

Smart Sensing Systems for the Daily Drive

Sarfraz Nawaz*, Christos Efstratiou[§], and Cecilia Mascolo*

* *University of Cambridge*

[§] *University of Kent*

In recent years, there has been explosive growth in smartphone penetration. Today, more than 90 percent of the world's population is covered by cellular networks, and nearly one quarter of the population uses smartphones.¹ This growth is set to continue as inexpensive devices become available and mobile phone users switch over to smartphones. This makes smartphones an ideal vehicle for large-scale sensing systems, because while a PC sitting on a desk doesn't know much about its user, a phone accompanying its user can learn about the user through a range of embedded sensors. The persistent connectivity offered by smartphones also lets users access information from almost anywhere.

Combining these two capabilities together can lead to interesting applications, especially when the user is away from home or work, without access to traditional devices. For example, consider a user traveling by car. He or she might wonder, "Are there any disruptions on my regular route that might cause delays? Will I be able to find parking when I arrive?" Furthermore, he or she might want to automatically send a message to someone saying, "I'm running late due to traffic." Smartphones can help, enhancing the driving experience.

In particular, we developed two smartphone-based prototype systems that address this scenario. Our first system, ParkSense, can sense on-street parking-space occupancy when coupled with electronic parking payment systems. Our second system can sense and recognize a user's repeated car journeys to provide personalized alerts to the user. Both systems aim to minimize the impact of sensing tasks on the device's lifetime so the user can continue using his or her smartphone for its primary purpose.

Sensing On-Street Parking

In metropolitan city centers and dense urban areas around the world, parking spaces are often a scarce resource. During busy periods of the day, drivers can spend a significant amount of time searching for available spots. Although off-street parking garages can usually estimate the amount of available spaces by counting incoming and outgoing vehicles, the lack of information on the availability of on-street spaces usually leads drivers to cruise around looking for vacant spots. This results in congestion, wasted fuel, and pollution in dense urban areas. A study of major downtown areas has shown that 30 percent of the traffic in these congested areas is due to cruising vehicles.²

This situation can be alleviated by providing real-time occupancy information to cruising drivers. One approach to capture space occupancy is to embed an occupancy sensor in each parking space, as was done for the SFpark pilot program in San Francisco (<http://sfpark.org>). However, significant costs are involved in instrumenting large cities.

Instead, we decided to exploit smartphone applications that let users pay for parking, such as PaybyPhone (<http://paybyphone.co.uk>). Such applications have recently seen wider adoption. For example, over 9,000 on-street parking spaces managed by the local city council in central London accept only electronic payments, and a smartphone application is the predominant means of payment. However, the old model of prepayment for parking in coarse-grained time slots is still prevalent, even with electronic payments. Thus, while these applications can accurately capture the time a driver occupies a space, the main challenge is to discover when that space is vacated—in most cases, long before the payment has expired.

Motivated by this trend, we developed a smartphone-based sensing system, called ParkSense,³ that can detect when a driver vacates a parking spot previously paid for through the smartphone, with minimal impact on the lifetime of driver's mobile device. When coupled with payment systems, this approach can provide a complete system for sensing on-street parking-space occupancy at the urban scale in a cost-effective manner.

Detecting that the user has vacated a parking spot can be decomposed into two subproblems:

- sensing that the user has returned to the parked vehicle location, and
- detecting that the user is in a moving vehicle after returning to the parking location.

The requirement for minimal-impact sensing prevents us from using traditional location sensors available on

smartphones, such as GPS or network-based location sensing, due to their high energy consumption. ParkSense thus relies only on the smartphone's Wi-Fi interface.

Similar to traditional systems (see Figure 1a), smartphone-based parking payment applications prompt the user to input a location number associated with the set of spaces for payment (see Figure 1b). These location numbers are displayed on nearby signs (see Figure 1c). A user must complete the payment before leaving the parked vehicle. ParkSense performs several Wi-Fi scans to capture nearby Wi-Fi access points while the user waits for the payment process to complete. ParkSense can then detect if a user has returned to parked car location by periodically comparing observed access-point Service Set Identifiers with those captured during the payment process. This coarse-grained location sensing only involves performing Wi-Fi scans. It is energy efficient because it doesn't require communication with external databases, as with network-based location sensing or tracking satellite signals using GPS.

Figure 1. Parking payment systems: (a) traditional and (b) modern parking payment systems. Location numbers needed for payment are displayed on (c) nearby signs

Sensing that a user has started driving after returning to the parked location involves performing continuous Wi-Fi scans. ParkSense computes the dissimilarity D between successive scans as $D = 1 - J$. The Jaccard index J is merely the ratio of the number of same access points observed in both scans to the total number of access points observed in the two scans. //What is the differences between the two underlined statements? **First one refers to number of APs in the intersection of two scans, second one refers to number of APs in union of two scans**// The dissimilarity D serves as proxy for the user's speed, which is compared to the normalized GPS speed (see Figure 2). As the comparison shows, the dissimilarity between successive scans increases as the vehicle speeds up, and it decreases as the vehicle slows down. A vehicle moving in an urban environment moves in a start-stop fashion, as shown in Figure 2, so ParkSense uses these large variations in dissimilarity D to detect that the user is moving at high speed and has thus vacated his or her parking space.

Figure 2. Comparison of a normalized vehicle speed and the dissimilarity D between successive Wi-Fi scans. The dissimilarity between successive scans increases as the vehicle speeds up, and it decreases as the vehicle slows down.

In our trials of the prototype system, ParkSense successfully detected when a user had left a parking spot 34 times out of 41 possible instances, at several locations and with different users. Furthermore, ParkSense achieved this using a fraction of the energy costs of traditional location sensors. For example, it consumed approximately 11 percent of the energy that network-based location sensing consumes, and approximately 21 percent of the energy that GPS consumes.

However, there were certain cases in which ParkSense couldn't detect that the driver had vacated a parking space. Examples included when a user drove away from the parking space but couldn't make progress due to traffic, and when a driver drove away in a vehicle parked and paid for by someone else (say, a family member). Another example was when a user returned to the vicinity of the parked vehicle but rode away in another vehicle. Although these types of situations are rare, their effect could be further mitigated by providing street- or block-level coarse-grained space availability to drivers looking for parking. //How would this address these particular issues? **Coarse grained availability (high, medium or low) can mask small number of errors**//

ParkSense addresses an important problem of sensing on-street parking space occupancy, but several other issues must be overcome to operate such systems at an urban scale. For example, just providing occupancy information can lead to urban-scale dynamics in which all drivers move to regions with high space availability only to find full spaces by the time they arrive in these regions. It is also important to understand the supply-and-demand conditions under which such recommender systems can properly function.

Sensing Repeated Journeys

For a smartphone to issue a personalized traffic alert to a user, it must first be able to learn and recognize routes that the user frequently drives. It can then monitor these routes using a traffic-advisory service and notify the user of any disruptions reported on these routes. After learning the routes, if a routine journey were to take longer than usual, the phone could also send a message to a preconfigured set of contacts, informing them of the delay.

However, learning and recognizing these routine journeys requires continuous passive sensing that precludes traditional location services available on smartphones, such as GPS or network-based location sensing, due to

their high energy consumption. Although it is possible to require the user to plug his or her phone in the charging port of the car for energy-hungry location sensing during journeys, it is a stringent constraint that the user might not always remember.

Google Now, the context-sensing core of the Android OS, learns a user's home and work locations and then informs the user of travel times between these based on the shortest route (www.google.co.uk/landing/now). However, there is usually more than one route available between different locations, and which one is relevant to the user can only be determined with continuous sensing during journeys, which most applications appear to avoid to reduce their impact on battery lifetime. Significantly reducing the energy consumption of sensing routine journeys on smartphones thus seems to be the key for making the aforementioned applications a reality.

Our second system achieves this using low-energy microelectromechanical systems (MEMS) accelerometers and gyroscopes embedded in smartphones.⁴ The key insight that this system relies on is the fact when a car travels along a particular route, it executes a certain set of maneuvers (turns and bends) that form a unique signature of the route. These turns and bends can be easily detected by the phone carried inside the vehicle using its embedded gyroscope.

During a repeated journey along the same route, the phone observes the same unique signature, as the car executes the same sequence of turns. Because the phone can be carried in any arbitrary orientation in the car during different journeys, we propose a mechanism that uses the smartphone's accelerometer to sense the direction of gravity; the gyroscope is then used to measure angular velocity around an axis parallel to gravity. This lets our system capture a route's unique signature as variations in angular velocity.

Figure 3 shows two such angular speed patterns captured by the smartphone carried in a vehicle from a 12km long route in a typical urban city when the route is driven on two different days. The upper red trace is from a journey that took place on a weekend, and the lower blue trace is from a second journey on the same route on a weekday with evening traffic. It shows that angular speed variations originating from repeated journeys along the same route bear strong similarity.

Figure 3. Sensing repeated journeys. Two angular speed patterns captured by a smartphone carried in a vehicle from a 12km long route in a typical urban city, when the route is driven on two different days.

The nonuniform scaling of time axis in Figure 3 arises due to the vehicle traveling at different speeds along different road segments of the route. We use dynamic time warping (DTW) to compute the similarity between these angular speed traces under this nonuniform scaling. For example, Figure 4 shows the mapping created by DTW between samples of the two angular velocity traces shown in Figure 3. During a bootstrap phase, our system captures angular variations from all car journeys the user takes. These are then clustered using pairwise DTW similarities, and clusters containing numerous members are identified as routes that the user frequently takes. Using open-ended DTW, our system can recognize a previously identified frequent route during the course of the journey. This lets the system capture geographic trajectories of frequently used routes by switching on location sensing during any of the repeated journeys after the bootstrap phase.

Figure 4. Dynamic time warping (DTW) comparison. The mapping created by DTW between samples of the two angular velocity traces shown in Figure 3

In a prototype evaluation of 43 journeys originating from over 400km of driving under normal conditions, our system achieved similar performance as one that used GPS trajectories to identify repeated journeys. Our prototype implementation on a typical Android smartphone and a microcontroller board that emulated a motion coprocessor achieved an energy consumption of an order of magnitude less than the GPS-based system. Motion coprocessors are common in modern devices, such as Apple's iPhone and Google's Nexus smartphones, for continuous sensing using inertial sensors. However, software APIs that access these aren't yet available to application developers. We thus relied on an external microcontroller board to demonstrate the energy efficiency of our prototype system. However, the popularity of continuous motion-sensing applications means that it is likely that software APIs for motion coprocessors will become available in near future.

Driver Analytics

We now discuss other areas in which smart sensing can be used in cars for novel car-based applications. One such area is driver analytics—that is, gaining insight into a user's driving behavior using sensor data, which is of particular interest to vehicle insurance providers. Until now, insurers have relied on factors such as the driver's age, experience, and profession to assess risk. However, smart sensing can offer detailed sensor data on speed,

acceleration, and braking, which can be used to adjust these risk models to individual users and driving styles.

The vehicle insurance industry has incentivized users with reduced premiums to use either black-box devices in cars or smartphone apps to collect this data. Both of these approaches offer advantages as well as challenges. Black-box devices are free from energy and arbitrary device orientation issues. However, the use of dedicated hardware adds costs and presents a barrier to wide spread adoption.

Another challenge with this approach is vehicles that are shared between different users—such as a family car. In this case, robust driver identification algorithms are required to accurately label trips driven by different drivers. A smartphone-application-based approach is suitable for widespread adoption; however, high energy consumption and arbitrary device orientation are still major hurdles. Current applications thus require the user to attach the smartphone to a phone holder and only collect a sample of several miles of driving, rather than perform continuous sensing. Addressing these challenges can potentially lead to significant impact on vehicle insurance industry and other commercial applications of in-car sensing.

Smart in-car sensing systems present an emerging and exciting research area. Drivers increasingly carrying smartphones, with their plethora of embedded sensors, creates new opportunities for in-car sensing. Such in-car smart sensing systems can immensely improve the driving experience by reducing travel times and offering novel context-based applications. In addition, these systems can help with urban-scale issues of congestion and pollution and offer business opportunities based on new streams of sensor data. These opportunities present new challenges, including energy efficiency and novel knowledge extraction from existing low-level sensors, which the research community must tackle to unlock the immense potential of in-car smart sensing.

References

1. “2 Billion Consumers Worldwide to Get Smartphones by 2016,” *eMarketer*, 11 Dec. 2014; www.emarketer.com/Article/2-Billion-Consumers-Worldwide-Smartphones-by-2016/1011694.
2. D.C. Shoup, “Cruising for Parking,” *Transport Policy*, vol. 13, no. 6, 2006, pp. 479–486.
3. S. Nawaz, C. Efstratiou, and C. Mascolo, “ParkSense: A Smartphone Based Sensing System for On-Street Parking,” *Proc. 19th Ann. Int’l Conf. Mobile Computing & Networking (MobiCom)*, 2013, pp. 75–86.
4. S. Nawaz and C. Mascolo, “Mining Users’ Significant Driving Routes with Low-Power Sensors,” *Proc. 12th ACM Conf. Embedded Network Sensor Systems (SenSys)*, 2014, pp. 236–250.

//Please supply author photo//

Sarfraz Nawaz is a senior research associate in Computer Laboratory at University of Cambridge. His research interests lie in mobile computing and wireless sensor networks. He is particularly interested in applications of mobile computing in and around vehicles and transportation systems. He received a PhD in Computer Science and Engineering from University of New South Wales. Contact him at sarfraz.nawaz@cl.cam.ac.uk.

Christos Efstratiou is a lecturer in ubiquitous computing at the School of Engineering and Digital Arts at the University of Kent U.K. His research interests lie in the areas of mobile and ubiquitous computing and wireless sensing. His recent work considers the challenges of using mobile and wearable sensing technologies to detect personal and group activity behaviour, and their potential applications in the domains of architectural design, healthcare, and wellbeing. Efstratiou received a PhD in context-aware and adaptive mobile systems from Lancaster University. Contact him at c.efstratiou@kent.ac.uk

Cecilia Mascolo is a Professor of Mobile Systems in the Computer Laboratory, University of Cambridge. Prior joining Cambridge in 2008, she has been a faculty member in the Department of Computer Science at University College London. She holds a PhD and an MSc degree from the University of Bologna.

Her research interests are in human mobility modelling, mobile and sensor systems and networking and spatio-temporal data analysis. She has published in a number of top tier conferences and journals in the area and her investigator experience spans more than twenty projects funded by EPSRC, MRC, EU, Google, Intel, Microsoft and Samsung. She has served as organizing and programme committee member of over fifty

mobile, sensor systems and networking conferences and workshops. She sits (or has sat) on the editorial boards of IEEE Internet Computing, Pervasive Computing, Transactions on Mobile Computing and ACM Transactions on Sensor Networks. More details at www.cl.cam.ac.uk/users/cm542. Contact her at cecilia.mascolo@cl.cam.ac.uk.