Participatory Bluetooth Sensing: A Method for Acquiring Spatio-Temporal Data about Participant Mobility and Interactions at Large Scale Events

Arkadiusz Stopczynski, Jakob Eg Larsen, Sune Lehmann, Lukasz Dynowski, Marcos Fuentes
Technical University of Denmark
Department of Applied Mathematics and Computer Science, Section for Cognitive Systems
Matematiktorvet, Building 303B, DK-2800 Kgs. Lyngby, Denmark.
{*arks,jel,slj*}*@imm.dtu.dk*

***Abstract*—Acquisition of data to capture human mobility and interactions during large-scale events is a challenging task. In this paper we discuss a mobile sensing method for mapping the mobility of crowds at large scale events using a participatory Bluetooth sensing approach. This non-invasive technique for collecting spatio-temporal data about participant mobility and social interactions uses the capabilities of Bluetooth capable smartphones carried by participants. As a proof-of-concept we present a field study with deployment of the method in a large music festival with 130 000 participants where a small subset of participants installed Bluetooth sensing apps on their personal smartphones. Our software module uses location and Bluetooth scans to utilize smartphones as provisional scanners that are present with higher frequency in regions with high density of participants. We discuss the initial results obtained and outline opportunities and challenges introduced by this methodology along with opportunities for future pervasive systems and applications.**

***Keywords*-data acquisition; participatory sensing; spatio-temporal data; mobility; interactions; large-scale events;**

I. Introduction

Over recent years, mobile phones have become increasingly ubiquitous. Beyond the availability of apps that enable a multitude of uses for smartphones, a smartphone has also become an important research instrument for measuring human mobility, behavior, and interaction patterns. Smartphones have been used to collect sensor data to measure spatial co-occurances and interactions [1] [2] and the embedded wireless technologies (Bluetooth, WiFi, Cellular), allow using mobile phones as a proxy for tracking human mobility [3] [4].

Human mobility at massive events is a particularly interesting and challenging topic for a number of reasons. Firstly, large-scale events take place in areas with a limited infrastructure, both in terms of mobility and communication. Secondly, with a large number of participants and very high density of individuals, the data needed to accurately describe mobility- and social interaction patterns requires higher resolution of the spatio-temporal data, both in terms of geo-spatial accuracy and frequency. During massive events, tracking provides insight into participant behavior in terms of mapping out particular points-of-interests and event preferences, such as artist and genre preferences at a music festival. In addition spatio-temporal data obtained from such events can be used to model the social interactions in this particular context by studying spatial co-occurrences of participants.

Smartphones are equipped with a Bluetooth transceiver designed for short-range ad-hoc wireless communication with headphones and other computing devices. However, Bluetooth enabled devices can also be used to obtain spatio-temporal data regarding presence of individuals, as each Bluetooth device is characterized by a unique identifier (a MAC address). The discovery process involved in inter-connectivity of Bluetooth devices enables detection of multiple Bluetooth devices in proximity. As a mobile phone is carried by a person, the unique Bluetooth identifier acts as a proxy for the presence of the person.

Whereas previous approaches to Bluetooth sensing has typically relied on a number of Bluetooth scanners situated in fixed locations, this paper describes an architecture for the Android mobile operating system that enables participatory Bluetooth sensing for data acquisition. An implementation was deployed in a field trial at a large scale music festival event, allowing insights into the feasibility of acquiring accurate spatio-temporal data using participatory Bluetooth sensing.

II. Related Work

Bluetooth sensing has been used as a way to measure human behavior in several different domains. Those can be roughly categorized into 1) fixed Bluetooth scanning devices tracking individuals on various scales and 2) participant-carried scanners (phones or specialized hardware) used for detecting people proximity in small-scale settings.

Scanners in fixed locations have been used to estimate queue length (waiting time) in airport security areas [5] or to quantify mobility metrics for a rural interstate highway [6], as examples of a small and medium scale Bluetooth scanning. Large scale tracking using Bluetooth scanners has

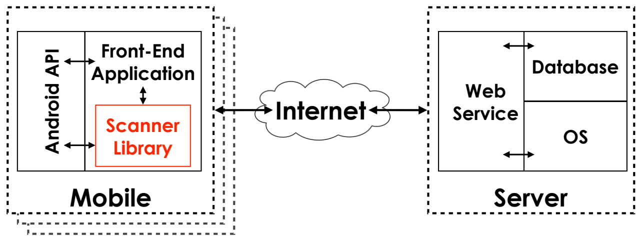

Figure 1. Architecture of the participatory Bluetooth sensing system with a scanner library to be integrated into a mobile application and the server back-end with a database.

been used in Cityware project [7] (city-scale) and in [8] and [9] for festival-scale events.

Mobile approaches to Bluetooth scanning has been employed primarily for detecting face-to-face meetings in a constrained group of people with Bluetooth discovery as a proximity indicator [1] [10]. In other projects the aforementioned method has been used to collect audience data. At the Ghent Street festival located in Belgium, 22 USB Bluetooth transceivers were situated at fixed locations around the city center to discover audience mobility [8].

More recently mobile participatory sensing [11] has been applied in a variety of different domains including public health, transportation [12], urban planning, and environment [13], with the aim to obtain high resolution and quality data. Dutta et al. [14] discuss a system for acquiring data on air quality in urban areas by deploying a handheld monitoring device tethering with a mobile phone to upload GPS annotated air quality data. Whitney and Lipford [15] discuss participatory sensing for community building, exemplified with applications providing bus and parking information. Collective noise mapping of cities is enabled by NoiseTube – a mobile app turning a smartphone into a noise sensor [16].

III. Data Acquisition with Participatory Bluetooth Sensing

We have developed a system that uses coordinated GPS and Bluetooth scans to convert participants smartphones into provisional scanners. Assuming that users of the app are representative of the general participant population, the sampling performed by these sensors is adaptive in the sense that sensors will be present in proportion to the density of individuals, providing high temporal and spatial resolution in more crowded areas. The participatory Bluetooth sensing technique relies on existing technology and infrastructure, that is, the already-deployed smartphones carried by participants, such as at a large scale event. The data acquisition is done by a software component we developed for the Android mobile operating system. The component collects and sends spatio-temporal data to an external repository. The overall architecture of the system is shown in Figure 1.

The software is deployed as a module without a graphical user interface (GUI) and referred to as the scanner library. The intention is that the scanner library can be imported into an Android application and run as a background service with minimum impact on the performance and resources. This is critical, as power resources are likely scarce at large scale events, such as music festivals. Therefore the sampling frequency is adapted to minimize the consumption of energy on the smartphone. A scan performed by the scanner library includes discovery of Bluetooth devices in proximity, update of the current location (GPS), and reading the current battery level using standard Android APIs. All data is stored locally on the phone and uploaded to a server database in periodic batches controlled by a scheduler module. If battery level is low the scheduler will decrease sampling and postpone data upload to preserve battery power.

IV. Field Study

We carried out an initial deployment of the participatory Bluetooth sensing method in order to analyze the applicability of the method with respect to acquire spatio-temporal data on participant mobility and interactions. The field study was carried out at the Roskilde Festival, which is an annual large-scale music festival located south of the city of Roskilde in Denmark. The festival is attended by approximately 130 000 participants and runs over 8 days total, where the first 4 days constitute the so-called warm-up phase with only a few bands playing and the latter 4 days constituting the main festival with more than 150 concerts on six different stages. The festival is situated on a large field outside the city, with very limited fixed infrastructure. Each year electricity, networking, water, and other infrastructure has to be deployed to parts of the area. Similarly, the large camping grounds where the festival participants stay also have limited infrastructure.

The deployment of the participatory Bluetooth sensing involved three different Android applications that were created with the authors specifically for this festival event and made available for download through Google Play. Decibel shown in Figure 2(a) was a social noise measuring app using the embedded smartphone microphone to measure sound levels and share the results with other users. Hide & Seek shown in Figure 2(b) was a social game where participants could create a personal profile with the objective of finding and meeting people at the festival through a set of different clues that were provided through the mobile app. MusicNerd shown in Figure 2(c) allowed the user to browse the current and upcoming concerts and helped the user get to the concert on time by estimating walking distance to the stage based on GPS location and Bluetooth to estimate how crowded the area was. Each application contained the scanner library implementing the data acquisition needed for the participatory Bluetooth sensing. The applications used the data collected by the library to extend the functionality: all apps were location-aware and additionally MusicNerd used the estimation of the crowd density (based on Bluetooth scans) to predict walking time.

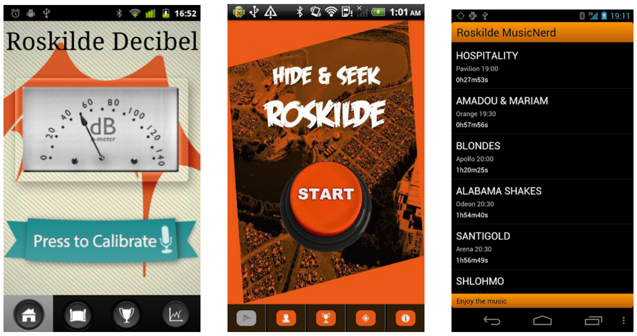

(a) Decibel (b) Hide & Seek (c) MusicNerd

Figure 2. Screenshots from the three Android Apps used as part of the participatory Bluetooth sensing, including (a) a noise measuring app, (b) a social game (Hide & Seek), and (c) a music program scheduling app.

Data acquisition started once the user accepted the conditions presented in the application. To preserve the anonymity of the participants the obtained MAC addresses were one-way hashed when sent to the server. In addition the participants had to actively switch on Bluetooth discoverability and location services on their smartphones in order for the data collection to occur. The applications were mentioned in a sub page in the official festival app, but unfortunately any other promotion of the mobile apps was not possible due to restrictions from the festival organizers. Thus the uptake of the apps depended entirely on participants actively searching for the festival related apps in Google Play, them finding it though the official app, or through word-of-mouth.

V. Results

A total of 155 smartphones participated in the field study and submitted at least one data point (location or Bluetooth scan) from any of the frontend applications. These data came from the festival as well as other places (people installing the frontend apps before entering the festival area or visiting the festival and left early without uninstalling the apps). Before limiting the discoveries to the festival 16 818 Bluetooth data points were collected, coming from 6 389 discoveries, and submitted by 155 scanners. 4 942 unique devices were discovered. The uptake of the applications were limited as the apps were not advertised beyond mentioned in the official festival app for Android.

Limiting the data to the festival is trivial in the temporal domain, as the time can be extracted directly from the Bluetooth timestamps, with the assumption that the mobile phone clocks are synchronized with the mobile phone network. To recover the location of the Bluetooth scans, independent streams of Bluetooth and location data need to synchronized based the timestamps from the smartphone clock. The software collected the data from both sensors independently, the decision was made because a) users can enable and disable those sensors separately b) the location fix can take significantly longer to obtain than a Bluetooth scan or may even fail entirely. Although it is possible to force location updates when the Bluetooth scanning is performed, we believe that for the participatory Bluetooth scanning it should be assumed that the location and Bluetooth are only roughly synchronized and constitute two independent data streams.

Determining the location for a Bluetooth scan in the unsynchronized streams case is a question of deciding upon a threshold: how distant in time the location and Bluetooth scans can be to take the location of the scan as known. Choosing 240 seconds and only accepting location updates with reported accuracy better than $30m$, we were able to assign location to 5 707 Bluetooth scans (out of the initial 16 818). Increasing the threshold only marginally improves this result. Many Bluetooth scans came from devices that never submitted their location, thus they are disregarded in further analysis. Assuming a threshold of 240 seconds, 5 705 unique scans came from 77 scanners and included 2 407 unique discovered devices. This constitutes the festival dataset used for further analysis[1].

A. Spatial and Temporal Coverage

The whole festival area, including concerts space and camping grounds covers roughly $2.23km^2$, where the main concert grounds cover about $0.265km^2$. Assuming $10m$ effective Bluetooth scan radius, over the whole festival 19.2% of the whole area was visited at least once and for the main concert grounds this number is 60.3%.

In order to analyse the data, we bin the whole area into $10m \times 10m$ spatial bins. We then assume that a Bluetooth scan happening in a bin covers the entire bin, and nothing beyond it. Assuming a rectangular shape of the festival grounds, we divide it into 160×160 spatial bins. Not all of these bins are accessible to the general public, so the effective area is smaller than that. We also bin the festival period into 1 296 10 minute temporal bins. 10 minutes is the minimal interval between scans on a single device, set due to battery constraints.

The spatial coverage is visualized in Figure 3. The data is smoothed using kernel density estimation and is based on the number of scans happening in the spatial bins over the course of the whole festival. The temporal coverage of the festival is presented in Figure 4, where each Bluetooth scan is visualized as a red line. To combine spatial and temporal coverage, we plot the histogram of the spatial coverage in Figure 5. We can see that in over half of the 1 296 temporal bins there was at least one spatial bin with a Bluetooth scan happening.

[1]The dataset including location and Bluetooth traces is available in a de-identified form at: https://github.com/h0pbeat/roskilde-2012-bluetooth-sensing-data

Figure 3. Kernel density overlay based on the number of scans performed in $10m \times 10m$ spatial bins over the course of the whole festival. We can clearly see the main festival grounds, with the center at the main stage area.

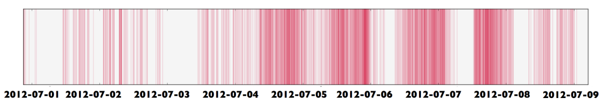

Figure 4. Scans coverage of the festival from Saturday evening until Sunday evening. We can see much increased scans intensity on Wednesday (the day before the opening of the main festival area), Thursday, Friday, and Saturday.

Users of the system demonstrated different mobility patterns: with a different number of total scans, they visited different numbers of spatial bins. This is summarized in Figure 6, where the x-axis contains the number of scans submitted, the y-axis the number of visited spatial bins over the whole festival, and the bubble size indicates the time span of users activity (the time between the first and last datapoint for a given scanner).

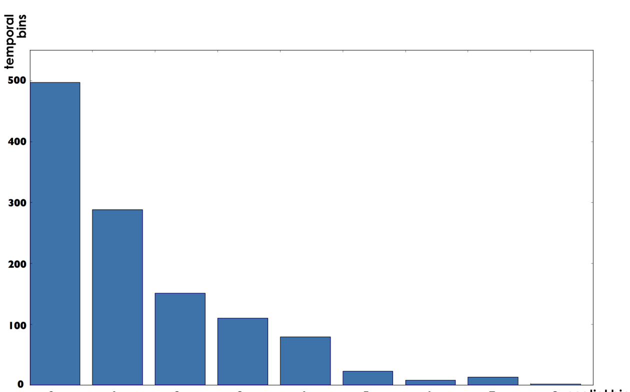

Figure 5. Histogram of the spatial coverage, showing how often a given spatial coverage occurred (number of spatial bins with at least one scan). One spatial bin was covered in almost 300 time bins.

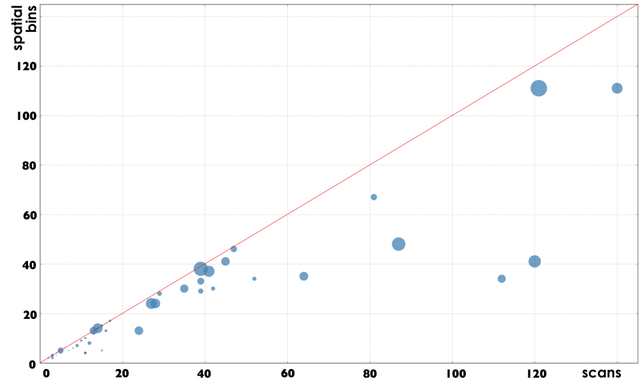

Figure 6. Scans and number of spatial bins covered. On the x-axis number of scans, on the y-axis number of covered spatial bins, on the z-axis (size of the bubble) duration of the users activity. We can see that users present different patterns of mobility: with similar number of scans they visited different number of spatial bins. The high mobility is not an effect of jumpy location inaccuracy, as we have validated this by increasing the size of the spatial bins.

B. *Extrapolated Coverage*

As the uptake of the applications was limited, we consider the proportions of coverage that can be obtained with more participants. To perform the extrapolation we define the following features to describe the users that participated in the field study: the time when the first scan was performed, time span of the activity, number of scans, burstiness parameter [17] (here used to capture the breaks in the scans, due to technical problems with the apps, phones running out of power, and users switching them off), and spatial mobility. Based on data from the participants in the field study, we simulate the spatial coverage with 200 and 2 000 participants taking part in the study as scanners (2 000 corresponds to about 1.5% of the festival population). The spatial bins in which they perform scans are chosen from the smoothed spatial coverage, as shown in Figure 3. Based on the features, we run 100 simulations to obtain the potential spatial coverage as visualized in Figure 7 (200 users) and Figure 8 (2 000 users). We can compare this figure with Figure 5 and notice significantly increased spatial coverage, with the main peak in the 2 000 extrapolation around 35 spatial bins covered.

VI. Discussion

Existing work on human mobility tracking at large scale events have typically relied on deployment of a Bluetooth sensing infrastructure involving a number of scanners situated in fixed locations. This approach has several disadvantages. Namely, costs associated with the deployment and maintenance of the infrastructure (especially in non-urban areas) and that a scanner can only cover a limited area (typically $10-20m$). Additionally, this technique is costly to adapt in constantly changing environments, where any

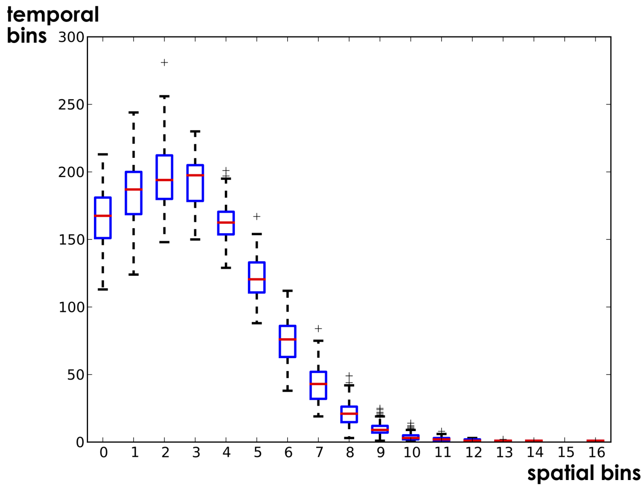

Figure 7. Extrapolated spatial coverage. Histogram of the spatial coverage, showing how often a given spatial coverage happened (number of spatial bins with at least one scan). Whiskers show the lowest datum still within 1.5 IQR of the lower quartile, and the highest datum still within 1.5 IQR of the upper quartile. $N = 200$ users and $K = 100$ simulations.

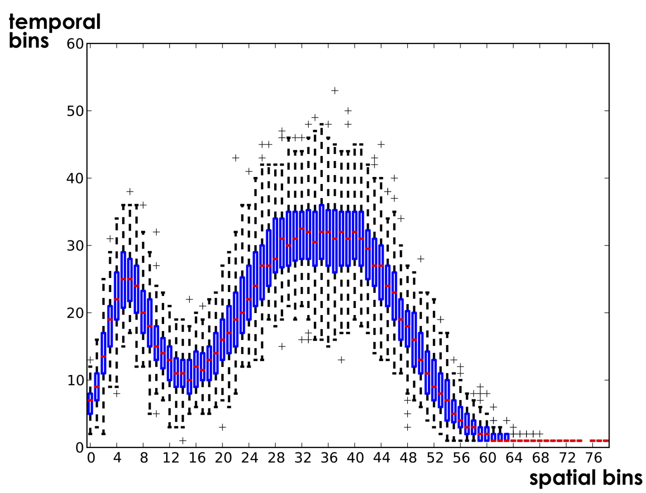

Figure 8. Extrapolated spatial coverage for $N = 2000$ users and $K = 100$ simulations.

change requires a full reinstallation of the equipment. An advantage of fixed Bluetooth sensing is that the deployed scanners can scan continuously (typically twice a minute) leading to a relatively high temporal resolution, but disadvantages include limited scalability and limited spatial coverage.

For participatory Bluetooth sensing the temporal resolution and spatial coverage rely entirely on the number of participants adopting the apps needed to constitute the sensing infrastructure. In our field study the limited uptake meant that we obtained a larger spatial coverage of the festival area compared to a fixed Bluetooth sensing scenario. However, since the sampling rate on the participant devices was low due to smartphone energy consumption concerns, the temporal coverage was lower in a fixed Bluetooth sensing scenario.

Ideally both high temporal resolution and spatial coverage would be maintained. By extrapolating the data that was obtained during the field study we suggest an increase in temporal resolution and spatial coverage could be obtained with more participants. The official festival app was available for both iOS and Android, and provided an interactive guide to the festival including an interactive map, music program, news, events, and artist information. During the festival the Android app was downloaded 6 700 times. Thus had the participatory Bluetooth sensing module been integrated into the official app that would potentially have lead to a substantial temporal resolution and spatial coverage of the festival. The acquired data could be analysed further in terms of studying participant mobility patterns, social interaction by spatial co-occurrences, as well as participant music preferences by studying participant presence in particular temporal and spatial bins aligned with festival events. Such detailed knowledge about participant behaviors would be relevant in terms of festival evaluation and planning.

Our focus has been on the spatio-temporal data acquisition enabled by participatory Bluetooth sensing. However, the infrastructure constituted by the participatory Bluetooth sensing architecture is also interesting from a pervasive applications and services perspective. This particular field study involved three different Android apps that each represented different application domains. The Decibel app was piggybacking noise measurement data along with the spatio-temporal data thereby enabling collective data collection and noise maps to be available for the participants. While the Hide & Seek game was facilitating social interaction, the MusicNerd app used the combination of spatio-temporal and Bluetooth data to guide participants to the next events based on the collective sensing data.

The spatio-temporal data can give the festival organizers unique insights into festival participant behavior for evaluation and planning purposes, by which festival participants will gain indirectly. But the festival participants can also gain immediately from participation through the use of the different types of participatory pervasive apps like the ones mentioned above, which underlines the importance of the unique value created by the apps for the participants in order to facilitate sufficient uptake.

Over the last years smartphones have become much more widespread. Previously the market was dominated by feature phones and Symbian smartphones, which both typically allowed Bluetooth to be configured to discoverable mode indefinitely. More recently Android has become the most widely adopted smartphone platform, but previous to Android OS version 4 Bluetooth could only be configured into discoverable mode for a limited time (300 seconds) meaning that potentially the overall availability of Bluetooth discoverable devices have decreased over the past years.

However, Android is currently the fastest growing mobile platform and the latest version does in fact allow Bluetooth to be configured to discoverable mode indefinitely. In addition, the latest Bluetooth standard (4.0) currently being adopted has the Bluetooth Low Energy (BLE) feature that has a lower power consumption compared to previous Bluetooth versions, thereby making Bluetooth feasible as a personal beacon technology in modern personal devices, including smartphones, tablets, wearable computing devices, and sensors. With these developments we envision Bluetooth being a relevant proxy for sensing people.

VII. Conclusions

We have carried out a proof-of-concept field study that demonstrates how participatory Bluetooth sensing performed at a large-scale music festival event. It has been demonstrated how spatio-temporal data can be acquired in order to get insights into human mobility and social interactions at a large scale event. We have discussed the necessary uptake of the participatory clients (mobile apps) in order to obtain different levels of spatial and temporal coverage of the area and suggest that participatory Bluetooth sensing can be a reliable, scalable, and lower cost alternative to an approach with Bluetooth scanners in fixed locations.

We have discussed how participatory Bluetooth sensing can provide an infrastructure for pervasive applications in large scale settings. Specific pervasive applications used during the field study included, participatory noise data collection, a social game, and a festival program with map guidance. In real-life settings the success of the technique depends on the value of the apps for the participants and thus the uptake. Extrapolation of the data acquired indicate that a more detailed mapping of participant mobility and interaction could have been obtained by integrating participatory Bluetooth sensing into the official festival app.

Acknowledgment

We would like to thank the festival organizers and the participants using our apps during the festival. This work funded in part by the SensibleDTU project (The Villum Foundation).

References

[1] N. Eagle and A. Pentland *Reality mining: sensing complex social systems* Personal and Ubiq. Computing 10(4), 255–268, 2006.

[2] B.S. Jensen, J.E. Larsen, K. Jensen, J. Larsen, and L.K. Hansen *Estimating human predictability from mobile sensor data* In Proc. of the IEEE Int. Workshop on Machine Learning for Signal Processing, 196–201, 2010.

[3] C. Song, Z. Qu, N. Blumm, and A.L. Barabási *Limits of predictability in human mobility* Science 327(5968), 1018–1021, 2010.

[4] M.C. Gonzalez, C.A. Hidalgo, and A.L. Barabási *Understanding individual human mobility patterns* Nature 453(7196), 779–782, 2008.

[5] D. Bullock, R. Haseman, J. Wasson, and R. Spitler *Automated Measurement of Wait Times at Airport Security* Transportation Research Record: Journal of the Transportation Research Board 2177(-1), 60–68, 2010.

[6] R. Haseman, J.S. Wasson, and D.M. Bullock *Real-Time Measurement of Travel Time Delay in Work Zones and Evaluation Metrics Using Bluetooth Probe Tracking* Transportation Research Record: Journal of the Transportation Research Board 2169(-1), 40–53, 2010.

[7] V. Kostakos and E. ONeill *Cityware: Urban computing to bridge online and real-world social networks* Handbook of research on urban informatics: The practice and promise of the real-time city, 195–204, 2008.

[8] M. Versichele, T. Neutens, M. Delafontaine, and N. Van de Weghe *The use of Bluetooth for analysing spatiotemporal dynamics of human movement at mass events: A case study of the Ghent Festivities* Applied Geography 32(2), 208–220, 2012.

[9] J.E. Larsen, P. Sapiezynski, A. Stopczynski, R. Theodorsen, and M. Mørup *Crowds, Bluetooth and Rock'n'Roll: Understanding Music Festival Participant Behavior* XXXIII Sunbelt Social Networks Conference, 2013.

[10] N. Aharony, W. Pan, C. Ip, I. Khayal, and A. Pentland *Social fMRI: Investigating and shaping social mechanisms in the real world* Pervasive and Mobile Computing, 2011.

[11] S. Reddy, V. Samanta, J. Burke, D. Estrin, M. Hansen, and M. Srivastava *MobiSense – mobile network services for coordinated participatory sensing* In Proc. of the Int. Symposium on Autonomous Decentralized Systems, 1–6, 2009.

[12] A. Thiagarajan, J. Biagioni, T. Gerlich, and J. Eriksson *Cooperative transit tracking using smart-phones* In Proc. of the 8th ACM Conf. on Embedded Networked Sensor Systems, 85–98, 2010.

[13] J. Burke, D. Estrin, M. Hansen, A. Parker, N. Ramanathan, S. Reddy, and M.B. Srivastava *Participatory Sensing*. In the Proc. of the Int. Workshop on World-Sensor-Web, 2006.

[14] P. Dutta, P.M. Aoki, N. Kumar, A. Mainwaring, C. Myers, W. Willett, and A. Woodruff *Common Sense: participatory urban sensing using a network of handheld air quality monitors*, In Proc. of the 7th ACM Conf. on Embedded Networked Sensor Systems, 349–350, 2009.

[15] M. Whitney and H.R. Lipford *Participatory sensing for community building* CHI'11 Extended Abstracts on Human Factors in Computing Systems, 1321–1326, 2011.

[16] N. Maisonneuve, M. Stevens, M.E. Niessen, P. Hanappe, and L. Steels *Citizen noise pollution monitoring* In Proc. of the 10th Annual Int. Conf. on Digital Government Research: Social Networks: Making Connections between Citizens, Data and Government, 96–103, 2009.

[17] K-I. Goh and A-L. Barabási *Burstiness and memory in complex systems* EPL (Europhysics Letters) 81(4) 48002, 2008.