Mobile Context Inference Using Low-Cost Sensors

Evan Welbourne[1], Jonathan Lester[2], Anthony LaMarca[3], and Gaetano Borriello[1,3]

[1] Department of Computer Science &
Engineering
[2] Department of Electrical Engineering
University of Washington
Box 352350, Seattle, WA 98195 USA
{evanwel, jlester}@u.washington.edu

[3] Intel Research Seattle
1100 NE 45th Street, Suite 600
Seattle, WA 98105 USA
{anthony.lamarca,
gaetano.borriello}@intel.com

Abstract. In this paper, we introduce a compact system for fusing location data with data from simple, low-cost, non-location sensors to infer a user's place and situational context. Specifically, the system senses location with a GSM cell phone and a WiFi-enabled mobile device (each running Place Lab), and collects additional sensor data using a 2" x 1" sensor board that contains a set of common sensors (e.g. accelerometers, barometric pressure sensors) and is attached to the mobile device. Our chief contribution is a multi-sensor system design that provides indoor-outdoor location information, and which models the capabilities and form factor of future cell phones. With two basic examples, we demonstrate that even using fairly primitive sensor processing and fusion algorithms we can leverage the synergy between our location and non-location sensors to unlock new possibilities for mobile context inference. We conclude by discussing directions for future work.

1 Introduction

The near ubiquitous deployment of cell phones and other ever more advanced mobile platforms in our daily lives has spurred an increased interest in mobile context inference. In the past few years, the decreasing cost of on-device location technology has led to a number of systems that use location traces to learn a person's significant places and his daily patterns of residence and transit among those places [1, 2, 3, 7, 17]. In other recent work, researchers have developed mobile systems that use simple, low-cost, non-location sensors to infer a user's (or a device's) situational context (e.g. activity, whether or not the device is being carried) [4, 5, 6, 19]. Some of these systems have used their sensors to infer high-level context about a user's place and situation, but only in simple ways and with limited flexibility. For example, Patterson et. al. used GPS logs and instantaneous velocity to recognize transit activities such as driving a car or riding a bus, but could only do so with a fairly regular GPS signal, and by using a large amount of a priori information (e.g. street maps and bus schedules).

Our goal is to truly fuse location and non-location sensors and to leverage the synergy between them to enable a wider variety of high-level mobile context inference. In this paper, we introduce a system that is a significant next step in this direction. The system uses Place Lab [18] to provide indoor-outdoor location information with an average resolution of 20 - 30 meters in covered areas. To collect additional sensor data we use the 2” x 1” multimodal sensor board (MSB), which contains a variety of common non-location sensors (e.g. accelerometers, barometric pressure sensors). We illustrate the fundamental advantages of our system by combining location context with simple movement detection in two representative examples. In particular, we show that like previous systems, our system can:

- Classify a mode of transit (as driving, walking, or running) when the user is moving;
- Extract significant places within a user’s daily movements.

However, unlike previous systems, our system can also:

- Classify mode of transit without GPS, precise velocity information, or learned knowledge of transit routines;
- Classify or highlight significant places based on the activity that occurs there.

Our principal contribution is to show how the problem of mobile context inference can be tackled using technologies that are or will be readily available in common computing platforms that are becoming truly ubiquitous, such as cell phones and wrist-watches. By using Place Lab as our location system we alleviate the need for costly infrastructure as well as for uncommon hardware, instead relying on existing infrastructure (cell phone towers and wireless access points) and commodity hardware (WiFi cards/chips and GSM receivers). We use sensors that vendors could easily incorporate into their mobile platforms; in fact, several of them have already done so. For example, Pantech’s PH-S6500 phone is equipped with six motion sensors that are used to detect how fast and how far its owner has walked, while new IBM ThinkPads are equipped with accelerometers that measure forces so as to park the hard drive in the event that the laptop is dropped. With our centralized system design, we also collect and fuse location and non-location sensor data without requiring sensors or computation to be distributed around the body.

The remainder of this paper is organized as follows. In the next section we review recent work in mobile context inference. In section 3, we describe our system and its components in detail. Then, in section 4, we present two examples that demonstrate how our system can fuse location and non-location sensor data to expand the range of possibilities for mobile context inference. Before concluding in section 6, we discuss directions for future work.

2 Related Work

There has been a substantial amount of work in mobile context inference since the early 1990’s. In the past five years, research has focused increasingly on systems that could be deployed ubiquitously in the near future by leveraging existing mobile infrastructure (e.g. commodity GPS devices, cell phones). Specifically, recent projects have tended toward one of two areas: location-based user modeling or situational context awareness using simple, non-location sensors.

A number of location-based projects have used long-term GPS or GSM traces of a user's daily movements (on the order of weeks) to generate a probabilistic model of that user's residence and travel patterns [1, 3, 4, 17]. The resulting systems can help determine where a user is and where she is going, but fall short of being able to determine other kinds of equally important context (e.g. what she is doing at that place). Moreover, these schemes are limited by the drawbacks of the location technology used: GPS-based systems have problems in urban canyons and indoors, while GSM cellular location data is often too coarse to allow distinction of neighboring places. Kang et. al. [10] circumvented these problems by using a WiFi-based location system (Place Lab) and were able to perform indoor place extraction with a location resolution of 20 - 30 meters [18]. We show how to augment Kang's method with basic movement detection using a common 3-axis accelerometer that costs about $10 US.

Most situational context-awareness projects have focused on human activity inference (e.g. walking, speaking, riding a bicycle) using multiple simple, body worn sensors [5, 6, 19]. This approach allows determination of what a user is doing at a particular point in the day, but falls short of being able to frame that activity in the context of his larger patterns of daily movement. Schmidt et. al. built a system with a multi-sensor board and GSM cell phone that is quite similar to ours, but focused much less on location.

A few systems [1, 4, 8] have combined place and situational context, but have been limited both in approach and in the technology used. Both Patterson and Liao [1, 7, 8] use incoming GPS data to infer mode of transit (e.g. foot, car, bus), but rely on a large corpus of learned travel patterns, and on prior knowledge of street maps and bus schedules. Marmasse and Schmandt process GPS and accelerometer data in real-time to distinguish transit activities such as walking and driving between places, but require that GPS signal be continually available and that sensors be distributed across the user's body rather than centralized on one device.

We overcome these limitations by using a single location system that works both indoors and outdoors, and by fusing location information with information from a collection of simple, non-location sensors. The use of these additional sensors can simplify the handling of location information by reducing or eliminating the need for a priori knowledge and extensive training. We also group all additional sensors on a centralized, worn device.

3 System Components and Sensors

Our system consists of three components: a GSM cell phone, a waist-worn, WiFi-enabled mobile device, and a multi-modal sensor board (MSB). The cell phone provides GSM data for location and could also be used to take photos and video during experience sampling method (ESM) studies. The mobile device provides WiFi data for location and is also the central processing and storage unit. The MSB captures a variety of sensor data and is discussed in more detail below. This design was developed with the foresight that the phone and mobile device functions would eventually be merged into a GSM PDA phone with WiFi.

The cell phone and the mobile device both run Place Lab software and communicate via Bluetooth; fusion and logging of location data is performed on the mobile device. The MSB is attached via USB or Compact Flash to the mobile device, which is responsible for both logging and processing its sensor data in real-time.

Our first implementation used a Nokia 6600 cell phone and a laptop with the MSB. After conducting our initial experiments (see section 4), we also assembled a more compact version of the system that uses a Nokia 6600, the MSB, and an iPaq hx4705. We plan to use the iPaq-based system in user studies because the iPaq hx4705 is smaller and less intrusive than most other mobile devices.

3.1 Multi-Modal Sensor Board

The MSB (Figure 1) is a 2” x 1” sensor platform we built to allow us to experiment with different sensors. The MSB contains seven sensors (listed along with sampling rates in Table 1) and is capable of running independently, collecting data from all seven sensors, and then relaying that information to a handheld, laptop, or to Intel’s iMote where it can be processed or stored. The iMote is an ARM7 Bluetooth sensor node which can be attached to the MSB [11, 12]. The MSB has an on-board ATmel ATMega128L microprocessor running at 7.3728 MHz which is capable of performing only simple data analysis before relaying the data to the host device.

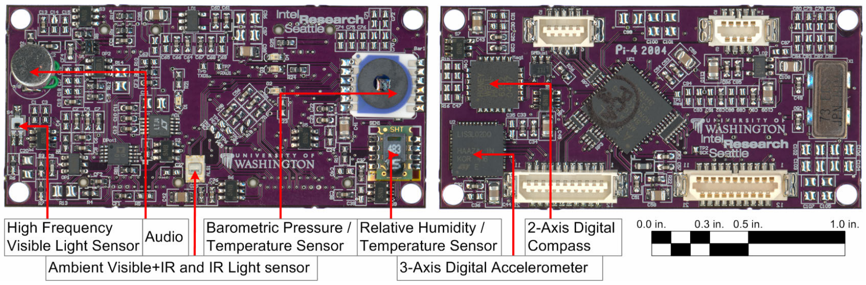

Fig. 1. Photograph of the top and bottom of the MSB

Table 1. The sensors available on the MSB and their sampling rates

Manufacturer	Part No.	Description	Sampling Rate
Panasonic	WM-61A	Analog Electric Microphone	15,360 Hz
Osram	SFH-3410	Analog Visible Light Phototransistor	549 Hz
STMicro	LIS3L02DS	3-Axis Digital Accelerometer	548 Hz
Honeywell	HMC6352	2-Axis Digital Compass	30 Hz
Intersema	MS5534AP	Digital Barometer / Temperature	14 Hz
TAOS	TSL2550	Digital Ambient (IR and Visible+IR) Light	5 Hz
Sensirion	SHT15	Digtial Humidity / Temperature	2 Hz

In our current setup we connect the MSB directly to a laptop or iPaq though either a USB or Compact Flash connection. The sensors available on the MSB were chosen to simplify future development; a majority of the sensors are simple digital components that could be easily integrated into devices such as cell phones. This allows us to experiment with a rich sensor set and to easily determine which sensors are appropriate for our goals.

4 Mobile Context Inference Examples

To demonstrate how our system can fuse location and non-location sensor data to infer high-level context and to enable new types of mobile context inference, we apply it to two common problems in location-based user modeling. Specifically, we show that like related systems [1, 2, 3, 4, 7, 8, 10, 17] our system can perform basic mode of transit inference and extract significant places from a user's daily movements. Further, we show that we can infer walking vs. driving without GPS, precise velocity information, or prior knowledge of transportation routines; we also present an example of sub-place identification and classification using a primitive form of activity recognition (movement detection). Both examples are presented and discussed below in more detail.

4.1 Mode of Transit

Using time, location, accelerometer data, and a very simple classification scheme, we can distinguish three modes of transit: walking, running, and riding in a vehicle. With an average location history of 25 to 120 seconds (shorter histories for areas with dense AP coverage, longer histories for sparser areas) we can classify a user as not in transit if she is settled within a geometric region or in transit if she is not settled within any geometric region. We can then use the location history to compute a rough estimate of the user's speed, which is classified as "vehicle speed" if she is in transit and consistently faster than 9 m/s (approximately 20 miles/hour). We can further distinguish the mode of transportation using the accelerometer data.

Previous work in bio-mechanics has shown that the frequencies of interest for walking motions are below 10Hz [13]. Using the accelerometer data from the MSB, we can perform a simple sum of the FFT coefficients from 0.5Hz to 3Hz (which contains the first harmonic of the walking and running motions) to get a rough idea of the movements being experienced. Since we only want a rough result and intend to perform this calculation on devices with limited power and computation, we use the Goertzel algorithm [14] which allows us to compute DFT coefficients very efficiently with a 2^{nd} order filter. We store two seconds worth of acceleration magnitudes in a buffer and compute a binary movement/no-movement estimate three times a second with a simple threshold of the sum (a sum of greater than 10 indicates movement and a sum of less than 10 indicates no movement).

Figure 3, shows some sample data where we have the recorded acceleration, the binary movement/no-movement estimate, and the estimated speed for a user in an area densely populated with 802.11 access points (Seattle's university district). At the beginning of the recording the user was walking, switched to running around the 150 second mark, and then got into and drove a car down a busy street from the 225 second mark until the 450 second mark. While driving, the user stopped at traffic lights near the 260 second and 325 second marks, pulled into a parking lot around the 450 second mark, and finally parked near the 460 second mark.

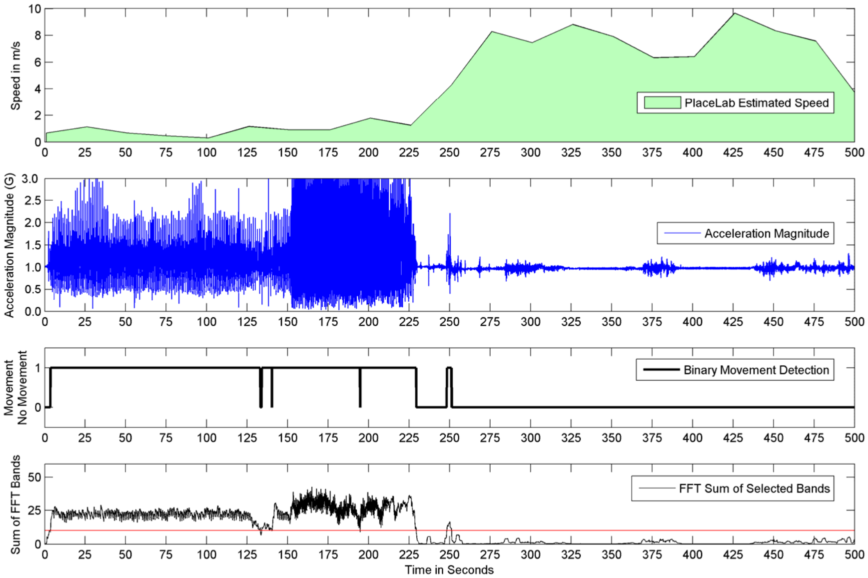

Fig. 3. The sample sensor trace with the user walking for ~150 seconds, then running until the ~225 second mark, then entering a car and driving for the remainder of the recording (with stops at the ~260 and ~325 second marks, and parking in a lot starting at the ~450 second mark). The summed frequency components of the acceleration are used along with acceleration magnitude and the "in transit" classification from the location history to determine whether the user is walking, running, or driving/riding in a vehicle. The top graph shows the PlaceLab estimated speed, calculated every 25 seconds. The 2nd graph shows the acceleration magnitude recorded from our sensor board, in G's. The 3rd graph shows the output of our binary movement (1) / no movement (0) threshold. The 4th graph shows the raw Goertzel sum calculated from the acceleration, the threshold at 10 is shown as a solid line across this graph. Points above the threshold are classified as 1, moving, and points below the line are classified as not moving, 0.

Throughout this recording, the user was classified as in transit, and his estimated speed was calculated every 25 seconds. We can see that for almost the entire trace, the user's estimated speed was below 9 m/s so that estimated speed alone couldn't provide a reliable mode of transit classification. Thus, given that the user is in transit, we look at the binary movement/no-movement estimate and acceleration magnitude to distinguish mode of transit. We can see that for the walking and running phases, the binary movement/no-movement estimate (the black line in the 3rd graph in Figure 3,valued at 0 or 1) was almost always 1, while for the driving phase the movement/no-movement estimate was almost always 0. From this we an infer that the user was walking or running for the first 225 seconds, and in a vehicle for the remainder of the trace (because he was in transit but classified as exhibiting no movement). By looking at the acceleration magnitude for the first 225 seconds, it is also clear that the user was walking for the first 150 seconds, and running for the next 75.

While this overly simple method for inferring mode of transit inference may not be as accurate as other machine learning techniques (and could certainly be improved upon – by adding audio data for example), it shows the promise of our multi-sensor approach. In the future, we expect to be able to distinguish among different motor vehicles, such as cars, busses, and motorcycles, by their characteristic acceleration patterns and acoustic signatures; and by leveraging more advancing machine learning techniques [20].

4.2 Significant Places

We can also combine place extraction with activity recognition to perform more sophisticated inference about significant places. For example, we could "label" a place with the set of activities that occurs there or extract only those places where a particular activity occurs. Our simple example is an enhancement to Kang's place extraction method and uses information about the user's movements within a given place. Specifically, if a user has spent 5 or more minutes within a place and our movement-detection algorithm reports that he has not moved, then we decide that he is at a desk or some other significant, fixed station. In this case, we might label that place as one where a stationary activity occurs. Alternatively, we can "zoom-in" on that location by taking a WiFi fingerprint [15]; this provides us with a way to distinguish high-resolution (2 - 5 meter) sub-places that are currently finer-grained than Place Lab's average resolution (20 - 30 meter) allows.

To test our "zoom-in" enhancement to Kang's method, we recorded a day's worth of place and movement data for one student on a university campus. A visualization of the collected data (Figure 4) illustrates how our algorithm can extract multiple indoor places and sub-places. This basic example demonstrates how using an indoor-outdoor location technology in combination with a collection of simple, non-location sensors allows us to explore the relationship between activity and place.

5 Future Directions

The preceding examples have demonstrated how our system can fuse location and activity to solve common problems addressed by previous mobile context inference systems. We have also shown how our system could unlock new possibilities for mobile context inference. In this section we discuss possible directions for future work with our system.

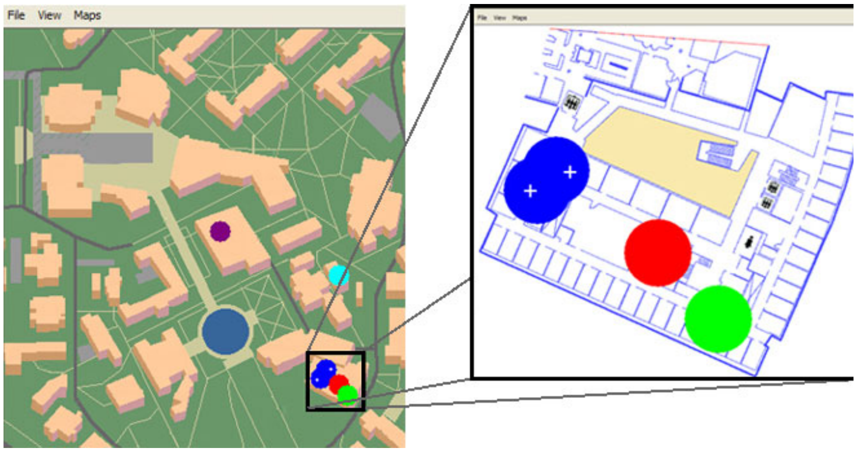

Figure 4: The places and sub-places extracted from a day's worth of data collection using our movement-detector-enhanced place extraction algorithm. The zoomed view shows three places within the same building (left to right): a student area, a hardware lab, and a faculty office. In the student area, the student spent time sitting at a desk and on a couch - both of these positions were recorded as sub-places (denoted '+' and depicted at an arbitrary position inside the place) and stamped with a WiFi fingerprint. No sub-places were identified in the hardware lab because the student was walking back and forth between tables, and no sub-places were recorded in the faculty office because the student spent only a short time there.

5.1 Infrastructure Assist

One key benefit of fusing location and non-location sensors is that we can use the result to assist the underlying location infrastructure. For example, mode of transit could be used to select an appropriate motion model for a location particle filter which would in turn improve Place Lab's accuracy. Collecting fingerprints of sub-places that are identified with activity (e.g. sitting) is another way to distinguish places at a higher resolution than what Place Lab currently offers. It is likely that other improvements can be made to Place Lab using data from more sensors. For example, digital compass data could be used to further augment a motion model by adding directional information.

5.2 Activity and Place

The meanings of place and activity are often closely tied. For example, it is often the case that a place is defined by the activities that occur there (e.g. a conference room is a place where meetings are held). Our system provides us with a rare opportunity to work with place and indoor activity – usually where most activities occur, and where the most time is spent. Some ideas in this area that may be worthy of investigation include the labeling of places with the activities that occur there, and work with the idea of a “mobile place” (e.g. a bus or train).

5.3 User Studies

We should be able to use our compact iPaq-based system in a number of user studies that rely on place and situational context. For example, we plan to use the iPaq-based system for a long term study in which data is collected from users during daily life, and annotated using the experience sampling method. We could also use this system to study the effects of place and activity on interruptability and prompting.

6 Conclusions

We have presented and described a new system for mobile context inference that fuses location and non-location sensors, and is designed for deployment on future computing platforms that will be truly ubiquitous (e.g. cell phones, PDAs). Our system has the primary advantage that it uses Place Lab for both indoor and outdoor location using existing infrastructure; and by using low-cost, simple sensors (the 3-axis accelerometer on the MSB costs about $10 US) that are already appearing on many emerging mobile platforms, we can widen the possibilities for mobile context inference and strengthen the robustness of our inference algorithms. We presented two examples of how our system can solve common problems (inferring mode of transit and extracting significant places), and unlock new possibilities unavailable to previous systems. Finally, we have discussed a number of promising areas for future work.

References

1. Patterson, D., Liao L., Fox, D., and Kautz, H., “Inferring High-Level Behavior from Low-Level Sensors”, The Fifth International Conference on Ubiquitous Computing, 2003.
2. Ashbrook, D., and Starner, T., “Using GPS to Learn Significant Locations and Predict Movement Across Multiple Users”, Personal and Ubiquitous Computing, v.7 n.5, pp. 275 – 286, 2003.
3. Marmasse, N., and Schmandt, C., “A User-Centered Location Model”, Personal and Ubiquitous Computing, pp. 318 – 321, 2002.
4. Marmasse, N., Schmandt, C., and Spectre, D., “WatchMe: Communication and Awareness Between Members of a Closely-Knit Group”, The Sixth International Conference on Ubiquitous Computing, , pp. 214 – 231, 2004.

5. Schmidt, A., Aidoo, K., Takaluoma, A., Tuomela, U., Laerhoven, K., and Velde, W. “Advanced Interaction in Context”, Proceedings of the 1st international symposium on Handheld and Ubiquitous Computing, pp. 89 – 101, 1999.
6. Bao, L., and Intille, S., “Activity Recognition From User-Annotated Acceleration Data”, Proceedings of the Second International Conference on Pervasive Computing, 2004.
7. Patterson., D., Liao, L., Gajos, K., Collier, M., Livic, N., Olson, K., Wang, S., Fox, D., and Kautz, H., “Opportunity Knocks: a System to Provide Cognitive Assistance with Transportation Services”, The Sixth International Conference on Ubiquitous Computing, 2004.
8. Liao, L., Fox, D., and Kautz, H., “Learning and Inferring Transportation Routines”, Proc. of the National Conference on Artificial Intelligence, 2004.
9. Stäger, M., Lukowicz, P., Perera, N., Büren, T., Tröster, G., and Starner, T., “SoundButton: Design of a Low Power Wearable Audio Classification System”, Seventh IEEE International Symposium on Wearable Computers, pp. 12 – 17, 2003.
10. Kang, J., Welbourne, W., Stewart, B., and Borriello, G., “Extracting places from traces of locations”, Proceedings of the 2nd ACM international workshop on Wireless mobile applications and services on WLAN hotspots, pp. 110 – 118, 2004.
11. Hill, J., et al., "The platforms enabling wireless sensor networks", Communications of the ACM, pp 41-46 , vol 47, no. 6, 2004.
12. Culler, D., and Mulder, H., "Smart Sensors to Network the World", Scientific American, pp 84-91, 2004.
13. Winter, D.: Biomechanics and Motor Control of Human Movement (2nd ed.). New York: Wiley (1990).
14. G. Goertzel, “An Algorithm for the Evaluation of Finite Trigonometric Series”, Amer. Math. Month., vol. 65, pp. 34-35, Jan., 1958.
15. P. Bahl and V. N. Padmanabhan, "RADAR: An RF-Based In-Building User Location and Tracking System", Proc. IEEE Infocom, March 2000.
16. LaMarca A., Chawathe Y., Consolvo S., Hightower J, Smith I., Scott J., Sohn T., Howard J., Hughes J., Potter F., Tabert J., Powledge P., Borriello G., Schilit B. “Place Lab: Device Positioning Using Radio Beacons in the Wild”, Intel Research Technical Report: IRS-TR-04-016.
17. Laasonen, K., Raento, M., and Toivonen, H., “Adaptive On-Device Location Recognition”, In Proceedings of the 2nd International Conference on Pervasive Computting, April 2004.
18. LaMarca, A., et al., “Place Lab: Device Positioning Using Radio Beacons in the Wild.” International Conference on Pervasive Computing (Pervasive 2005), Munich, Germany, May 2005 (to appear)

19. Siewiorek, D., et. al., “SenSay: A Context-Aware Mobile Phone”, IEEE International Symposium on Wearable Computers (ISWC 2003), New York, New York 2003.
20. Brunette, W., et. Al., "Some Sensor Network Elements for Ubiquitous Computing", The Fourth International Conference on Information Processing in Sensor Networks (IPSN 2005), Los Angeles, CA, 2005 (to appear).