ParkNet: A Mobile Sensor Network for Harvesting Real Time Vehicular Parking Information

Suhas Mathur, Sanjit Kaul, Marco Gruteser, Wade Trappe
WINLAB, Rutgers University, 671 Route 1 South, North Brunswick, NJ, USA
{suhas, sanjit, gruteser, trappe}@winlab.rutgers.edu

ABSTRACT

This paper describes the architecture and design of ParkNet, a mobile sensor network consisting of vehicles, which collects and disseminates real-time information about the availability of parking spaces in urban areas. We outline the broad challenges in real-time data collection and consumption by a mobile sensor network, as well as the design issues specific to ParkNet. Sensor nodes in ParkNet are a combination of privately owned and city owned vehicles and employ low cost ultrasonic sensors to detect the presence of vacant parking spots as they drive by. We present early results on the performance of our sensor platform.

Categories and Subject Descriptors

C.2.1 [**Computer Communication Networks**]: Distributed networks, Wireless communication

General Terms

Design, Algorithms, Experimentation, Measurement

1. INTRODUCTION

The importance of better roadways and parking systems in urban areas has recently been recognized as one of the most important avenues for betterment of urban infrastructure. In a study released in 2007, [1] it was estimated that congested roads in the U.S. cost $78 billion annually in the form of 4.2 billion lost hours and 2.9 billion gallons of wasted gas. A recent study [2] indicated that up to 45% of traffic on some streets in New York City is generated by automobiles circling the block. In [3], researchers found that over the course of a year in just one small business district of Los Angeles, vehicles looking for parking created the equivalent of 38 trips around the world, burning 47,000 gallons of gasoline and producing 730 tons of carbon dioxide.

The efficient use of parking resources in urban areas requires that real-time information about the availability of parking spaces be made available to automobiles on the street that are looking for parking. This in turn requires that parking spaces be monitored and users have a way of accessing this data while on the road.

Prior work that attempts to address the problem of parking in urban areas can be broadly classified into two categories. The first body of work aims to reduce the total amount of traffic on the streets in urban areas using mechanisms ranging from dynamic pricing of parking spaces [4] to fine grained congestion control via pricing [5]. The second body of work aims to monitor parking spaces by detecting the presence of parked vehicles over parking spots using sensors [6–8]. However, to the best of our knowledge, all efforts in this category rely on fixed sensors installed by municipalities in the ground or in parking meters. This necessitates a large fixed cost of installation and operation in order to cover parking spaces at a city-wide level. For e.g., the SF-park project [8] aims to cover only 25% of the street parking spots in San Francisco at a cost of several millions in funding. Further, this approach is limited by the fact that an end user must posses a device capable of connecting to the internet to access parking information. However, projects such as [8] highlight the magnitiude of the problem in large cities and the governerment's dedication to long term investments in a smart parking infrastructure.

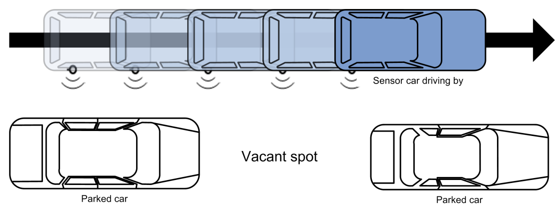

Figure 1: Ultrasonic sensor fitted on the side of a car detects parked cars and vacant spots.

The ParkNet project aims to disseminate information about parking space availability in a real-time manner by employing on-road vehicles to detect vacant parking spots as they drive by (Figure 1). In order to sense vacant spots, vehicles utilize an ultrasonic sensor mounted on the vehicle's side and a GPS receiver that notes the corresponding location. Ultrasonic sensors are widely used in robotics and industrial applications as 'range finders' or 'position estimators'. Their widespread use has meant that even high-performance sensors have a very low cost compared to RADAR or LIDAR sensors. In fact a number of new automobiles are already fitted with one or more ultrasonic sensors for automated parallel parking or various distance warning applications [9, 10] and increasing numbers of new vehicles are expected to be equipped them. The aim of ParkNet is to exploit the presence of on-board sensors to provide information about available parking spots to vehicles looking for parking, in a manner that is as real-time and as reliable as possible.

2. RESEARCH CHALLENGES

The goal of gathering, processing and disseminating sensor data for immediate use is an ambitious one. Apart from the problem of obtaining reliable sensing data, which carries significant challenges by itself, the dissemination of this information in a timely fashion poses a number of interesting

MobiHoc S3 '09, May 18, 2009, New Orleans, Louisiana, USA.
ACM 978-1-60558-521-5/09/05.

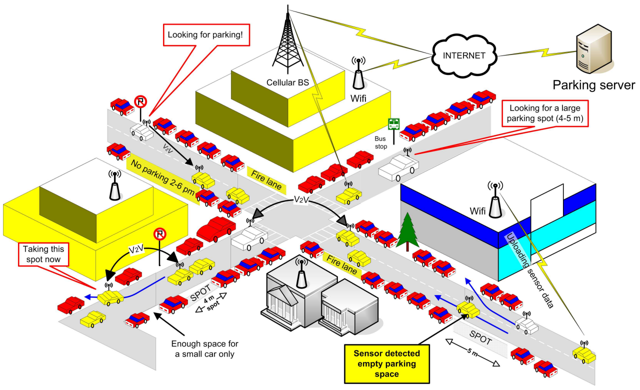

Figure 2: Yellow cars are are equipped with sensors, white cars are looking for a parking space and red cars are parked on the roadside. A vacant space may not be a legal parking spot or may not be large enough for some vehicles. This figure is best seen in color.

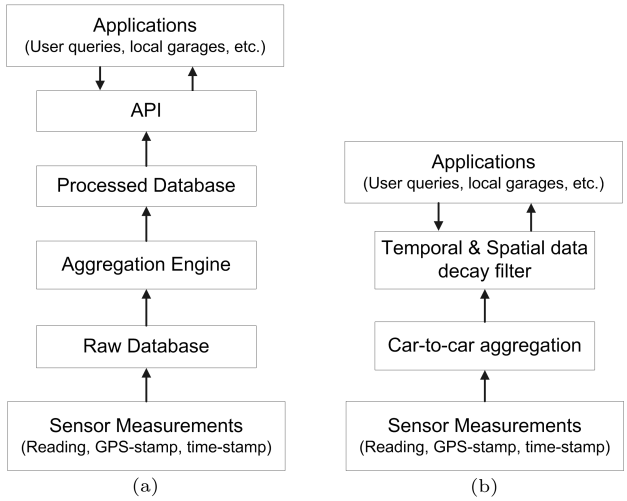

Figure 3: (a) Centralized and (b) distributed architectures of ParkNet

networking and communications challenges. For instance, a network architecture in which all sensor vehicles send their measurements to a central server over a cellular data uplink affords a cleaner design but involves a cost for data transfer and requires end-users to posses a cellular data connection as well. We discuss the sensing and communications challenges in detail below.

2.1 Sensing

Reliability of data: In the ideal case, we would like each vehicle equipped with a sensor to accurately detect the location and size of vacant parking spots automatically as it drives by. However, the problem is complicated by the fact that sensing parked cars requires that the sensing car be driving in the right-most lane. Therefore, the sensing mechanism requires knowledge about the lane in which the sensor car is driving at any given time in order to avoid collecting spurious measurements. More often than not, this information is not possible to deduce using commercial GPS receivers. We propose to employ the use of machine learning techniques to resolve this hurdle. When multiple vehicles report measurements on the same street, the reports of some of the vehicles can be filtered out as coming from the wrong lanes. In addition, as the system learns the sensing profile of parked vehicles, it will get better at distinguishing between parked vehicles and vehicles passing by in other lanes. As the status of parking spaces changes with time, the design of an algorithm for extracting reliable information by aggregating data from multiple sensor vehicles is a non-trivial task.

Speed and range constraints: A second problem specific to ultrasonic sensing is that the speed of sound limits the frequency with which a sensor can measure range for a given fixed maximum range. This is because the sensor has no way of differentiating between different return pulses as they carry no ID information or encoded sequence number and so the sensor must wait for the echo pulse to return or it must timeout before sending a new pulse. Our sensing platform, for example, provides a maximum range of 6.42 m and makes measurements only once per 50 ms. At a speed of 35 mph (typical speed limit on streets with side parking), a stretch of 4 m accommodates about 5 measurements. Further, since the sensing vehicle moves forward by the time the reflected sound waves return to the sensor, this implies a constraint on the vehicular speeds at which reliable sensor measurements can be made. Otherwise the sensor will miss the echo. We believe this problem can be alleviated to some extent by using a sensor with an appropriate beam width. Since each vehicle knows its own speed and the capabilities of its sensor, this would allow us to assign a reliability score to its measurements that we expect will aid in the aggregation of data from multiple sensors.

Learning true parking spots: Finally, measurements from an ultrasonic sensor cannot differentiate between a legal parking spot and an empty space where parking may not be allowed (say, because of a no-parking sign, a fire hydrant or simply an entrance to a driveway). We aim to address this problem by having the system adaptively learn the locations of legal parking spaces with time. This can be based on the simple observation that spaces that are occupied more often are more likely to be legal parking spaces.

2.2 Communications

Centralized architecture: We first consider a baseline architecture (Figure 3(a)) in which every car is assumed to possess a cellular data uplink available, as in [11]. As sensor data is collected, it is immediately uploaded onto a central server over the Internet. An end user that wishes to query the system must also have a way of accessing this server over the internet (say, using a smartphone with a cellular data plan). Sensor measurements from each sensor vehicle are time-stamped and location-stamped and are sent to the server where data from various mobile sensors is continuously aggregated and processed as it arrives. The database provides an API to applications for accessing data. For e.g., queries arriving from end-users seeking parking may be replied to by providing the closest 3 or 4 vacant spots, each with it reliability score, or parking garages may be supplied with the overall state of availability of street parking so they can adjust their prices dynamically.

Distributed architecture: We next consider an architecture aimed at allowing end users without cellular data connections to obtain relevant parking information by processing queries within the local network of vehicles. In this archi-

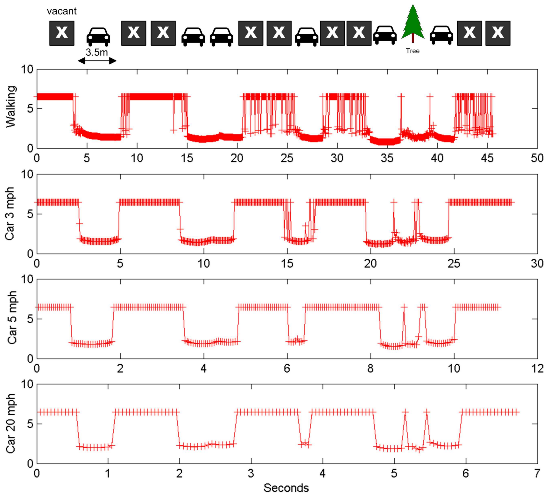

Figure 4: Raw traces for range measurements (in meters) made in a parking lot with well defined parking spots, walking and in a car at speeds of 3, 5 and 20 mph.

tecture (Figure 3(b)), vehicles exchange time- and location-stamped sensor data and each such vehicle, in addition to serving as a sensor, also serves as an in-network storage node. Queries submitted by users are broadcast to other vehicles in the area, possibly over more than one hop and replies are sent back to the origin of the query as quickly as possible. This is a challenging architecture as it requires communications over an ad-hoc vehicle-to-vehicle network. However, it is well worth exploring because of two reasons: (i) The cost of a cellular data connection has still not approached a level at which all users participating in our system can be assumed to have a subscription, and (ii) a concerted research effort is under way to develop vehicle-to-vehicle communication systems and vehicles are expected to soon be equipped with the first generation of such systems. The marginal cost of deploying a system for exchanging sensor data over such a communication system would be small and most implementation might infact be possible in software, thereby avoiding the high fixed cost of a cellular connection. Parking information traveling through the network has a temporal and spatial relevance. That is, the significance of the data decays in time as well as with distance from the location at which the measurements were made. Therefore, we must devise a method that lets each vehicle decide what data to share with other vehicles - it is necessary to discard less relevant data in order to make the system scalable, otherwise the amount of data in the network would grow without bound. However, the manner in which data must be assigned relevance is not obvious, and further, the manner in which data with lower levels of relevance should be discarded by vehicular nodes needs to be careful designed. We refer to this as a 'temporal and spatial decay filter' in figure 3(b).

WiFi APs and phones: A further variable in the design of the communications architecture of ParkNet is the possibility of opportunistically leveraging the presence of fixed WiFi connections that are commonplace in office and residential areas. There has been a significant research effort aimed at exploring the use of fixed 802.11 for mobile vehicular communications and the protocol optimizations that this entails (connection setup time etc.) Additionally, new mobile phones are increasingly equipped with a WiFi in-

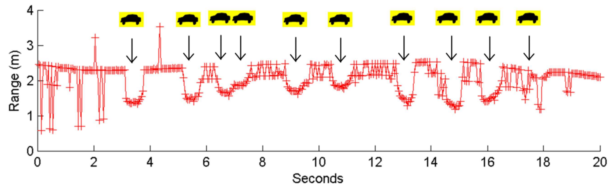

Figure 5: Raw trace from a drive-by at 25 mph in a residential area with no well defined parking spots. All but the last car on the street are clearly visible in the measurements.

terface. This introduces the possibility of employing users' mobile phones in a peer-to-peer manner to submit queries and exchange sensor data without always needing dedicated vehicular radios.

Finally, a parallel goal of ParkNet is to better understand and address the challenges in the development of vehicular networks and their integration with heterogeneous wireless networks in urban areas, particularly 802.11 wireless LANs and cellular data networks.

3. SENSING PERFORMANCE - EARLY RESULTS

Our sensor platform consists of a single ultrasonic transducer [12] that emits pulses of sound waves at a frequency 42 Khz every 50 ms. The sensor provides a single range reading every cycle, expressed in inches, and is capable of measuring ranges between 6-255 inches. The sensor provides data in serial format that can be accessed through a serial terminal on a laptop. Figure 4 shows traces from our sensor corresponding to a row of parked cars in a outdoor parking lot. Figure 5 shows a trace collected in a residential area without clearly demarcated parking. Our early experimental data indicates that the placement and beam-width of the sensor are important parameters and must be properly chosen. Further, as expected, sensing performance degrades in crowded areas with speed of the vehicle. Results obtained in open areas such as parking lots are much cleaner as automobiles are the only objects present.

4. REFERENCES

[1] "*Texas Transportation Institute*, Urban mobility report," 2007.
[2] "*Transportation Alternatives*, No Vacancy: Park slope's parking problem and how to fix it," 2007.
[3] D. Shoup, "Cruising for parking," *Access*, vol. 30, pp. 16–22, 2007.
[4] ——, "The Price of Parking on Great Streets," in *Planetizen Contemporary Debates in Urban Planning, Washington, DC*, Island Press, pp. 52-56., 2007.
[5] "Driving change in stockholm, IBM How it works series, Oct 2008."
[6] "http://www.streetlinenetworks.com."
[7] "http://www.vehiclesense.com."
[8] "http://www.sfmta.com/cms/pproj/sfparkover.htm."
[9] "Parkpilot, http://www.bosch.com."
[10] "Electronic parking assistant, http://www.hella.com."
[11] J. Eriksson, L. Girod, B. Hull, R. Newton, S. Madden, and H. Balakrishnan, "The Pothole Patrol: Using a Mobile Sensor Network for Road Surface Monitoring," in *The Sixth Annual International conference on Mobile Systems, Applications and Services (MobiSys 2008)*, Breckenridge, U.S.A., June 2008.
[12] "http://www.maxbotix.com."