Sentient Bikes for Collecting Mobility Traces in Opportunistic Networks

Marcelo Pias⋆, Kuang Xu⋄, Pan Hui†, Jon Crowcroft⋆,
Guang-Hua Yang⋄, Victor O.K. Li⋄ and Salman Taherian⋆

⋆ University of Cambridge ⋄ The University of Hong Kong
† Deutsche Telekom Laboratories/TU-Berlin

ABSTRACT

We study the problem of building low-cost city-wide location tracking systems with the intention to provide a platform for large-scale human mobility data collection. We take the city-bikes as our first target, motivated by several social networking applications, and choose Cambridge in the UK as our pilot case. We highlight the main application requirements, discuss the practical design issues, and propose a system architecture based on a hybrid sensor network.

1. INTRODUCTION

Recent years have seen the proliferation of powerful portable devices that combined with human mobility enable a new networking environment for store-and-forward and hop-by-hop *opportunistic communications* within the cities. These opportunistic networks are different to the traditional communication networks in which end-to-end communication links are assumed. Hence successful research and design of mobile communication systems require an understanding of human mobility, and this requires access to mobility data. However, the largest publicly available human contact traces contain about only 100 sparsely connected nodes. Therefore it is imperative to find practical and effective solutions to collect large-scale human mobility traces.

Several bike-sharing schemes (often called *community bicycle program*) have been introduced in some European cities (Barcelona, Lyon and Paris[1]). These schemes attracted tens of thousands of subscribers within few months. A cyclist produces far less carbon emissions, thus helps to reduce our carbon footprint. Bicycles are inexpensive when compared to vehicles, and they provide a means of physical exercise for their users. Thus, we believe that mobility traces from cyclists would be a valuable resource for the research community in mobile systems design.

Starting from here, we ask the question whether a low-cost location tracking system for large-scale traffic of city-bikes is possible. We would instrument bikes with sensors, and track them continuously. Despite technical details, a bottom-up practical and privacy-driven approach is needed. We motivate our study by proposing the following city-wide social application senarios:

Social networking *Exploring encounters*.
With the continuous locations of bikes, a cyclist can link up with others he passes by, and establish new or strengthen his social connections. For example, the positioning system can be linked to social networking websites such as Facebook and Twitter. If the system finds two registered participants in close proximity it updates their status in Facebook. Also, the positioning system can calculate the participants' cycled distances and provide incentives (e.g. virtual credit) for those who travelled the most in their daily cycling. This enables comparisons of travel profiles among friends and other cyclists.

Security *Bike theft*.
The public bike-sharing schemes have been receiving considerable attention in Europe. To check out and return a bike, a user swipes his payment/membership card at the bike stations spread throughout a city. To promote the use of bikes instead of cars, the programme's organisers offer affordable rental to their users. It has been reported, however, that the rental scheme in Paris has run into problems just 18 months after its launch. Over half of the original fleet of 15,000 custom-made bicycles have disappeared, and presumed stolen. Since the launch of this scheme nearly all the original bicycles have been replaced at a cost of 400 Euros each. We hope that there is a practical, city-wide location tracking system that can effectively help authorities tackle bike theft in community bicycle programs. And if we can build such a system at a low cost (say 10% of the Paris bike replacement cost), that would be a step forward.

In this paper we propose a general platform for large-scale collection of human mobility data with a particular focus on the security scenario. We want to target a city-wide deployment, with low-cost devices, and we do not require high accuracy for the position. The design of such a system needs a number of considerations. There are several possible competing approaches.

The contribution of this paper is a discussion of the practical problems of deploying a low-cost bike positioning system. We propose a system architecture that relies on a hybrid sensor network (an infrastructure with mobile sensor nodes). We advocate that the best approach for us to understand wireless sensor system is building them (bottom-up practical experimentation). Thus, we take an application-driven system design approach that favors simple solutions to our issues. We share the same view as the one put forward in [11] where we design systems and protocols for one specific scenario after another, and after a few solutions, we look for generality

[1] http://en.wikipedia.org/wiki/Community_bicycle_program.

ACM HotPlanet '09 Krakow, Poland

across the specific. The bicycle network of Cambridge (UK) is our application domain.

2. APPLICATION REQUIREMENTS

We choose Cambridge as our city since it has the highest level of cycling in the UK with one in four residents cycling to work[2]. The city-bikes need to be constantly tracked from a central location. We equip each of the bikes with a tiny battery-powered sensor device that can collect sufficient data to estimate location. The location update rate is in the order of tens of seconds. The bike sensor device needs to continuously send sensor data to back-end servers. The system requirements are as follow:

- **Sensing:** location of bikes (e.g WGS84 coordinates) within a city map should be computed within an accuracy of 50 metres. Inertial MEMS sensors can be integrated to improve the accuracy of location estimates through dead reckoning. Such sensors can also assist in optimising the system power consumption, for instance, by switching off sensor components when the bike is not in motion.

- **Communication:** the network coverage is around 2.5 km x 2.5 km (Cambridge city centre). The infrastructure should have back-end servers, and around 30 fixed nodes. We aim at instrumenting approximately 200 bikes.

- **Real-time versus data logging:** the location is updated every 100 metres for real-time tracking. However, finer-grained sensor data should be logged on the device for post-processing and reliability.

- **Sensor system packaging:** small size and compact form factor to be fitted on a suitable part of the bike (e.g frame or hub) or inside the light case.

- **Battery life:** it should last for at least a month on average commuting journeys (an hour per day). Energy harvesting can be exploited from the bike's motion (wheel dynamo system) to increase battery lifetime.

- **Security:** it should provide an effective way to avoid damaging to sensor devices, likely to be the case when addressing bike thefts.

- **Cost of the system:** Ideally the cost should be around 35 US dollars for a sensor device to be placed on the bike.

3. RELATED SYSTEMS

The most commonly available location technology today is the Global Positioning System (GPS). Although accurate and very effective in open environments, GPS alone does not support two-way data communications, which is required by a tracking application. We anticipate, however, the use of GPS integrated with wireless communication in some of our systems.

The city-bike location discussed here concerns mostly outdoor location systems and techniques. Although a number of suitable indoor location systems exist, in this section we discuss only those that employ RF signals in their engine.

At first distributed location algorithms (e.g DV-hop [10] and four-stage [14]) seem favourable to our city-wide positioning. But the majority of these techniques lack practical deployments and evaluations. They have been analysed mostly in simulations. However, there are a number of practical deployments of location systems that rely on an infrastructure of servers for information aggregation and position calculation, and they are mostly limited to single hop localisation from a set of known access points (APs). Such approaches improve accuracy, reliability, and dependability of the system.

The wide adoption of WiFi-enabled mobile devices (e.g. smart phones, PDA) and the rapid deployment of WiFi access points make WiFi localisation attractive. The RADAR project [1] pioneered indoor WiFi location, using WiFi "fingerprints" previously collected at known locations inside a building to identify the location of a user's device down to 2-3 metres (median accuracy). The fingerprint algorithm assumes that at a location point, a user's WiFi device may receive beacons from different access points with certain signal strengths; this set of APs and their associated signal strength represents a *fingerprint* that should characterise that position. The fingerprint algorithm is split into two stages. In the first, the radio map at different locations are built by using methods such as war-driving in open spaces, and in the second stage (position phase), a device performs a scan of its environment, calculates its relative location with regard to the surrounding APs (such as in [12]), and compares this with its known radio map . The techniques that address the issue of APs that have been deployed after the war-driving radio mapping are discussed in [3].

UCSD's Active Campus project [6] employs WiFi to locate devices inside and outside buildings based on a simplistic algorithm that relies on known positions of access points on a university campus. Place Lab [7] introduced wide-area WiFi location, showing a median accuracy ranging between 15 and 60 meters and high coverage. Place Lab depends on *war-driving* data collected by a variety of users as they move naturally throughout a region. Place Lab is intended for use in metropolitan-scale deployment, and it is useful for users with smart devices. But if we want the end devices to be low cost, it may not be an appropriate choice.

The BikeNet project [4] addressed environment monitoring applications with various types of sensors attached to bikes. It also studied the sensors for inter-bike networking through data muling. Unlike BikeNet, this work focuses purely on the tracking of objects using scalable and low cost mobile sensor devices. We approach these issues with a dual-function system, where bikes are tracked around a city area, and this data communicated wirelessly to an infrastructure. Bicing is a bicycle-sharing programme in Barcelona. In [5], the authors based on a six-week dataset of observations, introduced a notion that the digital footprints of cyclists could be used to uncover human behavior patterns and city dynamics. However, they do not have information on where the bikes travel between bike stations. Our work provides the system to upload locations of bikes in real time, which serves as complementary data to their work. In addition, this can increase the operational efficiency of bike-sharing programmes and reduce the rate of bike theft.

4. THE SENTIENT CITY-BIKE

The Sentient City-Bike system is designed according to an event-based layered architecture. City-bike mobile sensor nodes publish event sensor data that are pushed towards trackside servers. The sensor data mainly comprises signal strength of WiFi AP beacons. Inertial sensor data may be collected in selected subsets of sensor nodes. This can be used for strapdown inertial navigation techniques.

To reduce costs, the majority of the mobile node devices have limited processing power, memory and battery lifetime. The trackside servers are sufficiently powerful to perform the location es-

[2]Cambridge City Council: http://www.cambridge.gov.uk

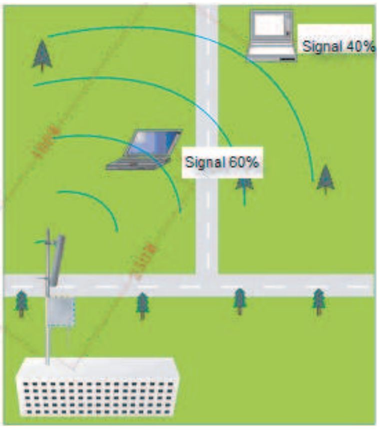

Figure 1: Measurement setup (AP 4000-MR-LR)

timation by correlating the AP information with the available AP database using some of the techniques discussed in Section 3. The system architecture is split into three subsystems: the communication infrastructure, the central server, and mobile nodes.

4.1 Communication infrastructures

We assume there is a gateway that delivers data wirelessly between the bike sensor system and trackside servers on the infrastructure. We were presented with a number of industry standard wireless systems (Bluetooth, Zigbee, Wifi). Bluetooth requires device discovery and connection establishment which causes problems for real-time data collection. Zigbee is designed for low power and low sampling rate (250 Kbps) and short range (up to 30m) applications. Our practical attempts with ZigBee radio from various makers failed mostly because of a lack of robustness.

WiFi is more on the expensive side exhibiting higher power consumption compared to the other two options. But WiFi power consumption has improved with recent developments for smart phones (iPhone and Google Android). Manufacturers (e.g. Marvell and Atheros) have developed low power chips, and Gain Span (Intel spin off) has claimed a WiFi sensor node that can last for years on a single AA charge. Of course, this assumes an efficient scheduling of tasks on the sensor node (sampling, processing, data logging and data transmission). Another benefit of WiFi is its ability to support data communication with an extended range. Data rates of 54Mbit/s in 802.11a/g and beyond in 802.11n exceed the required throughput of general tracking system. The extra bandwidth can be exploited by other applications.

We consider WiFi Mesh that is a communications network made up of WiFi radio nodes organised in a mesh topology. Mesh networks often consist of clients, routers and gateway. In our system, the mobile nodes are the mesh clients, and the infrastructure trackside nodes form the mesh routers and gateways. Table 1 lists a few options of mesh routers. We carried out some initial tests to understand the communication coverage of a mesh router. We used an AP4000-MR-LR (24dbm) with 18dbi flat board antenna in an open space (Figure 1). The antenna was placed on the roof of a ten-story building. A remote user carrying a WiFi device was able to watch a movie in real-time at a distance of 250m, and surf the web at 400m. The communication range was 800m between any two mesh backhaul router. These preliminary results are encouraging.

4.2 Central Server

The central server is used for long-term storage and computation. Since it has powerful resources, we leave the computation of location and other sensor-related analysis to this subsystem. It offers a web interface to end-users with enhanced map applications. The server provides web services to make location and sensor derived information available in social networking websites (Facebook, or Twitter for quick location updates).

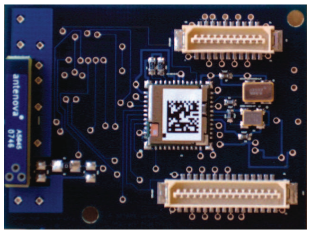

Figure 2: Cambridge Sensor Kit (CSK) WiFi board

4.3 Mobile sensor node

Our system has two types of lightweight and small physical size nodes. The basic node has essential sensing capabilities. Super nodes are more powerful devices in terms of processing, sensors and memory. We equip a fraction of super nodes (between 15% and 30%) with GPS receivers and other types of sensors. We select those nodes with potential to be active throughout the coverage area. They should be able to do complementary fingerprint data collection, so that the AP central database is kept up to date. This is an on-going task to complement the initial fingerprint data collection as part of an existing infrastrucutre/database. Super nodes equipped with GPS should also provide additional calibration data to verify and further improve the accuracy of location information gathered from the other fraction of nodes.

The sensor nodes together with lightweight batteries (10 grams) will be placed inside the plastic bike light casing, frame or hub. Some of the nodes may include inertial on-board MEMS sensors (accelerometers, gyroscopes, and magnetometers). Such inertial data can be fused with WiFi Fingerprint information (basic and super nodes) and GPS data (super nodes) through an Extended Kalman Filter (EKF). Techniques similar to the ones employed in strapdown inertial navigation systems can be used here to improve the accuracy of the location estimates.

We looked at various WiFi-based sensor systems (GainSpan, G2 Microsystems) but chose to discuss the Cambridge Sensor Kit (CSK), and mobile phones as first candidates for mobile nodes.

4.3.1 *Cambridge Sensor Kit (CSK)*

The UCAM-WSB100 [3] is a low power WiFi-based sensor board designed in the Computer Lab (University Cambridge). This board is part of the Cambridge Sensor Kit (CSK) that has been developed for a sports sensing application [4]. The CSK WSB100 is compatible with the Crossbow Imote2 processor board. It has a 12-channel ADC and supports 802.11 b/g with power management. We expect the CSK WiFi system to work adequately in our city-bike application. And if customisations are needed they can be made since the CSK hardware design and Linux-based software are open source.

We did some preliminary tests to examine the WiFi connectivity of the CSK. The sensor kit was placed in a bicycle light casing, and WiFi signal strength measurements were taken as a cyclist approached and passed a WiFi base station in an urban environment (Figure 4). In this setup, we used an off-the-shelf D-Link WiFi AP with a single 9dB antenna, elevated to the 1.2m off the ground. Figure 3 shows these measurements as the cyclist approached and

[3] http://imote2-linux.wiki.sourceforge.net/UCAM-WSB100
[4] SESAME: http://www.sesame.ucl.ac.uk

AP Name	Tx Power	Rx Sensitivity	Frequency Band	Aerial Gain	Price (USD)
ORiNOCO AP400-MR	20dbm (radius coverage up to 300m)	NA	2.4G (client) 5.47-5.725G (backhaul)	8/10/18dbi	\$800 + \$200 (aerial)
Motorola IAP4300	27dbm	-100dBm@1Mbps	2.4GHz(client and backhaul)	8/10/18dbi	\$2000

Table 1: Options of mesh routers

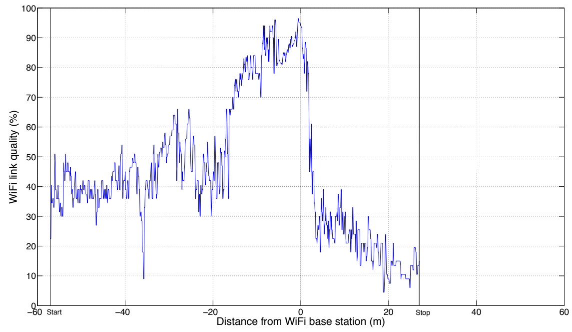

Figure 3: CSK WiFi Signal Strength

Figure 4: Bike Experiment Setup

passed the base station (located at the zero metre landmark). The 'start' and 'stop' labels on the diagram indicate the direction of the motion. As evident, the WiFi signal strength improves as the cyclist approaches the base station, and it drops when the cyclist passes the base station. The rate of increase and decrease in the signal strength is uneven, as the WiFi signal changes characteristics from line-of-sight connectivity to non-line-of-sight connectivity. In the former setting, there is a clear signal path between the base station and the sensor system (light case facing the base station). Whereas in the latter setting, the human body blocks and substantially attenuates the main signal propagating from the base station to the CSK sensor node.

4.3.2 *Phone as Mobile Sensor Node*

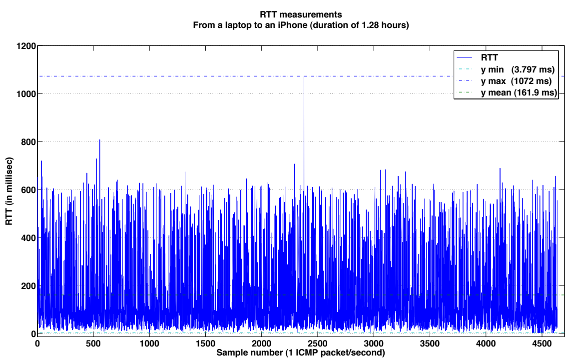

Figure 5: iPhone RTT measurements.

Given the fact that almost everybody has a mobile phone, we also consider using mobile phone devices as our sensor nodes. The latest smartphones (iPhone, Google Android G1) include a number of the sensor systems we discussed above (GPS receivers, accelerometers). It is expected that other types of sensors will be added in future phones. We tested an iPhone in the same urban setup as the one described for the CSK (above). We *ping* the iPhone from a laptop. The mean RTT value was around 200 ms at 250 m distance. But we observed that when pinging the basestation from the iPhone the mean RTT dropped to around a sensible value (3ms). This suggests the iPhone OS process/task scheduling is giving low priority to IP protocol processing when the foreground application is not any network function. We did not observe packet loss within the 250 metres range. In another experiment we collected data by running ping from a steady laptop against the iPhone, all connected to a local basestation. Figure 5 shows the RTT measurements from a laptop to an iPhone with WLAN for an hour (a sample every second). The mean RTT was 150 ms, and the difference between minimum (3 ms) and maximum (1000 ms) is worrying.

4.4 WiFi Coverage

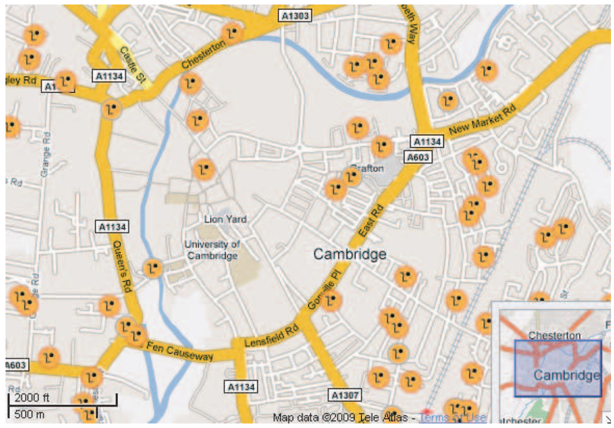

Figure 6: FON WiFi Access Points in Cambridge

WiFi Fingerprint scheme uses not only public but also private APs, whenever these are available. The scheme does not require proprietary support from the WiFi infrastructure vendors. The major effort needed is the war-driving to build a database of known APs and their locations. We expect this work to take a month or so. Cambridge is a University town with a number of residences, and we expect a high density of WiFi APs. Figure 6 shows the FON WiFi/AP coverage in Cambridge. We can see that even if we only rely on the FON APs, the city is reasonably well covered.

4.5 Power consumption

The issue is building a low-cost location system that would not require frequent charges of battery. This poses a challenge to the form factor and size of the device since continuous power consumption at high data rates requires a larger battery capacity, which is proportional to battery physical size. WiFi is commonly regarded as power hungry. But recent low voltage hardware designs have taken the consumption down.

The trade-off between data rate, power consumption and com-

Chipset	IEEE MAC	RX Pw	TX Pw	Data Rate	Energy/bit
Chipcon CC2420	802.15.4	19.6mA x 1.8V $\approx$ 35.28mW	17.04 mA x 1.8V $\approx$ 30.67 mW	250 Kbps	TX: 122 nJ/bit
Marvell 88W8686	802.11b/g	174.87mA x (1.8 and 3.0V) $\approx$ 320.61 mW	257.05 mA x (1.8V and 3.0V) $\approx$ 593.55 mW	11 Mbps	TX: 53 nJ/bit

Table 2: Preliminary power measurements

munication range is a key issue to be addressed. We did some back-of-the-envelope calculations that may be used as an indication of the power draw for the Marvell 8686 chip used in both CSK and iPhone. We argue that any power consumption (cost) analysis should consider the data rate (benefit). Table 2 makes a comparison between IEEE 802.15.4 and 802.11 based on a metric that correlates cost/benefit (**energy per bit**):

$$E_{bit} = \frac{P}{b}$$

where P is power consumption and b bit rate. It can be seen that the energy per bit for the CC2420 (IEEE 802.15.4) is twice as high than the Marvell 8686 (IEEE 802.11 b/g). This is due to the fact that the 802.11 achieves much higher data rates (higher benefit). This is very conservative because we assume the 802.11 is operating at 11 Mbit/s, where it can achieve up to 54 Mbit/s (802.11g). Both systems have similar power budgets while in power down modes. The time cost for switching from the idle to sleeping modes (power down) is likely to be different though, with higher latencies expected for the IEEE 802.11 b/g.

We believe that the real issue is the design of an improved duty cycle scheme for the sensor node system - a compromise in the scheduling of tasks (sampling, processing, logging, and data transmission). This issue needs to be further investigated in our practical application scenario. In addition, we expect to be able to convert some of the mechanical energy from the bike's motion, and store it in the sensor node system [9]. For example, we can put the sensor in the wheel-hub to make it powered by a dynamo system. Whether this energy will be sufficient to power up our entire system remains to be investigated.

4.6 Synchronisation of heterogeneous systems

Data samples collected from the various city-bike sensor nodes should be correlated for either real-time monitoring or post-processing analysis. This brings the issue of synchronisation of different systems. The complexity of the problem tends to increase with strict accuracy requirements. In practical terms, we deal with a hybrid wireless network, partly formed by mobile sensor nodes on city-bikes that exchange information with infrastructure trackside servers. The synchronisation is needed across these heterogeneous systems for real-time, and post-processing analysis. A number of time synchronisation protocols for wireless sensor networks have been proposed in the literature [13]. However we seek simple solutions to synchronise the clocks of the systems relative to a common reference time (say world time). Time synchronisation will be required within 100 ms (10 samples/second sampling).

We anticipate that sensor data collection should also benefit from this synchronisation. For instance, data collection may have to be smarter in downloading data from crowds of bicycles, for instance, during the morning rush hour in Cambridge. Even if there is reasonable network bandwidth available (say 54 MBps in 802.11g), the burst of data may overwhelm the system in such situations. Thus sensor nodes may have to postpone data download to the next data collection point simply because the current spot (say an AP close to a traffic light) is overwhelmed. To work in this solution space, we require synchronisation of systems across sensor nodes, and infrastructure nodes.

4.7 Privacy-compliant system design

We anticipate that there will be a number of technical (system-related as discussed above) and non-technical challenges (social and legal) when a city-wide location system is deployed. This requires the expertise from engineers, social scientists and lawyers. From a city planner's view, a location system is likely to serve different purposes, mostly related to the tracking of people and vehicles in the city so that better facilities and services can be rolled out. Most of such services require collection, storage and processing of a huge amount of personal data. This brings the issues of privacy, data protection and trust - all crucial for the development of location-based applications. A central authority has access to the location information of each individual. This poses an obstacle to the deployment of such a system as users may be discouraged in taking part [2]. Incentives and privacy measures should be introduced to motivate the users. Users should trust our systems, and it is crucial that we make transparent how the user's data will be processed and used. Once trust is established, incentives for user adoption of our systems will be introduced with the integration of user location data to online social networking (e.g. Facebook and Twitter).

However, we understand that early adoption of privacy issues in the design of our low-cost city-wide localisation system should be encouraged. Otherwise, later corrections through regulatory measures can be expensive to implement. In this context, privacy should not be dealt with as an add-on function to the system, e.g. by data filtering and minimisation measures. An approach here is to take an iterative design with user feedback using the social science technique of realised scenarios proposed by Martin et al. [8]. Realised scenarios are real technical installations, combined with social situations of an application, on a low level of technical sophistication and social realism. With this technique the users acceptance of different modes of data collection and processing can be modelled as a give-and-get game. It could be determined what type and amount of sensitive information users are *willing to give* to the technical system with respect to the *reward* they expect to receive. Instead of an overall scepticism towards these technologies, this approach could lead to a scaled and more situation-aware analysis, which will not require an overall consensus.

5. COST AND SECURITY

We want to start with 200 sentient city-bikes instrumented with CSK sensor nodes at unit cost of around US$35, which gives a total of US$8,000 in sensor devices. We also plan to use 30 WiFi access points that can cover the Cambridge city centre (2.5 km x 2.5 km). The ORiNOCO AP400-MR AP costsUS$1,000 per unit, totalling US$30,000. This is a high initial setup cost which is likely to be amortised when more bikes are instrumented in this scheme. Another option is to use inexpensive wireless routers (e.g. LinkSys WRT54G) at US$50, and 30 routers should bring the cost down to US$1,500. Table 3 gives the cost breakdown. In addition to the infrastructure, the software development (i.e. mobile sensor node, AP, and war-drive) is another considerable cost. However, by resorting to the location database of SKYHOOK [5], this cost item is

[5]SKYHOOK, http://www.skyhookwireless.com/, provides public APIs to access its location database.

negligible.

These figures are independent of the number of bikes. With these estimates the cost per bike is US$53, which is almost twice of the intended cost of US$35 as discussed in Section 2. But we expect the unit cost to decrease significantly when the system scales up. Nevertheless, US$53 remains a fraction of the 400 Euros cost per bike stolen on the Paris city-bike program.

Mobile node with CSK	US$8,000
Infrastructure	US$1,500
Central server	US$600
Basic cost in total	US$11,600

Table 3: Cost breakdown

For the security of devices and bikes, we plan to install the device inside the hub of the bike's wheel, making it sufficiently hard to remove without damaging the bike. Below we discuss several options to secure bikes against thefts:

- For the bike-sharing scheme, most systems require a user to return a bike to a *well-known* place (e.g. bike kiosk). We can place a software-triggered alarm in the server or in the WiFi APs nearby. If the user cycles outside an area of normal cycling the alarm is fired and the system reports the latest known location.

- Thieves typically take bikes somewhere to remove various distinctive markers for re-selling. The system may provide hints to the location of such professional thieves' garages. This information may assist the police authorities in reaching those places.

- The tracking device can be linked to distance or time-based charging for instance, via surcharging a credit card if a malicious user wants the bike to go more than a short way within a short time, then the bill can be set prohibitive. In a stronger approach for this situation, it is interesting to instrument a bike with a immobilizer controlled by the tracking device, to keep the bike from further proceeding.

6. CONCLUSIONS

In this paper we made an attempt to discuss the design issues for practically experimenting with a location tracking system - the Sentient City-Bike. We looked at our city-bike application requirements, and discussed some of the key design choices and competing solutions. We favour a hybrid sensor network, with mobile sensor nodes communicating with track-side servers. Low power WiFi with optimised power management serves dual-function in our approach: location tracking and data communication. We also proposed potential social networking applications to engage large participation of users. Although application-driven research requires time and money for prototyping and data collection, we argue that large-scale mobility traces collection is a premise for successful research, design, and building of mobile systems.

Acknowledgements

We want to thank Jonathan Cameron for his effort in developing Linux device drivers for the CSK WiFi board, and George Coulouris for his feedback on the wireless sensor system. We are also grateful for the comments and suggestions received from the anonymous reviewers.

7. REFERENCES

[1] P. Bahl and V. N. Padmanabhan. RADAR: An in-building RF-based user location and tracking system. In *IEEE INFOCOM (2)*, pages 775–784, 2000.

[2] A. R. Beresford and F. Stajano. Location privacy in pervasive computing. *IEEE Pervasive Computing*, 2(1):46–55, 2003.

[3] Y.-C. Cheng, Y. Chawathe, A. LaMarca, and J. Krumm. Accuracy characterization for metropolitan-scale wi-fi localization. In *Proc. of MobiSys '05*, pages 233–245. ACM, 2005.

[4] S. B. Eisenman et al. The bikenet mobile sensing system for cyclist experience mapping. In *SenSys '07*, pages 87–101, New York, NY, USA, 2007. ACM.

[5] J. Froehlich, J. Neumann, and N. Oliver. Measuring the pulse of the city through shared bicycle programs. In *Proc. of UrbanSense08*, November 2008.

[6] B. G. Griswold et al. Using mobile technology to create opportunistic interaction on a university campus. In *Proceedings of UBICOMP Workshop on Supporting Spontaneous Interaction in Ubiquitous Computing Settings*, Munich, Germany, 2002.

[7] A. LaMarca et al. Place lab: Device positioning using radio beacons in the wild. In *Proceedings of PERVASIVE 2005*, Munich, Germany, 2005.

[8] M. Meister et al. Application scenarios for cooperation objects and their social, legal and ethical challenges. In *Proc. of the First International Conference on the Internet of Things*, pages 92–97, 2008.

[9] P. Mitcheson, E. Yeatman, G. Rao, A. Holmes, and T. Green. Energy harvesting from human and machine motion for wireless electronic devices. *Proceedings of IEEE*, 96(9):1457–1486, Sept 2008.

[10] D. Niculescu and B. Nath. Ad hoc positioning system (APS). In *Proc. of IEEE GLOBECOM*, 2001.

[11] B. Raman and C. K. Censor networks: A critique of "sensor networks" from a systems perspective. *ACM SIGCOMM Computer Communication Review*, 38(3):75–78, July 2008.

[12] J. Rekimoto, T. Miyaki, and T. Ishizawa. LifeTag: WiFi-based continuous location logging for life pattern analysis. *Lecture Notes in Computer Science*, 4718:35, 2007.

[13] K. Romer, P. Blum, and L. Meier. *Handbook of sensor networks: algorithms and architectures*, chapter Time synchronization and calibration in wireless sensor networks, pages 199–237. Wiley and Sons, 2005.

[14] C. Savarese, J. Rabaey, and K. Langendoen. Robust positioning algorithms for distributed ad-hoc wireless sensor networks. In *USENIX Technical Annual Conference*, pages 317–327, 2001.