampling in space, and also constrains the size and weight of the sensors. Although co-location with people means that samples will often be taken near a particular person, hence providing a good approximation of a person's exposure to pollution, co-location also means that a person's behavior (putting their phone in their pockets, riding in cars, remaining indoors vs. outdoors) will impact the readings of the sensors. Tracking a person's location also has enormous privacy implications. Ubiquity implies low cost and, coupled with size constraints, low-precision sensors. Embedding sensors into a ubiquitous device also implies a passive sensing model, in which the user can not be expected to perform any action to sense the environment, nor can they be expected to calibrate or otherwise maintain the sensor.

In this paper, we focus on a small piece of this puzzle: non-uniform sampling, and imprecision. We investigate the question, "Assuming well calibrated sensors, what precision can we expect from a network of sensors embedded in location aware cell phones?" We make a case for using a Gaussian Process noise model and show some early empirical and simulation results.

A. Problem formulation

Fundamentally, we are interested in measuring and characterizing the environment using sensors embedded in location-aware mobile phones. For the sake of concreteness, in this paper we focus on carbon monoxide, but we believe that these results will extend to many other environmental factors, including other gaseous pollutants, aerosol pollutants, radiation and network signal strength.

Since we are interested in modeling the environment as people experience it over time, we use a model with two spacial dimensions (people basically move two dimensionally), and a temporal dimension.

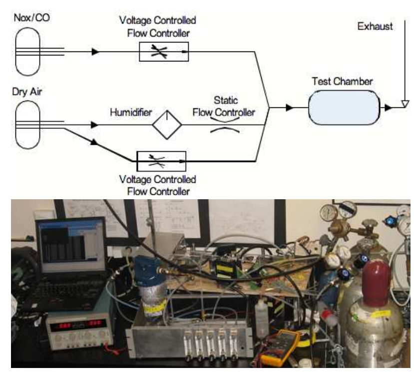

Fig. 1. The test chamber allows precise control of the concentration of toxic gases and fast response, which allows precise calibration and characterization of the sensors.

B. Data

In order to understand pollution sensors in greater detail, we have designed a series controlled laboratory experiments. To characterize the CO sensors we are using, we use two electronically controlled mass flow controllers, one attached to pure air, the other attached to 100ppm CO air. The output of the flow controllers is then pumped into a cylindrical chamber that contain six sensor and associated electronics. Finally, the gas is injected into the laboratory's exhaust system (see Figure 1). This setup allows us to precisely control the concentration and rate of flow of CO in the sampling chamber. The sensors and flow controller are monitored and controlled using a NI USB-6218 data acquisition module from National Instruments attached to a laptop.

II. A Gaussian noise model

Sensor noise is often well modeled with a Gaussian distribution. One reason for this is that Gaussian noise turns out to be a good model for a wide range of physical phenomenon, including the thermal noise in electronics.

The CO sensor that we use produces a very faint signal, which makes it vulnerable to ambient noise (e.g. the sensors receive and amplify this noise over the air), including AC power hum. Figure 2(a) shows the noise deviation from the mean of readings from the sensor before and

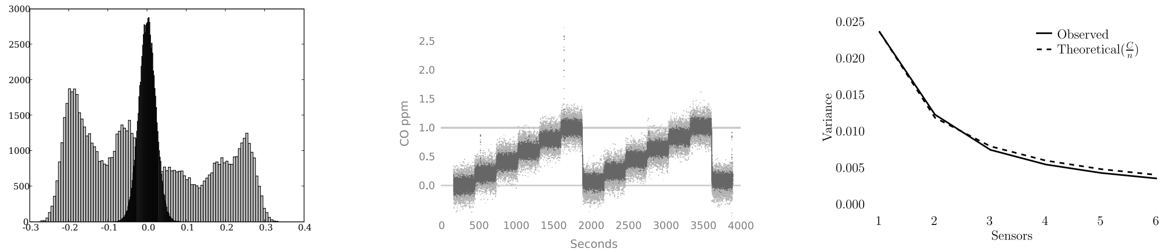

(a) Senor noise before and after filtering with a 60Hz (and harmonics) notch filter

(b) Sensor readings of concentration of CO(ppm) vs. time. Light dots show the readings from a single sensor. Dark dots show the average of six sensors.

(c) Variance of the average signal from a set of sensors (in ppm) vs. the number of sensors in the set. $\frac{C}{n}$ is show for reference.

Fig. 2. Empirical results with our CO sensors and test chamber.

after the 60Hz hum and its harmonics were removed using notch filters. The filtered noise is Gaussian, providing some empirical justification to assume a Gaussian noise process.

III. Empirical results

As the density of sensors at a given location increases, we can increase our precision by super-sampling, and averaging. For sensors with Gaussian noise (which our CO sensors exhibit) sampling in the same location, we expect the variance of the signal to be $\frac{C}{n}$ if we average the signals from n sensors with noise variance C. Note that when the noise is not Gaussian, the noise power will still decrease, but at a slower rate.

In Figure 2(b), we a experiment with six sensors in a chamber in which we can control the concentration of CO. In this case, we stepped the concentration of CO by 0.2ppm increments over an hour, and observed the response of the sensors. The light dots show the response of one sensor, and the dark dots show the averaged response of six sensors. Clearly the noise variance has decreased. Figure 2(c) show the variance of the signal versus the number of sensors averaged. The empirical results match the theoretical results closely!

IV. Gaussian processes

Using Gaussian process regression (GPR), we can also increase the precision of the system even when samples are not in the same location in space-time (a more realistic situation). The closer the samples are to one another, the greater the increase in the precision.

We should note that a GPR is appropriate not only because the *sensor noise* is Gaussian, but because process by which concentrations of gas mix and vary is also often modeled as Gaussian [3]. Modeled this way, we have the sum of two Gaussians, which is itself a Gaussian. More complex models might include inference of prevailing winds as well, but it remains to be seen if these complications are in fact necessary.

Gaussian process regression is a kernel method, and as such, shares many similarities with other kernel methods such as support vector machines (SVM). It is beyond the scope of this paper to describe the mathematics of GPR. Depending on the kernel, GPR can be as computationally efficient as SVM [4].

V. Learning curves

The amount that the precision of the system increases depends on the density of sampling. As the density of sampling increases, so does the precision.

To quantify this increase in precision for a given algorithm, it is typical to consider the "learning curve" of the algorithm. The learning curve shows the deviation of the true values of samples from the inferred function as the number of training examples increases for a given area. Sollich [5] provides some reasonably tight analytical bounds on the learning curves for GPR. In the future we will present an analysis of the learning curves under various model assumptions.

In Figure 3, we see simulation results in which the variance of the signal at a point decreases when nearby sensor's readings are also taken into account. In this simulation, we use a standard radial basis kernel, and the sensors are uniformly distributed within twice the scale of the kernel. This means that many of the points will be relatively far away from the point of interest, and will not contribute significantly to reducing the variance. Nonetheless, we can see that as the density near the point of interest increases, the variance decreases.

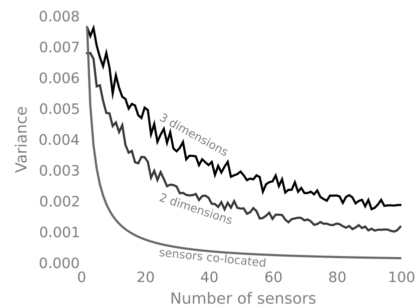

Fig. 3. Simulation of signal variance at a point when samples from different nearby sensors are also utilized vs. the number of nearby sensors. Variance is shown for two dimensional and three dimensional coordinates. For comparison, the variance is show for the case in which all of the sensors sample at the same point in space, as in Figure 2(c).

VI. Future work

This paper begins to explore one way in which mobility in sensors can be exploited to increase the usefulness and (in this case) precision of the sensing system. Although it examines super-sampling under (mostly) ideal situations, many questions remain to be answered. How does miscalibration impact these results? How do deviations from the Gaussian noise model impact the learning curves of the algorithms? How accurately can the system parameters be calibrated, and how does that impact precision? Is the (approximated) radial basis kernel the most appropriate covariance function? How should increased sample density be traded off with sampling in under-sampled locations, give limited resources to transmit samples?

Although we have also made some initial theoretical progress in automatically calibrating the bias of sensors in the sensing system using Gaussian process models [6], many questions also remain in this area. How does the automatic calibration hold up with a large, real data set. What is rate of drift of the calibration of the sensors? How much should we trade off calibrating vs. super-sampling? How can we infer the gain error of sensors?

Another significant obstacle to ubiquitous and personal sensing using mobile phones is obstruction of the airflow to the sensor (i.e. because the phone is in the user's pocket or purse). How can we detect this situation? Can it be compensated for, or do we need to discard samples taken in such a situation? In a related question, how can we detect indoor vs. outdoor environments. We have done some promising initial experiments using the microphone of the phone to classify the user's environment based on ambient noise, but these efforts need to be fleshed out.

Finally, many issues remain surround the end applications of the data. Can users be guided to safety in an emergency based on their position and the inferred position of a plume? How should data be visualized? How can it be anonymized while remaining sufficiently useful to various types of end users?

VII. Conclusion

Although many questions remain to be answered before we can build a working sensor system based on sensors integrated into mobile phones, we are encouraged by these results. We believe that mobile sensing has the potential to provide the platform for building the largest scientific instrument ever made: one with a dynamic range wide enough to construct an accurate image of the impact that humans have on their environment at a societal scale while also being able to examine an individual's exposure to a specific element at a specific place and time. Until now, no sensing system has been able to do this, and we believe that the potential benefits to society are enormous.

VIII. Acknowledgments

We would like thank the Intel CommonSense team for their help and support building prototype sensors. This work was partly funded by Intel Research and the National Science Foundation.

References

[1] AirNow, "AirNow: Quality of air means quality of life," http://www.airnow.gov/.

[2] R. Honicky, E. Brewer, E. Paulos, and R. White, "N-SMARTS: Networked Suite of Mobile Atmospheric Real-Time Sensors," in *Proceedings of Networked Systems for Developing Regions 2008*, 2008.

[3] M. R. Beychok, *Fundamentals of stack gass dispersion*. Milton R. Beychok, 2005.

[4] C. E. Rasmussen and C. K. I. Williams, *Gaussian processes for machine learning*. MIT Press, 2006.

[5] P. Sollich, "Learning curves for Gaussian processes," in *Proceedings of the 1998 conference on advances in neural information processing systems II*. Cambridge, MA, USA: MIT Press, 1999, pp. 344–350.

[6] R. Honicky, "Automatic calibration of sensor-phones using gaussian processes," EECS Department, UC Berkeley, Tech. Rep. UCB/EECS-2007-34, 2007.

An Implicit and User-Modifiable Urban Sensing Environment

Yasuyuki Ishida[1], Shin'ichi Konomi[2], Niwat Thepvilojanapong[1]
Ryohei Suzuki[3], Kaoru Sezaki[2] and Yoshito Tobe[1]

[1] Department of Information Systems and Multimedia Design, School of Engineering, Tokyo Denki University
2-2, Kanda-Nishiki-Cho, Chiyoda-Ku, Tokyo 101-8457, Japan
[2] Center for Spatial Information Science, the University of Tokyo, 5-1-5, Kashiwanoha, Kashiwa-shi, Chiba 277-8568, Japan
[3] Institute of Industrial Science, the University of Tokyo, 4-6-1, Komaba, Meguro-Ku, Tokyo 153-8505, Japan
Email: {yasu,wat}@u-netlab.jp, {konomi@iis,sezaki@iis,ryohei@mcl.iis}.u-tokyo.ac.jp, yoshito_tobe@osoite.jp

Abstract

Capturing useful data in a complex and dense urban space is an inherently challenging task. There is so much data people can capture in a city, yet they may fail to capture important information, some of which they don't even know that it exists. In this paper, we discuss an implicit approach to urban sensing and introduce an *implicit sensing system* that combines wearable sensors and mobile phone networks to support early-stage exploration of urban issues. An important technical issue that arose in the development of such sensing system is localization in indoor and urban canyon environments. We discuss the use of RFID tags as easy-to-deploy location reference points that could be installed, modified and reused by end-users. The evolution of *user-modifiable location infrastructure* should reflect and support implicit as well as explicit sensing that takes place in a city.

I. Introduction

Cities occupy just 2 percent of the Earth's surface, yet their inhabitants already consume 75 percent of the planet's natural resources for goods and services, and 80 percent of global carbon dioxide emission originate in towns and cities [4]. To improve city inhabitants' collective ability to understand the invisible impact of what they do in their everyday life on the global environmental issues, we could exploit personally-owned, location-aware sensors. In particular, GPS-equipped mobile phones together with portable/wearable sensors would allow urban pedestrians to easily capture various kinds of geographically indexed data, which could be voluntarily shared and aggregated to create sensor-powered maps [10] and support *participatory urbanism* [11] as well.

In this paper, we focus on two major limitations of mobile phone-based sensing by pedestrians. First, mobile phones often require users to *explicitly* capture data, which may not be a problem in goal-oriented sensing campaigns [2] that take place *after* environmental concerns arise. However, it is difficult to motivate capture *before* such concerns arise. Second, we need detailed location information to meaningfully interpret and use the sensed data; however, GNSS (Global Navigation Satellite System) technologies such as GPS do not work well in urban canyons as well as indoor and underground spaces.

We developed prototype mechanisms for *implicit* sensing and *user-modifiable* localization infrastructure, which together can alleviate these limitations to empower citizens. *Implicit* sensing is embedded in our everyday activities whose primary goal is not necessarily data collection. Implicit sensing can collect data *before* citizens perceive the need of focused observation. Such data can support early-stage exploration of urban issues.

Participatory Sensing [2] has developed tools and infrastructure for enabling public *campaigns* using networked mobile devices and sensors. Participatory Sensing requires users to *explicitly* register for the campaigns. Opportunistic Sensing [3] is also a system of people-centric urban sensing. In Opportunistic Sensing, users are required explicit procedure for selecting their interest. While the both Participatory Sensing and Opportunistic Sensing require explicit procedure of users, our proposed method is for *implicit* sensing which can extract previously unnoticed and unobservable atmosphere. Nokia's Sensor Planet [1] also proposed a platform for collecting and sharing sensor data for human centric sensing. However, for urban sensing, we need to consider about how to analyze and extract the meaningful information. EQUATOR e-science project [8] mapped carbon dioxide levels using mobile sensors. Our proposed system considered to extract environmental information by sensing human body as well as directly sensing the environment.

We first examine the data from a preliminary indoor experiment using pressure-aware slippers [15]. This informs the iterative design process of the WINFO+ system [5][12], which allows city-wide *implicit* sensing in the wild. Our preliminary experiments with WINFO+ suggest that such footwear-based sensing can reveal interesting information provided that there is a pervasive location infrastructure that 'seamlessly' covers a city. This leads to the discussions on user-modifiable, decentralized localization infrastructure for urban sensing, which can be developed by extending and integrating our RFID-based localization system [14]. We believe that WINFO+ together with the *user-modifiabl*e localization infrastructure allows people to collect meaningful data in a city.

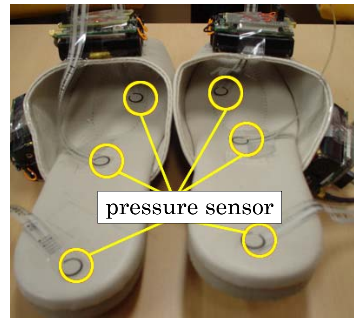

Fig. 1: Sensor-enabled slippers.

II. IMPLICIT SENSING: THE CASE OF SENSOR-ENABLED FOOTWEAR

To better understand the challenges and implications of implicit sensing, we have embedded networked sensors in footwear. We began by analyzing the data from pressure sensor-enabled slippers [15] focusing on pressure distribution and its correlation with a person's walking patterns. The prototype integrates normal slippers, Crossbow MICAz Motes, and three pressure sensors to wirelessly send pressure data to a server. The server then performs relevant signal processing. The three pressure sensors are embedded at the front, the center, and the back on the surface of the slippers (see Fig.1). We asked our subjects to walk with this prototype, and identified distinct signal patterns for the normal, shuffle, and forward-bending walking.

We call the period during which a pedestrian's foot contacts the ground an *epoch*. By extracting peak values from the front and rear sensors within an *epoch*, we can closely examine what goes on within each *epoch* and classify *epochs* into the following four groups:

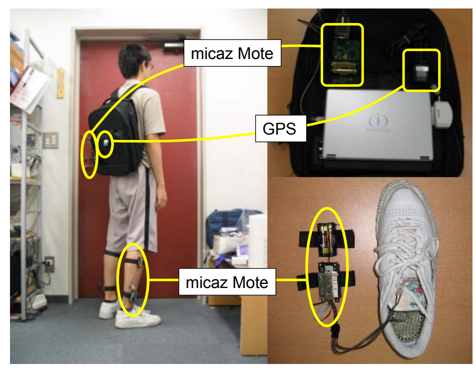

Fig. 2: Prototype of WINFO+ Client.

Group A: Peak values from the front and rear sensors are high. This suggests smooth movement in normal walking.
Group B: Peak values from the front and rear sensors are low and high, respectively. This suggests shuffle walking.
Group C: Peak values from the front and rear sensors are high and low, respectively. This suggests forward-bending walking.
Group D: Peak values from the front and rear sensors are low.

Overall, our basic data analysis suggests that small inexpensive sensors, if integrated in footwear, can capture what mobile phone-based sensors cannot easily capture. Interestingly, footwear devices can capture data without requiring a user to *explicitly* perform data capture operations. However, footwear-based *implicit* sensing is different from surveillance as the sensing is carried out through the users' personal devices and they must be able to fully control the process of (not) capturing, storing, and disclosing data.

The idea of integrating shoes and sensors [9] is not new. However, a city-wide urban sensing requires a durable and easy-to-use device, scalable and adaptive system architecture, and reliable positioning infrastructure that works both indoors and outdoors. Based on the basic analysis, we developed a footwear sensing system called WINFO+ [5]. It is based on a client-server model and composed of WINFO+ Client (WIC) and WINFO+ Server (WIS).

A WIC is a wearable device that consists of a personal computer, "probe shoes," a GPS receiver, and wireless interface (see Fig. 2). The personal computer wirelessly obtains the pressure data from the shoes. The data are tagged with the GPS timestamp and compressed by using the four epoch types. WICs then transmit the data to a

WIS, along with the latitude, longitude, epoch type, and timestamp information. In our prototype, WICs can communicate with a WIS virtually anywhere in a city by using the PHS (Personal Handy-phone System) technology.

WINFO+ is designed for *adaptive sensing* in diverse device, information and resource environments. WICs should acquire the right amount of information depending on their screen size and CPU power (*device adaptive* sensing). Also, WINFO+ should respond to dynamic behavior of data (*information adaptive* sensing). For example, we might need finer-grained data when the data change substantially either in temporal or spatial axes. Moreover, the system should be able to determine the frequency of sensing and transmission depending on the amount of battery power left (*resource adaptive* sensing).

The WIC prototype is easy to wear and designed to look socially acceptable in most public spaces. Using the prototype, we carried out small-scale experiments in real urban spaces. Three male graduate students wore the prototype and walked in a Central Tokyo area near Akihabara without drastically changing walking styles: they walked spontaneously along a street, crossed the street by using a pedestrian bridge, and stopped at a train station. Figs. 3 and 4 show sample data from one of the subjects. This experiment showed that footwear-based city-wide sensing can reveal characteristics of the surfaces on which pedestrians walk as well as wearers' walking habits. We also acknowledged the importance of location information in interpreting and using these sensed data.

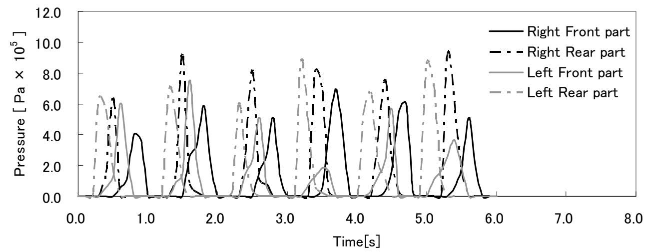

(a) Flat areas.

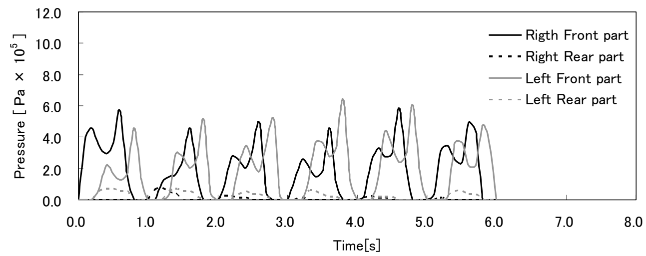

(b) Stairs (upward).

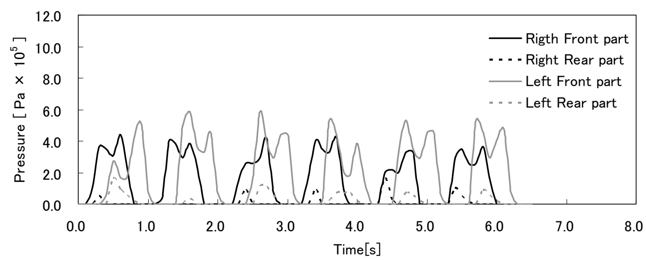

(c) Stairs (downward).

Fig. 3: Temporal change in pressure.

III. USER-MODIFIABLE LOCALIZATION INFRASTRUCTURE

GNSS technologies such as GPS do not work well in urban canyons, indoor/underground spaces, and so on. This can be problematic when people want to collect location-relevant sensor data in such spaces. As demonstrated in the following scenario, RFID tags can be used as location reference points for urban sensing, thereby complementing GNSS technologies:

> *Imagine an apartment complex that may have an air contamination problem. The residents can install RFID tags at their front doors to help citizen scientists collect air quality data in the building. Mobile sensor devices can obtain unique IDs from the tags, and then retrieve corresponding 3D location information by querying a database. The tags could also be used for implicit sensing. For example, mailmen's shoes capture and accumulate location-indexed pressure and temperature data in the building over a year, which may be later found useful to discuss remodeling of the building for elderly people.*

An important issue here is the motivation and the deployment cost to physically install the tags, measure their positions, and update the database. WiFi-based localization [7] require little deployment cost only when WiFi stations are already deployed in the environment. In contrast, RFID-based localization uses inexpensive RFID tags that can be easily deployed on demand at a wide variety of places.

In Japan, the government has shown keen interest in RFID location reference points [14] and already embedded about a hundred "intelligent benchmarks," which are equipped with passive RFID tags, in the city of Kobe. We surely need much more RFID reference points to fully cover a city-wide area: perhaps, millions of them (e.g., at 10-meter intervals).

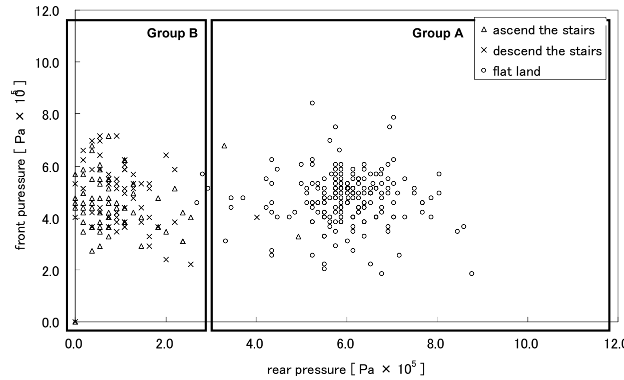

Fig. 4: Front-rear diagram.

Although government-initiated centralized deployment can be heavyweight and costly, they can hire professional land surveyors who have the skills to install high quality reference points in terms of physical robustness and information accuracy. An alternative approach is the citizen-initiated decentralized deployment that is more scalable in terms of the number of tags. We envision a hybrid, *user-modifiable* environment in which a small number of strategically allocated quality-assured tags (T1) and a large number of end-user tags (T2) coexist. In such a user-modifiable environment, we can reduce the overall deployment cost by reducing (1) the number of tags that must be installed and (2) the cost to install each tag.

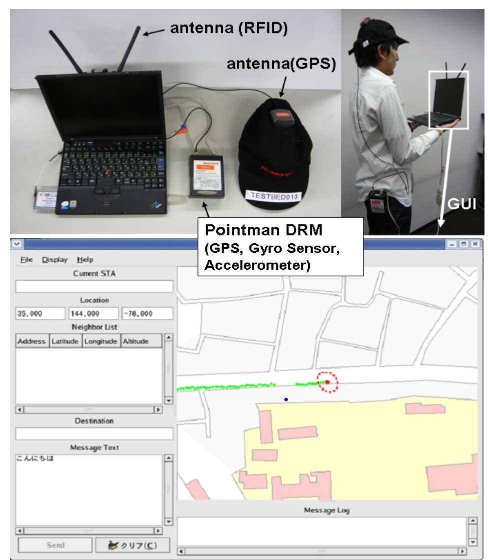

Fig. 5: P2P-based localization system.

As shown in Fig. 5, we have developed a P2P-based localization system that reduces the number of required tags. The pedestrian device can estimate its position using GPS, (active) RFID location reference points, dead reckoning modules (Honeywell GyroDRM™), and location information shared by colocated pedestrians. Research [6] shows that the combination of RFID, GPS, and dead reckoning can improve positioning accuracy in both indoor and outdoor environments even without such location information sharing. Our system uses GPS if the satellite signals are available. Otherwise, the system operates without GPS by obtaining location information from a nearby RFID tag. Even when the user's device is away from RFID tags, it can estimate the position by using dead reckoning modules. However, as pedestrians move and time passes by, the positioning error increases. In our positioning mechanism, colocated devices exchange their location estimation (along with relevant error estimation) with each other in order to cooperatively reduce the positioning error considering human mobility patterns. We carried out an experiment in a *54m x 63m* space on a university campus and verified the effectiveness of the cooperative location estimation.

To reduce the cost for installing each tag, we have developed a mechanism that automatically estimates the position of a newly installed tag by collecting location information from pedestrians who pass by the new tag [13]. This mechanism allows people to simply put a tag without manually updating the database. We developed a prototype and tested it on a university campus, and found that the location estimation error of a new tag quickly decreases and stays below *2 meters*.

These mechanisms together can support the ecology of location reference points by facilitating end-user deployment. As our scenario may suggest, existing environmental concerns could motivate end-users to install location reference points. However, implicit urban sensing without clear value proposition may not directly motivate end-user installation. We would like to understand end-user installation from an ecological perspective rather than a narrow scope of cost-benefit balance. For example, people may be able to reuse location reference points that were installed for some other purpose. Such practices cannot be prescribed, but could be facilitated by technological and social systems.

IV. DISCUSSION AND CONCLUSION

The combination of WINFO+ and the user-modifiable localization infrastructure enables implicit city-wide sensing that can reveal characteristics of ground surfaces as well as walking habits. Note that mobile phone clocks can provide a means to tag sensed data with timestamps when GPS is not available. Our approach complements mobile phone-based *explicit* sensing and allows people to capture some data *before* environmental concerns arise. Similar approaches could be used for other kinds of wearable sensors (e.g., heart rate, temperature, moisture, and blood pressure sensors).

Allowing for participatory contribution of both sensor data and location reference points can create exciting opportunities; however, it can also introduce issues around data quality, security, and privacy. How should we deal with variability of data quality, and "junk reference points" that could degrade the localization accuracy? Implicit sensing unobtrusively collects data from users' body areas as well as spaces inhabited by people. Capturing and sharing such data could cause serious privacy problems. We are currently exploring different approaches to address these issues.

WINFO+ exploits mobile phone communication (Personal Handy-phone System) for disseminating sensor data. Advances in mobile phone technologies may allow tighter and flexible integration of wearable sensing devices and mobile phones in the future. In addition, an increasing number of phones are integrated with RFID, Bluetooth, and 2D barcode technologies. We can design location reference points that exploit these technologies so as to complement GPS and facilitate location-indexed data collection by mobile phones.

Finally, it is important to acknowledge that supporting participatory urban sensing is more than just creating easy-to-capture, easy-to-share, and easy-to-modify environments. People need to have information and skills in order to meaningfully participate in collaborative sensing and sensemaking. This motivates a future work to design an integrated support environment for urban data practices.

REFERENCES

[1] E. Balandina & D. Trossen, "Nokia remote sensing platform middleware and demo application server: futures and user interface," Nokia Research Center, Helsinki, 2006.

[2] J. Burke, D. Estrin, M. Hansen, A. Parker, N. Ramanathan, S. Reddy & M.B. Srivastava, "Participatory Sensing," *SenSys'06 Workshop on World-Sensor-Web*, 2006.

[3] A. Campbell, S.B. Eisenman, N.D. Lane, E. Miluzzo & R.A. Peterson, "People-Centric Urban Sensing," *SenSys'06 Workshop on World-Sensor-Web*, 2006.

[4] Cities Alliance, *2006 Annual Report*, http://www.citiesalliance.org/publications/annual-report/2006-annual-report.html (Accessed: July 29, 2008).

[5] Y. Ishida, R. Suzuki, K. Sezaki, N. Thepvilojanapong & Y. Tobe, "Demo Abstract: WINFO+: Extracting Environmental Information Using Walking Signals," *Adjunct Proc. of Pervasive 2008,* pp.107-110, 2008.

[6] Kourogi, M., Sakata, N., Okuma, T., & Kurata, T., "Indoor/Outdoor Pedestrian Navigation with an Embedded GPS/RFID/Self-contained Sensor System," Proc. ICAT 2006, pp.1310-1321. Springer, 2006.

[7] A. LaMarca, Y. Chawathe, S. Consolvo, J. Hightower, I. Smith, J. Scott, T. Sohn, J. Howard, J. Hughes, F. Potter, J. Tabert, P. Powledge, G. Borriello & B. Schilit, "Place Lab: Device Positioning Using Radio Beacons in the Wild," *Proc. Pervasive 2005*, 2005.

[8] Milton, R. & Steed, A., "Mapping carbon monoxide using GPS tracked sensors," *Environmental monitoring and Assessment, vol. 124*, pp. 1-19, 2007.

[9] Morris, S.J. & Paradiso, J.A., "Shoe-Integrated Sensor System for Wireless Gait Analysis and Real-Time Feedback," Poroc. EMBS/BMES Conf., pp.2468-2469, 2002.

[10] S. Nath, J. Liu & F. Zhao, "SensorMap for Wide-Area Sensor Webs," *IEEE Computer Magazine, 40*(7), pp. 90-93, 2007.

[11] E. Paulos, R.J. Honicky, & E. Goodman, "Sensing Atmosphere," *SenSys'07 Workshop on Sensing on Everyday Mobile Phones in Support of Participatory Research,* 2007.

[12] K. Sasaki, U. Inoue & Y. Tobe, "WINFO: A human-assisted sensor network," *The second International Workshop on Networked Sensing Systems (INSS 2005)*, pp.186-190, 2005.

[13] O. Sangratanachaikul, S. Konomi & K. Sezaki, K., "An Easy-to-Deploy RFID Location System," *Adjunct Proc. of Pervasive 2008 (Late Breaking Results)*, pp.36-40, 2008.

[14] K. Sezaki & S. Konomi, "RFID-Based Positioning Systems for Enhancing Safety and Sense of Security in Japan," *Proc. Int'l Workshop Ubiquitous Pervasive and Internet Mapping (UPIMap 2006),* pp.194-200, 2006.

[15] Y. Uehara, T. Uchiyama, M. Mori, H. Saito, & Y. Tobe, "Always-on karte: a system for elderly people's healthcare using wireless sensors," *The third International Conference on Networked Sensing Systems (INSS 2006)*, pp.45-48, 2006.

Evaluating the iPhone as a Mobile Platform for People-Centric Sensing Applications

Emiliano Miluzzo, James M. H. Oakley, Hong Lu, Nicholas D. Lane,
Ronald A. Peterson, Andrew T. Campbell
Computer Science Department, Dartmouth College, Hanover, NH 03755, USA

Abstract

A number of mobile phones such as the Nokia N95 and Apple iPhone are being used by researchers to explore new people-centric sensing applications. These top-end phones include various sensors (e.g., accelerometer, proximity sensor, GPS, camera, microphone), radios (e.g., Bluetooth, WiFi, cellular), operating systems (e.g., Symbian, customized Mac OS X), and processors (e.g., 330 Mhz ARM, 412 Mhz ARM). While there is some data on the Nokia N95, little, however, is known about the ability of Apple's iPhone to support the necessary sensing, processing and communications needs of these emerging sensing applications. We present the first quantitative performance evaluation of the iPhone's sensors, localization engine, and networking stack while running CenceMe, a representative people-centric sensing application. We profile the performance of CenceMe running on the iPhone in terms of energy consumption and computational speed of its algorithms. One drawback of using the iPhone over the N95 is that it does not allow third party applications such as CenceMe to run as a background process, making continuous sensing problematic. The upside is that the iPhone offers a rich UI architecture, high computational capability, and an efficient application distribution system through Apple's App Store.

I. Introduction

New top-end mobile phones such as the Nokia N95 [7] and Apple iPhone [8] are enabling a new class of mobile, people-centric applications [1] [2] [3] [4] to emerge. While researchers have demonstrated that the Nokia N95, a Symbian-based mobile phone, can efficiently host sensing applications [2] [3] [4], little is known about the ability of the Apple iPhone to support people-centric sensing applications. Recently, Apple opened up to the deployment of third-party applications with the release of the iPhone SDK in March 2008. This, combined with its ease of use, rich UI, and efficient application distribution system through the Apple App Store makes the iPhone an appealing platform for development of new mobile applications. A natural question for our community is what are the trade-offs when implementing and deploying a sensing application using the iPhone; more specifically:

- How easy is it to program a sensing application on the iPhone?
- What are the pros and cons of the iPhone in comparison to other sensor capable mobile platforms?
- What is the energy profile when the iPhone's sensors, WiFi and cellular radio are involved in realizing the application?
- What is the processing performance of the iPhone when running signal processing algorithms such as fast fourier transform, a common tool used to interpret audio and accelerometer sensor data?

We address these questions in this paper. While the presentation of our results is limited due to space, we provide a short qualitative comparison of a number of devices used for mobile sensing including the Apple iPhone, Nokia N95, and Intel Mobile Sensing Platform (MSP) [11]. The main contribution of this paper is the observations and insights when running CenceMe, a representative people-centric sensing application [4], on the iPhone. Specifically, we quantitatively evaluate the iPhone's computational capability, energy profile, and localization accuracy. We believe this study will be useful to the growing community of iPhone developers, particularly, those interested in building people-centric sensing applications.

II. Comparison of Mobile Sensing Platforms: iPhone, N95, and MSP

In what follows, we present a short qualitative comparison of the Apple iPhone, Nokia N95 mobile phone, and the MSP. All these devices are actively being used in support of mobile sensing applications and systems development. The N95 is currently one of the top-end Nokia mobile phones equipped with an accelerometer and GPS, while the MSP is representative of the class of embedded devices used for human activity recognition research. A simple

TABLE I: Platform comparison for the Apple iPhone, Nokia N95, and Intel MSP

	iPhone	**Nokia N95**	**Intel MSP 430**
Processor	412 MHz ARM	330 MHz ARM	416 MHz Xscale
RAM	up to 70 MB	up to 128 MB	256 KB
ROM	20 MB	up to 160 MB	32 MB
Storage	up to 8GB/16GB	min-SD card (up to 8 GB)	mini-SD card (up to 8 GB)
Sensors	3-axis accel, mic, GPS	3-axis accel, mic, GPS	3-axis accel, mic, light, barometer, temp, IR, humidity, compass
Radio	WiFi	WiFi, Bluetooth	Bluetooth, Zigbee

comparison of some of the technical details of the three devices is reported in Table I. As shown in Table I, all three platforms present similar computational capabilities given similar processors, and large storage and ROM size. The RAM on the MSP is much smaller than on the iPhone and N95, which first and foremost are designed as mobile phones, hence the need to handle multiple processes at the same time including graphics computation. The MSP's short-range radio technology is flexible allowing the implementation of advanced pairing algorithms between nodes while the use of the iPhone and N95's short-range radio is limited to simple neighbor interactions. The main difference between the three devices is represented by the sensing capability; specifically, the MSP outshines both the iPhone and the N95 in terms of the number of available sensors. This is not surprising given that the MSP is an embedded purpose-built platform for activity recognition. However, even with a reduced set of on board sensors, the iPhone and N95 are powerful devices and capable of inferring human activites - for example, we have implemented the full featured CenceMe application on the N95 [4] as well as a version on the iPhone [10]. Providing mobile phones with more sensing capabilities (e.g., gyroscope) would greatly enhance the humans presence classification accuracy given the broader input to the classifiers feature vectors.

III. Programmability Characteristics of the iPhone

In what follows, we analyze the programmability characteristics of the iPhone. Any third-party application is handled by the iPhone OS using a sandboxing model which does not allow the application to access some of the iPhone functionalities (such as WiFi APIs or iTunes) for security reasons. A simplified version of a SQL database, namely *sqlite* [14], designed to run on resource constrained environments, is also supported as a means to ease application on-the-phone storage.

By following a systematic approach, we intend to answer the following question: what are the positive and negative aspects of the iPhone as a programmable platform? Although the iPhone presents a rich set of features making it potentially a good platform for the development of sensing applications, the iPhone SDK also provides some barriers in its current stage of development (i.e., iPhone SDK for iPhone OS 2.2). In what follow, we briefly discuss the pros and cons of the current iPhone development environment.

Advantages:

- *Programming Language*. The iPhone is programmed in Objective-C [12]. Objective-C is a superset of the C language, with some object oriented programming features. The advantage of Objective-C over other languages such as Symbian C++ adopted by Nokia, is that it is a quite simple language to learn and use. The iPhone APIs and emulator (which runs on desktop/laptop machines) make programmability, UI design, and code debugging an efficient process for developers.

- *APIs*. The APIs are well designed and documented, abstracting the developer from low level components. For example, the location engine API returns data transparently to the user regardless of whether the location coordinates come from WiFi, cellular triangulation, GPS, or a combination of sources. In addition, the accelerometer and microphone APIs are cleanly designed and make accessing these devices simple and strightforward to use.

- *Indoor Localization*. By using WiFi [6] and cellular triangulation to determine the location, the iPhone localization for indoor spaces is quite accurate, as discussed in Section IV. This is an important feature, for example, for mobile social networking applications considering that people spend a large amount of time indoors.

- *User Interface*. The iPhone experience is greatly enhanced by the Cocoa-Touch layer architecture [9] that provides for a good user experience. Combined with a powerful graphics framework, this makes the iPhone UI one of the best presentation layers of any mobile devices.

- *Application Distribution*. Apple provides an efficient way to distribute third-party applications to the public through the App Store [13]. Once an application is tested and approved by Apple, the application is posted on the App Store. After that the application can be downloaded and automatically installed on any iPhone.

Disadvantages:

- *Lack of Background Mode*. The main drawback of the iPhone is the lack of background mode to run third-party applications. This is enforced by Apple for security reasons. This means that anytime the phone goes into sleep mode or the user launches another application, the currently running third-party application is terminated. As a result of this design decision, sensing applications cannot provide continuous sensor data feeds. Therefore, applications can only generate intermittent data streams. This limits the effectiveness of continuous sensing application such as CenceMe. Apple's response to the lack of background capability is the *Push Notification Service* coming in the next releases of the SDK. With the push notification service, probes can be sent by the Apple backend servers, which, in turn, serve as relays for push messages sent by a sender host to a receiver iPhone. As the receiver iPhone is woken up by the probe the user is asked by the iPhone OS whether to let the application run in response to the probe message or not. It is worth noting that the Nokia N95 and Intel MSP support background mode and therefore support the implementation of continuous sensing applications.

- *Short-Range Radio API Limited Capability*. Currently, it is not possible to access directly the Bluetooth or WiFi radio stack APIs on the iPhone. The only way to exchange information between neighboring iPhones is by using the iPhone networking stack via the Bonjour service through WiFi. The short-range interactions of devices via this networking capability is therefore very limited and does not allow developers to build sophisticated pairing protocols.

IV. Performance Evaluation

In this section, we report some initial results from a number of experiments aimed at evaluating the iPhone computational capability, battery duration, and localization accuracy by using the original iPhone model (without GPS) and the new iPhone 3G (with GPS) running, respectively, iPhone OS 2.0 and 2.1.

Computational Capability. In order to evaluate the processing power of the iPhone we run a fast fourier transform (FFT) algorithm, which is part of the CenceMe software suite, and measure the iPhone computation time. The FFT computation evaluation is performed during normal CenceMe usage process [10]. The FFT implemented as part of the CenceMe application is the Kiss FFT [15], a well known open source high performance FFT library. We choose the FFT as a means to evaluate the iPhone under high computational load because the FFT is a common tool used in inference techniques applied to sensor data such as accelerometer and audio data streams. As shown in Figure 1, the iPhone computation time up to 4096 FFT points is below 60 msec even for a large number (i.e., 60000) of sampled events in time. Many sensor data analysis algorithms make use of 512 - 2048 FFT points calculation, which means that they could efficiently leverage the iPhone's computational power. Large data bins in time, up to 60000 samples in our experiment, could also be quite efficiently handled by the FFT on the iPhone in at most 200 msec.

Battery Lifetime. We perform some experiments to quantify the battery drain of the iPhone when running CenceMe compared to the baseline power usage without CenceMe. We set the screen saver to off so that the phone never goes into standby mode. The battery duration for different data upload rates when CenceMe is running is compared to the duration when CenceMe is not running, as shown in Figure 2(a). With the phone's standby mode off and running CenceMe continuously, the battery lifetime spans between 4 hours and 37 min to 7 hours according to the upload rate. We then turn the screen saver back on and set it to 5 minutes and run CenceMe with the following cycle: run for 5 minutes, let the screen saver go off, leave the screen saver up for 5 minutes, wake the phone up for 5 minutes, and so on, until the battery discharges completely. In this way, for the same upload rates, we obtain a phone usage time (meaning time available to operate CenceMe) between 4 hours 50 min and 5 hours 20 min. (Note, the battery maximum usage is between 10-11 hours). This battery duration is similar to the duration obtained with iPhone usage patterns comparable to the one of our experiment when running different applications than CenceMe. This is because the prevalent battery drain is due to the iPhone LCD screen rather than the networking activity for data transmission/reception operated by CenceMe.

Localization Accuracy. To evaluate the localization accuracy of both the old model iPhone (without GPS) and the

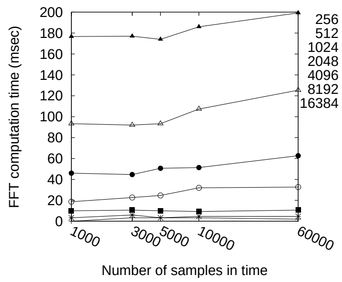

(a)

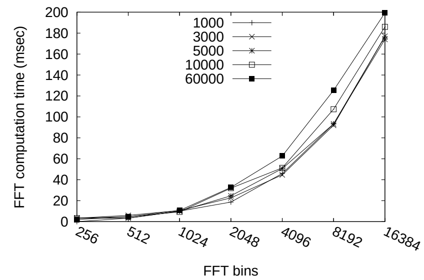

(b)

Fig. 1: FFT computation time as a function of (a) the number of samples in time while varying the FFT bin size (as shown in the legend) and (b) the FFT bin size while varying the number of samples in time (as shown in the legend).

TABLE II: Localization accuracy for different places in the Dartmouth Campus - Legend: **C.S.** = Cellular Signal; **A** = Old iPhone localization accuracy (m); **B** = iPhone 3G localization accuracy (m); **C** = Garmin GPS accuracy (m); **D** = Old iPhone-Garmin GPS localization difference (m); **E** = iPhone 3G-Garmin GPS localization difference (m)

	Location	**WiFi**	**C.S.**	**A**	**B**	**C**	**D**	**E**
1	Computer Science Bld indoor	good	good	83	22	N/A	N/A	N/A
2	Computer Science Bld outdoor	good	good	44	17	14	29	36
3	Library outdoor	good	good	17	9	8	0	1
4	Library indoor	good	mediocre	13	5	N/A	N/A	N/A
5	Golf course	none	good	759	17	5	45	1
6	Engineering Bld	weak	weak	95	17	5	14	0
7	Main St.	none	weak	179	47	11	5	4
8	The Green	none	good	323	47	5	24	2

iPhone 3G (with GPS) we carry out the following experiment: we walk in the Dartmouth College campus with both the iPhone models and a Garmin eTrex GPS device. We record the geographical coordinates from the Garmin and the iPhone devices at eight different locations. Eight clusters are shown on the maps in Figure 2(b) and Figure 2(c), each cluster indicating the location manually tagged by the person carrying the devices, the location reported by the Garmin, and the location indicated by the old model and 3G iPhones. The old model iPhone's localization engine uses WiFi [6] and cellular triangulation. Therefore, when WiFi and/or the cellular coverage is poor the resulting localization accuracy is low. This can be seen for locations associated with clusters 5 and 8 where there is poor WiFi and/or cellular coverage. In case of clusters 1 and 4, which are indoor locations where GPS performs poorly, the iPhone localization is more accurate given the high quality of the WiFi and cellular signal. The old iPhone model estimates an accuracy between 13 and 759 meters, as shown in column A of Table II. Dartmouth College is located in the small college town of Hanover, NH, and subquently not served by many cell towers. Clearly, the availability of more cell towers would allow the iPhone's localization triangulation algorithm to be more accurate. The actual distance difference between the old iPhone and Garmin GPS reported locations, as shown in column D of Table II, is 45 m at most, indicating that the iPhone localization algorithm uses a conservative approach to estimate its accuracy. The GPS boosts the localization accuracy of the iPhone 3G, being particularly effective where there is a lack of WiFi coverage or when the cellular signal is poor. This can be seen from columns B and E of Table II where, respectively, the error estimated by the iPhone 3G and the localization difference between the iPhone 3G and Garmin GPS are reported. It is evident how the iPhone 3G-Garmin GPS localization difference is smaller than when using the old iPhone model.

V. Related Work

There is a growing body of work on the evaluation of sensing platforms for embedded sensing systems. For example [16] discusses the technical details and performance evaluation of the Telos motes, a widely used sensing platform in wireless sensor networks research. With the increasing interest in people-centric sensing applications, where sensing devices are carried by individuals, new sensing platforms are being developed, evaluated, and reported

(a) (b) (c)

Fig. 2: (a) Battery duration with and without CenceMe running while the iPhone screen saver is set to off. - Localization accuracy for eight different locations in the Dartmouth campus of (b) the old iPhone (no GPS), and (c) the iPhone 3G (with GPS).

in the literature. While the Intel MSP is used for activity recognition [11] there is little available in the literature on the evaluation of mobile phones for sensing. In [3] and [4] the authors present some early performance studies when using mobile phones for sensing applications. In [3], the authors show some energy profiling of the Nokia N95. In [4], the authors present a detailed evaluation of the N95 in terms of programmability and computational performance including detailed energy considerations.

VI. Conclusion

In this paper, we presented an evaluation of the iPhone running the CenceMe application. We quantitatively evaluated the iPhone's computational capability, energy profile, and localization accuracy. We showed that the computational capability of the iPhone is sufficient to handle high load FFT calculations. We also showed that iPhone's cellular and WiFi assisted localization outperforms pure GPS-based localization, as long as cellular and WiFi coverage are present. We believe this study is useful to iPhone developers, particularly those interested in building people-centric sensing applications.

Acknowledgments

This work is supported in part by Intel Corp., NSF NCS-0631289, and ISTS at Dartmouth College. ISTS support is provided by the U.S. Department of Homeland Security under award 2006-CS-001-000001, and by award 60NANB6D6130 from the U.S. Department of Commerce. The authors would like to thank Daniel Peebles for the useful insights provided to this work.

References

[1] A. Kansal, M. Goraczko, and Feng Zhao, Building a sensor network of mobile phones. In *Proc. of IPSN 2007*, April 25-27, 2007.
[2] E. Agapie, et al., Seeing Our Signals: Combining location traces and web-based models for personal discovery. In *Proc. of HotMobile 2008*, Napa Valley, CA, USA, February 25-26, 2008.
[3] Shravan Gaonkar, Jack Li, Romit Roy Choudhury, Landon Cox, Al Schmidt, Micro-Blog: Sharing and Querying Content through Mobile Phones and Social Participation. In *Proc. of MobiSys 2008*, Brekenridge, CO, USA, June 17-20, 2008.
[4] Emiliano Miluzzo, et al., Sensing Meets Mobile Social Networks: The Design, Implementation and Evaluation of the CenceMe Application. In *Proc. of SenSys 2008*, Raleigh, NC, USA, Nov. 5 - 7, 2008.
[5] Andrew T. Campbell, et al., The Rise of People-Centric Sensing. In *IEEE Internet Computing Special Issue on Mesh Networks*, July/August 2008.
[6] Skyhook Wireless. http://www.skyhookwireless.com/
[7] Nokia N95 series. http://www.nseries.com/index.html
[8] Apple iPhone. http://www.apple.com/iphone/
[9] Apple iPhone Dev Center. http://developer.apple.com/iphone/
[10] CenceMe portal. http://www.cenceme.org.
[11] Tanzeem Choudhury, et al., The Mobile Sensing Platform: An Embedded Activity Recognition System. In *IEEE Pervasive Computing*, April-June 2008, Vol. 7, No. 2, pp. 32-41.
[12] Objective-C. http://developer.apple.com/documentation/Cocoa/Conceptual/ObjectiveC/.
[13] Apple App Store. http://www.apple.com/iphone/appstore/.
[14] SQLite. http://www.sqlite.org/.
[15] Kiss FFT library. http://sourceforge.net/projects/kissfft/.
[16] Joe Polastre, et al., Telos: Enabling Ultra-Low Power Wireless Reasearch. In *IPSN/SPOT 2005*, April 25-27, 2005, Los Angeles, USA.

MobileSense - Sensing Modes of Transportation in Studies of the Built Environment

Jonathan Lester[1], Phil Hurvitz[2], Rohit Chaudhri[3], Carl Hartung[3], Gaetano Borriello[3]

[1]Department of Electrical Engineering, University of Washington, Seattle, WA 98195, USA
Email:{jlester}@u.washington.edu
[2]2Department of Urban Design and Planning, University of Washington, Seattle, WA 98195, USA
Email:{phurvitz}@u.washington.edu
[3]Department of Computer Science and Engineering, University of Washington, Seattle, WA 98195, USA
Email:{rohitc,chartung,gaetano}@cs.washington.edu

Abstract

We discuss a study we conducted using the Mobile Sensing Platform and GPS information to record the activities and locations of 53 subjects, each of whom collected data for one week. Using this data we are developing methods for combining activity inference from sensors to infer the mode of transportation, label significant locations, and extract trips. Although our data collection was conducted using a non-consumer sensor platform, we discuss how our methods can translate onto existing mobile platforms such as the iPhone and Nokia N95, and how these platforms will enable us to study larger populations to draw more concrete conclusions about the relationship between the urban environment and people's activities.

I. Introduction

As sensing platforms have advanced to become more usable and more readily available, the range of potential users and applications continues to expand. At the same time, commoditization of sensors in mobile phones has increased their availability and provided researchers with opportunities to study much larger populations than in the past. One area of interest in these sensors has been activity inference: the ability to tell what activities a person is performing based upon sensor information. Activity inference provides the 'what' of a user's context while location sensors (such as cell-tower/WIFI localization and/or GPS) provide the 'where'.

This 'what' and 'where' information can be used by a number of mobile phone applications, from physical fitness and health monitoring, to recommendation systems, to studying environment and personal behavior. In this paper we focus on the last application area, though the techniques we develop are applicable to a number of different applications. We worked with colleagues in the College of Architecture and Urban Planning's program on the Built Environment who are interested in studying the effect the urban environment has on the types of activities people engage in. Some of the questions these researchers are interested in are: what associations exist between different types of activities and different kinds of urban environments; what characterizes the land use patterns where people spend long periods of time; and do the properties of origins and destinations differ between short and long trips or trips made by different modes of transportation. To help answer some of these questions, we conducted a study to validate that sensors can be a useful tool for studies of the built environment, and to develop methods for performing these analyses on a larger scale. We collected one week's worth of sensor and GPS data from 53 participants. In all we gathered approximately 2,900 hours of data and ~2,900 surveys via the Experience Sampling Method (ESM) to provide ground truth. In the rest of this paper, we discuss the data collection, the methods we are developing to analyze the data, our validation for our method's usefulness, and our future work of implementing these methods on mobile phones to expand these studies to hundreds of users.

II. Related Work

Mobile device-centric sensing has become an active area of research in the past few years. Several projects in the

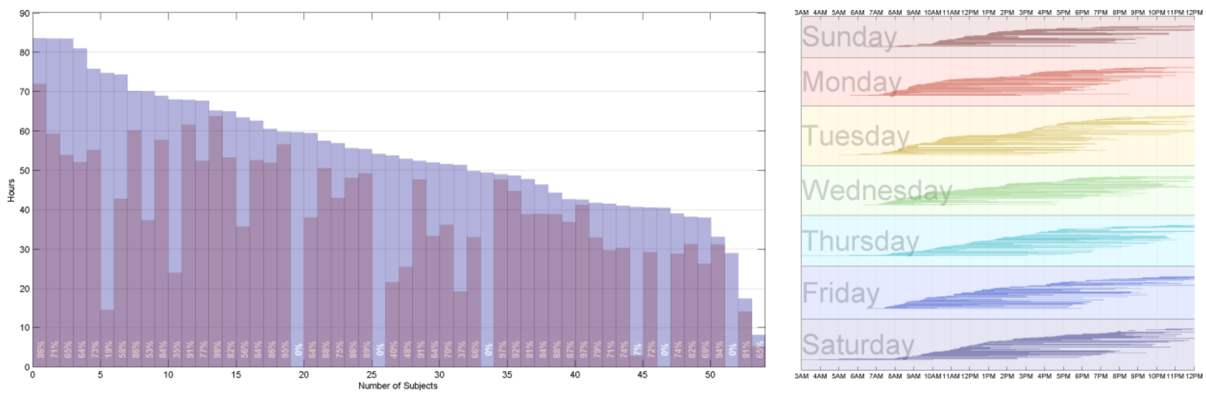

Fig. 1. (left) Data collected from each subject with the amount of time a GPS lock was available. (right) Each line indicates a trace of data collected from subjects showing good coverage throughout the day and days of the week.

academic research community and the industry are quite relevant to this work. Agapie et al. [5] present the Personal Environmental Impact Report (PEIR) system in which they leveraged mobile devices to collect time stamped location data and use this information to build models that estimate individuals' impact on environmental pollution levels. Miluzzo et al. discuss the CenceMe system in [6]. CenceMe leverages sensing capabilities of Mobile Devices to infer contextual information such as activity, availability, surroundings, etc., and injects this "Sensing Presence" into social networking applications like Facebook and MySpace. Liao et al. [8] used GPS data alone with a relational Markov network to infer the activities a person performs based upon the time of day, GPS trace, and their proximity to stores and restaurants.

The Smartraq study investigated the relationship between objectively measured physical actvity and objectively measured urban form; however, this relationship was based solely on subjects' home neighbrohood environment, rather than the complete spatial area of activity [7]. Mohan et al. discuss TrafficSense in [9], a framework in which they use sensors on mobile Smartphones to monitor road and traffic conditions, including mode of transportation.

III. Data Collection

For our data collection we wanted to collect GPS readings, accelerometer traces, and barometric pressure readings (which aid in detecting when people are moving up or down). Even though there are a few mobile phones such as the Nokia N95, and the new iPhone 3G, which have GPS receivers and accelerometers, these platforms are not yet perfectly suited to perform data collection necessary for our applications. While mobile phones excel in being robust, widely available, and relatively inexpensive devices; they often have restrictive security models, limited hardware access, poor documentation, and short battery lives. For example, the GPS receivers of the N95 and iPhone 3G are functional, but they do not provide data of an adequate quality for an initial experiment. These restrictions make it difficult to build data collection applications that need to run in the background for prolonged periods of time, a basic requirement for data collection in a research setting.

Instead we used the Mobile Sensing Platform (MSP) [3] which combines an Intel XScale processor with an accelerometer, barometric pressure sensor, light sensors, humidity sensors, microphone (not used in this experiment), GPS, and storage capacity. Subjects participating in the experiment were given an MSP inside a box with a belt clip, shown in Fig. 2, to be worn on their waist, and a cell phone for entering ground truth survey information.

Subjects were recruited to wear the device for a one-week period. During the week they would wear the MSP clipped on their waist and carry a Windows Mobile SMT phone wherever comfortable. The Smartphone was running the MyExperience ESM [2] sampling software, which would prompt them for a survey approximately every hour. These surveys asked questions about activity type, duration, purpose, and location (Fig. 2, middle). There were 53 participants with an average age of 32 and standard deviation of 11 years. Thirty-eight percent were female, and 86% were college graduates or had a post-graduate degree. On average each subject collected approximately 53 hours of data during a week. As subjects were allowed to turn off the device when they did not wish to be monitored, the amount of data collected varied among different subjects. Fig. 1 (left) shows a bar chart

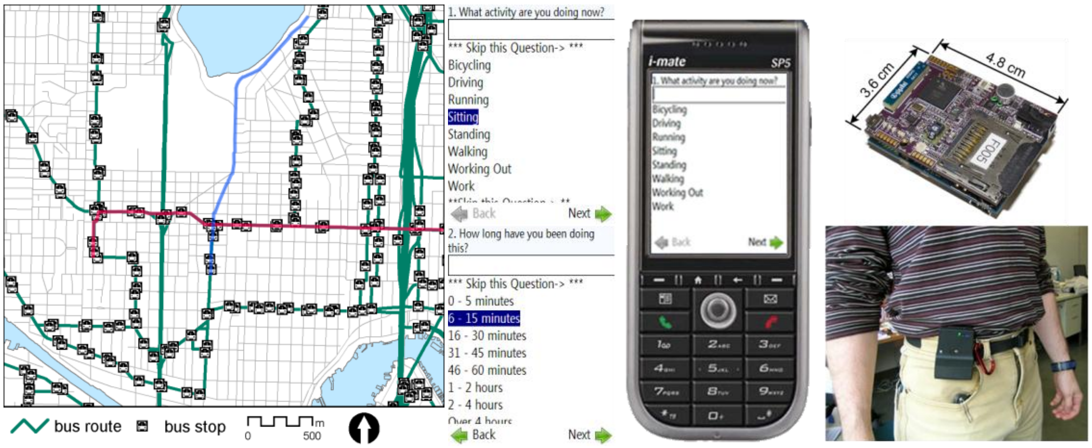

Fig. 2. (left) GIS Data showing street roads, bus routes, and bus stops; (middle-left) Questionnaires that prompted users to enter survey data about the activities they were doing; (middle-right) Picture of mobile phone users carried to enter surveys; (right) sensor platform (MSP) shown out of the box and inside the box

indicating the total number of hours of data collected from each subject in blue, along with the lengths of time when a GPS lock was available shown in purple. The percentage of time when a GPS lock was available is shown for each subject as well; on average a GPS lock was available 66% (std. dev. 28%) of the time. Also shown in Fig. 1 (right) is a trace showing the start and end time of a recording session as a horizontal line, for all the data collected, organized by day of the week. The data collected is spread fairly well across the days of the week as well as covering most of the hours of the day.

IV. Data Analysis

The data collected from our experiment serves as a test bed for the development of our analysis methods and models. Of particular importance to this research is the fusion of Geographical Information System (GIS) data sources with GPS readings. While typically one might assume that GIS data are difficult to obtain and use; in reality, these datasets generally exist for larger metropolitan areas and are frequently freely available from local governments. Because of the added value of geographic information, rather than focusing our efforts on extracting information from the raw GPS and sensor trace without any prior knowledge, we use the GIS data layers to provide the location of features such as roads, bus routes and bus stops, building outlines, and land use to bootstrap our inferences. While the analysis methods presented here used data from our MSP sensing device, we are also investigating how well these methods and the associated infrastructure transfer to mobile phone platforms for the evaluation of similar methods with a variety of data sources and usage scenarios.

A. Mode of Transportation

One inference we want to make is to predict transportation mode for personal trips. Given the sensor information and GPS traces, we would like to determine when someone is in a car, a bus, walking, or riding a bike. Walking and bicycling activities tend to involve periodic movements (pedaling and stepping) that we can identify using a step detector based on accelerometer readings. Pedaling on a bicycle is somewhat similar to walking, but often exhibits distinct accelerometer and velocity patterns. When GPS information is available, we can make use of GPS velocity estimates to help differentiate between walking and cycling. Bicycling tends to have interspersed periods of pedaling and coasting, so within a series of GPS+sensor records, if we detect periods of rapid motion with no stepping, the series is more likely cycling than walking.

For cars and buses we can similarly use some sensor information because car/bus rides look very different from walking, i.e. the lack of steps. However, because of similar velocity and accelerometer patterns, differentiating between bus and car trips can be difficult. In this case the GIS methods become very useful; we can overlay a GPS trace on a general street grid as well as a bus route layer, as shown in Fig. 2. The blue trace shows a route which diverges predefined bus routes, which results in a higher confidence that the trip was made by car. Likewise, the red trace shows a trip of sufficient duration following a bus route, which also includes stops at mapped bus stop

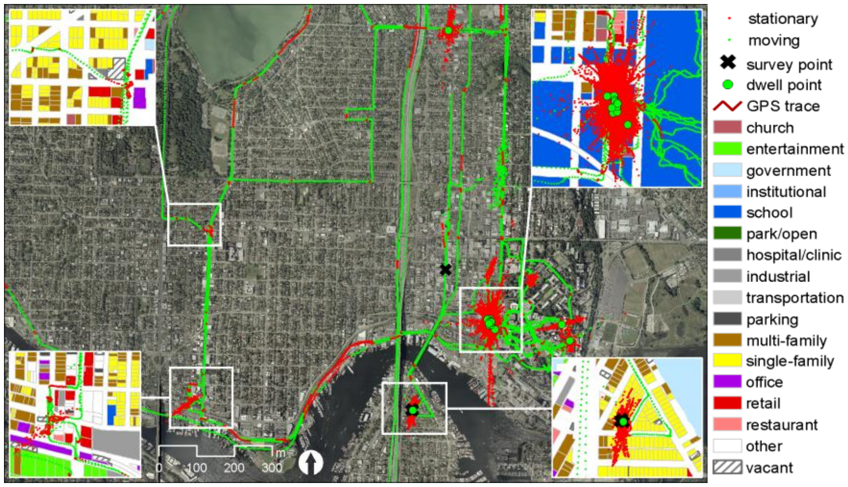

Fig. 3. GPS trace showing the path a person followed over several days. Stationary GPS points are shown in red, moving points are shown in green. Inserts show the GIS layer information for several points along the person's path

locations, so this trip was more likely to have been made by bus. However, because of the staccato motion of stop-and-go traffic, a car trip can also inadvertently appear to be a trip made by bus. Nevertheless, it may be possible to differentiate between car and bus trips based on the location of bus stops (versus random traffic stops).

B. Extracting Trips and Dwell Locations

A majority of people spend their time at several distinct locations throughout the day: home, work, retail and grocery stores, restaurants, etc. From the data we have collected, we are particularly interested in identifying locations where people spend a great deal of time, and associating these locations with information about the built environment obtained from GIS data sources. Fig. 3, shows a GPS trace from several days of one of our subjects. In this figure we've colored the trace with points where the person was stationary and moving, along with several locations where they dwelled for long periods of time. The combination of activity information, GPS location, and GIS data provides a very rich information base for the investigation of spatial patterns of personal behavior. For practitioners of urban design, transportation planning, and public health, knowledge of how people in a given area interact with the surrounding environment would be invaluable. Do people living or working in a particular area have more or less physical activity? Do they tend to go out for lunch? What restaurants do they visit, what grocery stores do they purchase from, and how frequently? Do they walk, bike, drive, or take public transportation to get to nearby restaurants and stores? Do these patterns of behavior vary with the socioeconomic status of individuals or neighborhoods?

C. Validation and Usefulness

In addition to the sensing and location information we collected from our subjects, we also asked them to answer hourly surveys asking questions about the type, duration, and location of activities they were performing. The purpose of these surveys was to provide some ground truth for evaluating the performance of activity inference. Ideally we would like to obtain fully labeled datasets; however, this is not feasible given the scale of data collection involved and, of course, cost limitations.

The surveys are used as testing data to verify that the automatically classified activity matches self-reported activity. However, unlike many statistical techniques, the ground truth here should be treated with some care; subjects may make mistakes in recording their activity or in estimating the true duration of continuous activity. For example, a subject walking home may have reported walking for 20 minutes but there may have been several minutes of window shopping, standing at crosswalks, etc. Due to these complications we have to be a bit more careful about how we incorporate the self-reported survey data for judging the accuracy of our methods. We are currently examining several methods and heuristics to determine which methods provide the best insight.

The question of usefulness for the GIS researchers is currently on an on-going research problem. As we complete more of the analysis of the data it will be possible to tease out more information from the data and determine if the researchers' hypotheses are borne out by the experimental data. While the data set is quite large, in order to make conclusions that would affect urban design or public health policy for long term effects, it will be necessary to isolate fewer variables and conduct more longitudinal studies.

V. Conclusion and Future Work

While there have been several sensing platforms that have reached a large audience of researchers [4], [3]; in general large scale deployments have been the exception rather than the rule. However, with the commoditization of sensors in mobile phones, the audience of available end users is much larger than would be possible for traditional research platforms to reach. While mobile phones do have some limitations, their wide adoption and ever growing capabilities make them ideal for the next generation of location-based sensing applications.

The current methods we are developing are based on data obtained from the MSP and continuous GPS (when available). However, we are also conducting several pilot studies using the Nokia N95 and iPhone. The accelerometer data retrieved from the N95 and iPhone are quite similar to the data obtained from the MSP. A few important differences are that while the MSP is be clipped on the belt and not usually moved, the N95/iPhone would often be manipulated to look up information or to make calls, which would introduce noise that would need to be taken into account in our inference models. Another important difference is that the location stack on the iPhone provides several levels of localization: large scale cell tower localization, medium scale WiFi localization, and high-precision GPS localization. Rather than require the iPhone to continuously gather precise GPS readings, we are instead trying to use more coarse grain location estimates in our algorithms to more intelligently decide when to switch to more accurate (and power hungry) localizations. If a person is sitting at work inside a large building it makes little sense to try to get a GPS lock, instead we might be satisfied with the cell tower localization and only decide to obtain WiFi or GPS localization if we believe the person is leaving the building.

The study described in this paper has provided a large amount of useful data which we will be able to use in the design of algorithms and applications for mobile phones. However, while we were able to collect a great deal of data in this study, in order to be able to answer questions of how the urban environment affects what people do, we need to collect data for longer time scales and from more individuals. With the increasing capabilities and adoption of mobile phones, developing methods that can be implemented on mobile phone platforms opens up the possibility of reaching large audiences. But, at the same time, custom devices that are readily available are still important and both can provide useful data that is applicable to both custom sensing platforms and commodity mobile phones.

References

[1] Consolvo, S., McDonald, D. W., Toscos, T., Chen, M. Y., Froehlich, J., Harrison, B., Klasnja, P., LaMarca, A., LeGrand, L., Libby, R., Smith, I., and Landay, J. A. 2008. Activity sensing in the wild: a field trial of ubifit garden. In Proc. of CHI 2008.

[2] Froehlich, J., Chen, M., Consolvo, S., Harrison, B., & Landay, J. 2007. MyExperience: A System for In Situ Tracing and Capturing of User Feedback on Mobile Phones. In Proc. of MobiSys 2007.

[3] MSP Research Challenge. http://seattle.intel-research.net/MSP

[4] TinyOS Community Forum | Hardware Designs. http://www.tinyos.net/scoop/special/hardware

[5] E. Agapie, G. Chen, D. Houston, E. Howard, J. Kim, M.Y. Mun, A. Mondschein, S. Reddy, R. Rosario, J. Ryder, A. Steiner, J. Burke, E. Estrin, M. Hansen, M. Rahimi. 2008. Seeing Our Signals: Combining location traces and web-based models for personal discovery. In Proc. of HotMobile 2008.

[6] Miluzzo, E., Lane, N., Eisenman, S., Campbell, A. 2007. CenceMe - Injecting Sensing Presence into Social Networking Applications. In Proc. of EuroSSC 2007.

[7] Frank, LD., Schmid, TL., Sallis, JF., Chapman, J., Saelens, BE. Linking objectively measured physical activity with objectively measured urban form: findings from the SMARTRAQ. American Journal Preventive Medicine 2005.

[8] Liao, L., Fox, D., and Kautz, H. 2005. Location-Based Activity Recognition using Relational Markov Networks. In Proc. of the Nineteenth International Joint Conference on Artificial Intelligence 2005.

[9] Mohan, P., Padmanabhan, V., Ramjee, R. 2008. TrafficSense – Rich Monitoring of Road and Traffic Conditions using Mobile Smartphones. Technical Report. MSR-TR-2008-59.

An Implementation of Mobile Sensing for Large-Scale Urban Monitoring

Teerayut Horanont[1], Ryosuke Shibasaki[1,2]

[1]Department of Civil Engineering, University of Tokyo, Meguro, Tokyo 153-8505, JAPAN
Email: teerayut@iis.u-tokyo.ac.jp
[2]Center for Spatial Information Science, University of Tokyo, Kashiwa, Chiba 277-8568, JAPAN
Email: shiba@csis.u-tokyo.ac.jp

Abstract

The extent of human flow in the urban space has reached unprecedented levels, and is therefore the focus of the present study. Unlike common practice on urban monitoring which utilizes cameras or sensors, this research aims to introduce a new platform of urban analysis and monitoring by using mobile sensing to recognize urban behaviour. Population density and its moving pattern can thus be visualized from the usage of the mobile phone.

The mobile sensing system is implemented on a web-based interface in order to maximize compatibility and interoperability. The aggregate mobile usages and antenna-mast positions are interpolated into grid-density surfaces. We then analyze urban patterns at a point of time to illustrate how people experience their city. Furthermore, the analytical results help to detect and to explain geographical 'hot spots' and clusterings of the unique characteristics of each urban space. Finally, we present visualizations of the results such as pseudo color and contour maps in order to demonstrate the urban dynamics.

I. Introduction

With the ubiquity and ever-increasing capabilities of mobile devices, cell phones and their locations could potentially become a powerful source to describe the pattern of the urban space. Traditionally, urban monitoring and analysis rely on a fixed location and a considerable amount of statistical data. This procedure does not permit the identification of multi-temporal events in wide areas. Another key problem is how to acquire and update the statistical data of the moving objects, for instance, humans and their activities in the whole city space.

In this research, the usage of mobile devices will be treated as a medium for data collection. Erlang data that represent a distribution of call duration in the Global System for Mobile Communication (GSM) network could be performed as aggregate-data sources to estimate the population density of a city. For the large scale monitoring, clusters of Erlang data from mobile base stations are excellent at providing indirect interpretations of spatial patterns of urban life and its temporal dynamic. This aspect is very useful in the view of public monitoring. It could potentially become a new way to extract or identify invisible problem spots from the complex urbanized areas. Furthermore, mobile sensing is potentially applicable for public-marketing analysis. The distribution of a population at different points of time in each city space could be an ideal source to help people decide a place for urban advertising or opening a shop. In addition, an exploration of mobile sensing data would give the urban planner a better understanding on flowing patterns of people at specific times of the day. If we look at broader contexts, then population transfers from city to city during special events or public holidays could also be determined.

Study area

The research takes place at Bangkok Metropolitan, Thailand. Bangkok city has a well developed mobile network and has a high degree of mobile usage. We received the support of the Erlang data and base station locations from Advanced Info Service PLC (AIS), a leading mobile operator in Thailand.

II. System Architecture

In the early research, the system was implemented as a fundamental tool for spatial exploration and visualization which permitted the data to be obtained, integrated and displayed quickly, easily and flexibly. Since the internet has now made distance virtually disappear, we questioned the traditional method of urban monitoring by enhancing the way how to instantaneously integrate and obtain large amounts of data via the network.

Back end application servers

To maximize compatibility and interoperability, open standards such as XML and web services are utilized for data exchange and sharing. Data analysis is assisted by the open source statistical package R (http://www.r-project.org), integrated into the PostgreSQL database via the PL/R procedural language.

A browser front end

The web browser we use as a universal front end, an Ajax mashup, is a hybrid web application which presents a rich UI to update and integrate contents asynchronously from multiple sources. This makes combining data easier, not only spatial data from the host server but also third-party sources from the services available on the internet. The calls can also be made directly to the third-party sources from the browser or back to the originating server, which acts as a proxy for the third-party contents.

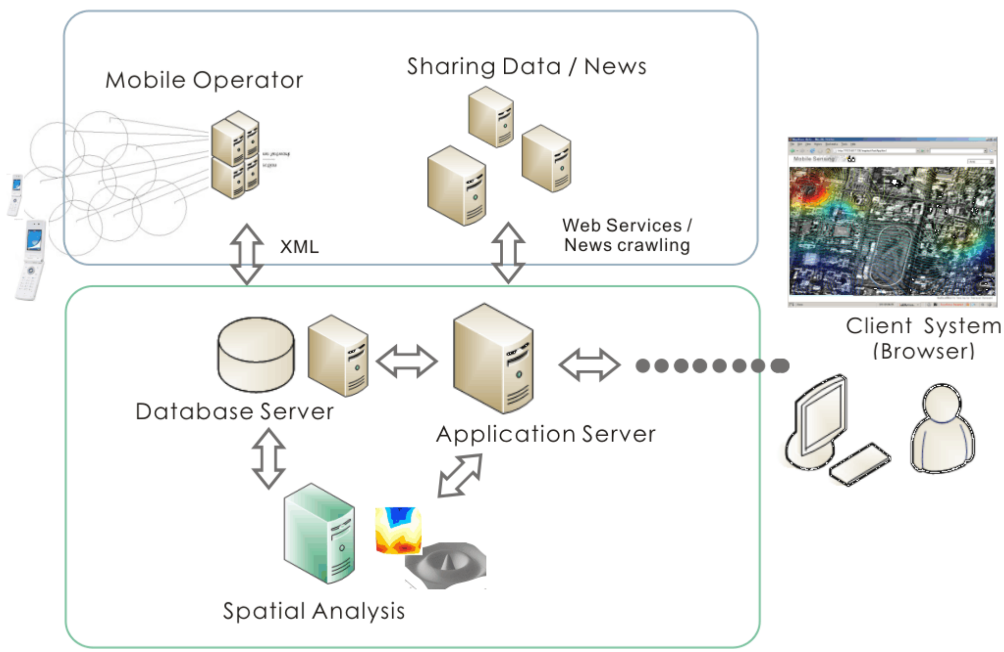

Fig. 1. System architecture of mobile sensing system

Data manipulation

We first simulate a connection with a mobile operator in our local environment instead of communicating directly to the mobile-operator system. The mobile-log data during the period from February to April 2008 that cover a part of the central Bangkok area are transformed and inputted into a database. The data established in the database mainly include cell-id, the base station geographic position, update time and Erlang data. The Erlang data which is calculated from call duration are performed as a sample distribution in order to estimate the population density of the whole area.

cellid	lat	lon	Start_time	erl
BKKC1	13.75697	100.5594	2008/03/01 9:00	33.98
BKKC2	13.75697	100.5594	2008/03/01 9:00	18.93
BKKC3	13.75697	100.5594	2008/03/01 9:00	33.17
PTWA1	13.75138	100.5402	2008/03/01 9:00	20.75
PTWA2	13.75138	100.5402	2008/03/01 9:00	17.93
PTWA3	13.75138	100.5402	2008/03/01 9:00	33.07

Table 1. Sample data from the Base Station Controller (BSC)

Population and prediction model

To present population data in a continuous space, we need interpolation techniques to generate a surface from discrete points. There are many interpolation techniques, each with its own weaknesses and strengths. In this paper, we introduce an inverse distance weighted (IDW) and ordinary kriging method to predict population density in our prototype system. In order to increase the resolution of the interpolation results, we calculate the weight in each sector cells separately by using voronoi tessellation to increase a discrete set of points.

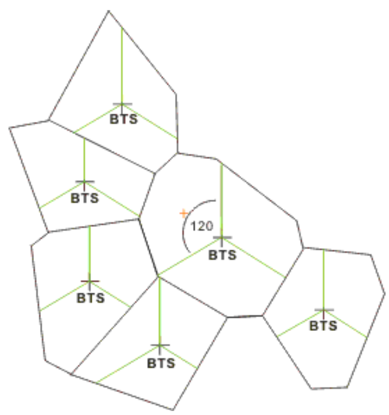

Fig. 2. Voronoi tessellation over sector base station

III. Results And Discussions

We start by presenting our first results on querying the cumulative usage of mobiles over an hour interval. Histogram and time series statistics can be retrieved from specific locations on a map and display in a time-plot based graph. The first two types of graphs, day and month, were produced by implementing a getStat webAPI in conjunction with Timeplot, a DHTML-based AJAX widget. The exploration leads to a better understanding of each area's activities during one day, as well as the difference characteristics between weekdays and weekends. In Figure 3(a), we pick up an office area in central Bangkok, and the time plot shows the trends of increasing activities from early morning up to the peak at noon, then decreasing gradually after 5.00 pm. Figure 3(b) and (c) illustrate the overall activity on a monthly basis. We can capture a weekly rhythm of this area which clearly defines high activities on the weekdays and appears to decrease on Saturday and Sunday. Figure 3(c) presents a pulse in April: we can see a week of flat, low activities since it was a long holiday in Thailand.

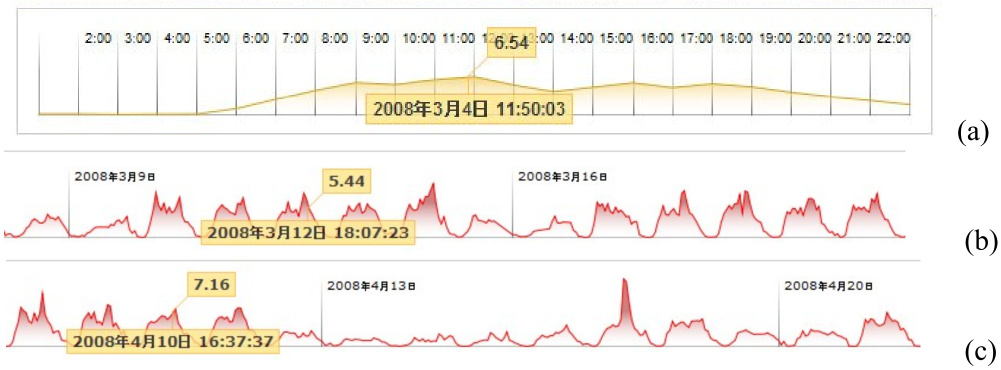

Fig. 3. Day and month statistics from cumulative mobile usages data

Temporal analysis of human flow

Another approach to explore the data is to generate a surface-flow pattern by interpolating the aggregate call traffic. Exploratory analysis of temporal data can give a clear view on how people flow into and out of the city throughout the day. Figure 4 shows the flow patterns in one local area from 6.00 am until 8.00 pm. This observation leads to the speculation on how one part of the central city is upscale, crowded and how long the area keeps busy until people move to another part of the city. It is extremely useful for the urban planner to figure out how type of land use, street network and other city landscapes could affect the flow and density of the urban area. Furthermore, results of the study not only provide a tool for area or zoning analysis but also could be used to specify hidden problems of the particular space over a period of time.

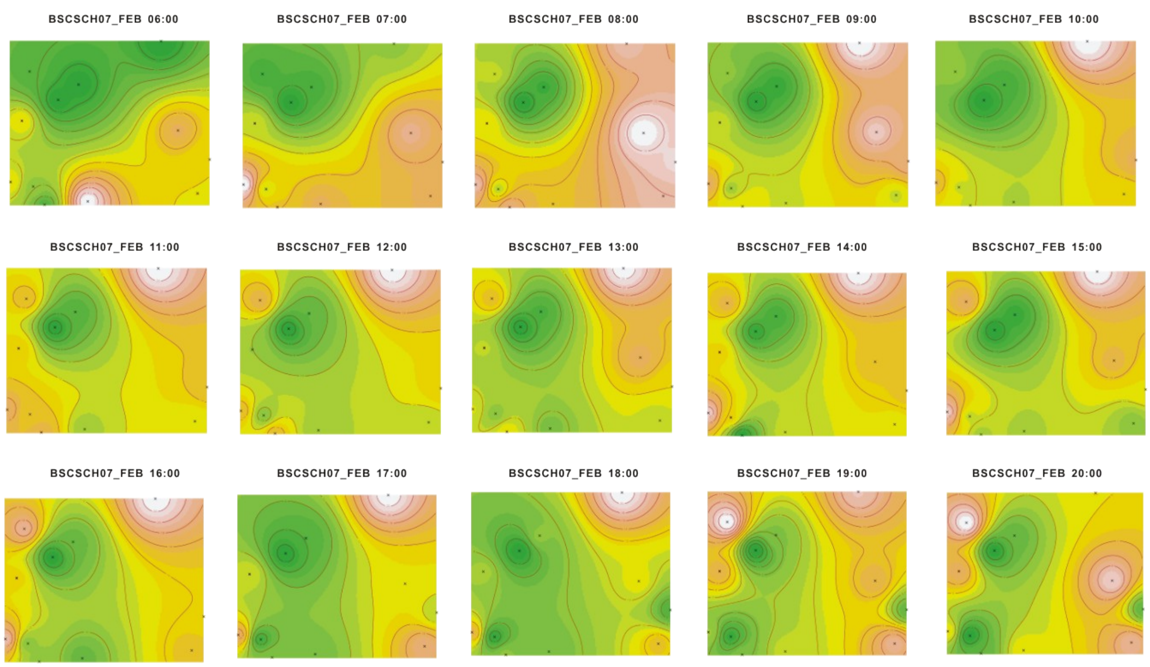

Fig. 4. The flow pattern from early morning to late evening in central Bangkok

Hot spot capturing

In order to clearly highlight the extreme values in distribution surfaces, we now generate the volume of population density with an overlay contour diagram using a specific color palate. In Figure 5, to illustrate how mobile density could reflect the real world daily activities, we capture Pathumwan area. This is one of the most active spots in Bangkok that has a mixed land use, for example high rise office buildings and a large-scale shopping complex. If we compare the same period of time at 1.00 pm. on Friday and Sunday, the result demonstrates that on Friday the activities are dense at the office area and, in contrast, a peak density moves towards the shopping area on Sunday. Obviously enough, this kind of hot spot extraction could be useful to capture some hidden aspects in the urban space. We are planning to collect a longer span of archive data and make a base urban signature in order to implement a real-time signature recognition and hot spot extraction.

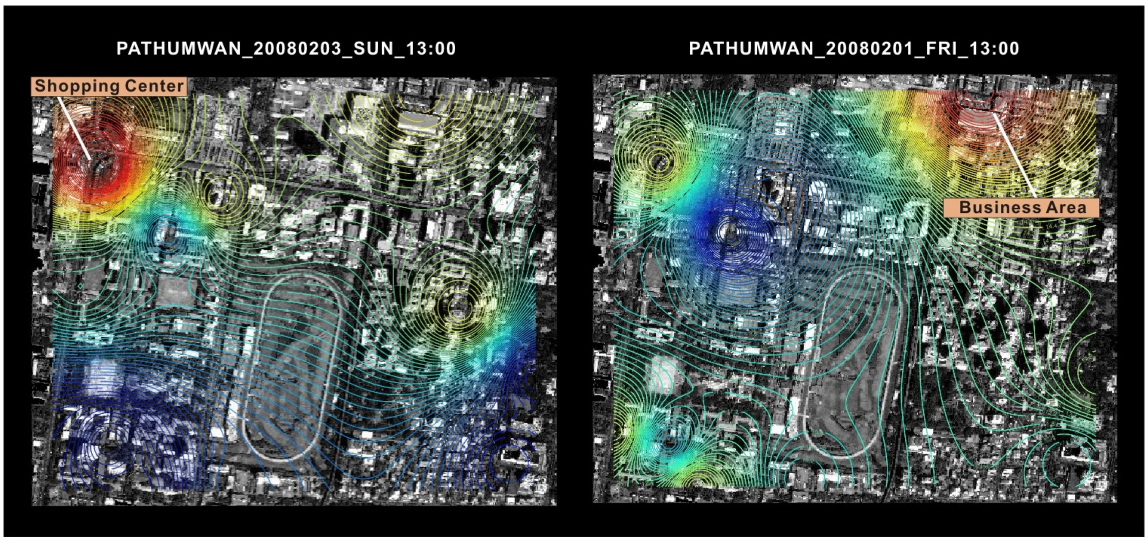

Fig. 5. Density contour in Pathumwan area on Friday and Sunday at 1.00 pm.

REFERENCES

[1] Boucher, N J. The cellular radio handbook: a reference for cellular system operation. 4th Ed.

[2] Hengl T. (2007). A Practical Guide to Geostatistical Mapping of Environmental Variables.

[3] Lance A. Waller and Carol A. (2004) Visualizing Spatial Data. Applied Spatial Statistics for Public Health Data. pp. 68-117.

[4] Mannings, R. (2006). The Invisible GIS: Technology Convergence to make future GI user friendly.

[5] Pashtan A. (2005). Mobile Web Services. Aware Networks. Illinois.

[6] Ratti, C. (2007). Go with the flow. The Economist Technology Quarterly, pp. 12-13.

[7] Ratti C., Pulselli R M., Williams S., Frenchman D. (2005) Mobile Landscapes: using location data from cell-phones for urban analysis. SENSEable City Laboratory, MIT.

[8] Working Group. (2005). Urban Mobility Initiative. European Conference of Transport Research Institutes: REPORT ECTRI number 2005-06.