Electronic Communications of the EASST
Volume 37 (2011)

Workshops der wissenschaftlichen Konferenz
Kommunikation in Verteilten Systemen 2011
(WowKiVS 2011)

Resource Management in Public Sensing – Challenges and Requirements

Daniel Fischer, Klaus Herrmann

10 pages

ECEASST Home Page: http://www.easst.org/eceasst/ ISSN 1863-2122

Resource Management in Public Sensing – Challenges and Requirements

Daniel Fischer, Klaus Herrmann

University of Stuttgart
Institute of Parallel and Distributed Systems (IPVS)
Universitätstr. 38, D-70569 Stuttgart
firstname.surname@ipvs.uni-stuttgart.de

Abstract: Public sensing is a new research area emerging from the fields of wireless sensor networks and mobile computing. The idea of public sensing is to use the huge number of heterogenous, uncontrollable sensors and system resources which are readily available in the mobile phones of users and in the environment to execute sensing tasks. For example, such sensors can be used to find the shortest route in a traffic navigation application. In order to plan, execute and adapt sensing tasks, applications need to be able to query for the available resources, e.g. density of certain sensors, in a given area. In this work, we investigate the challenges of providing such information and deduce requirements for a resource management system in public sensing. In addition, we propose and discuss concepts that cope with those challenges.

Keywords: Public Sensing; Urban Sensing

1 Introduction

In recent time, we observe an increasing availability of sensors that are embedded in every day devices carried by the general population. Typical examples are smart phones like the iPhone which includes sensors to measure acceleration, brightness, proximity, the magnetic field, as well as GPS, sound, and a video camera. Motivated by this, many researches nowadays advocated the use of such consumer devices with sensing capabilities for large-scale sensing tasks [CEL+08, AAB+07, CHK08, LEM+08]. Typically, this is called urban, participatory, or *public sensing*. Several examples for applications are air/noise pollution monitoring [KBRL09, MSN+09], large-scale discovery and classification of sound events [LPL+09], and discovering free parking lots [LSM10, CGM06]. The incentives for users to provide their resources and participate in such a system can vary. Typically, the services (like finding a free parking lot) would only be provided to those users who also run the sensing software and provide data, leading to a hight participation if services are attractive.

We expect that goals of different public sensing applications are very different and require gathering and processing of sensor data ranging from local to a global scale. Since the public sensing system cannot monitor every possible data at any time, it has to be instructed about what to measure where and in which quality depending on the application requirements. At the same time, the density, type, as well as quality of sensors may be very heterogeneous in time and space.

Therefore, we assume that future non-trivial public sensing applications include some kind of a planning component that configures the public sensing system depending on the currently available sensors as well as other resources such as local computing capabilities or network properties. Hence, planning application tasks efficiently and effectively requires knowledge about the available resources in the public sensing system. We call the entity providing this knowledge the *resource manager*. The resource manager comprises interfaces for making queries, an infrastructure for transferring the raw data to the requester and the protocols for efficiently collecting and relaying the data. In this paper, we focus on the requirements on its structure and protocols, and we sketch how these requirements can be met.

As an example, consider a car navigation applications which provides the fastest route from A to B for its users. This application relies on a measurement of traffic flow speed of individual street segments. However, the actual method to infer flow speed depends on the available sensors. For example, GPS tracking might be used. Another method may involve using cameras, internal car sensors, acceleration sensors or a mixture thereof. Hence, an application needs a view on the available resources to find the most suitable method for its sensing tasks (i.e. inferring traffic flow speed). This view is provided by the resource manager.

Existing architectures for public sensing [PMM07, PRS+06, LLEC09, HBZ+06, CKK+08, CEL+06] do not cope with the problem that applications have to use different plans depending on the available sensors in order to achieve their goal. They simply assume that the required sensors are in the target area and hence, do not require knowledge about the available resources in the public sensing system. Consequently, no work on resource management in public sensing has been done so far.

The ITU reports that by the end of 2009, there were more than 4.6 billion mobile phones in use [Int09]. Hence, if only a fraction of those users participate in public sensing, we have a huge number of sensors available. Therefore, scalability is an important issue to cope with. Furthermore, users will only participate if we avoid draining the battery of their devices and instead assign the role nodes play based on the (possibly very heterogeneous) available spare energy among devices. We also have to adapt to the mobility of participants (and hence, their devices) since we have no control over it. Finally, to allow applications to react to changes in the availability of resource fast, we need to provide resource information as timely as possible and be able to predict future resource availability with a certain accuracy.

In the following section (Section 2), we introduce our system model. Subsequently, in Section 3, we discuss challenges and requirements in providing the discussed resource manager. In Section 4 we propose concepts to solve those challenges and conclude this paper in Section 5 with a discussion of our work.

2 System Model

We assume that public sensing systems will be of massive scale. Possibly millions of heterogeneous mobile nodes (e.g. mobile phones, smartphones, laptops, ...) may provide sensor data readings. Mobile nodes sense different kinds of data at different periods of time, e.g. a user sitting in an office cannot sense traffic, and a temperature sensor may capture different readings depending on whether it is carried outside or inside a pocket.

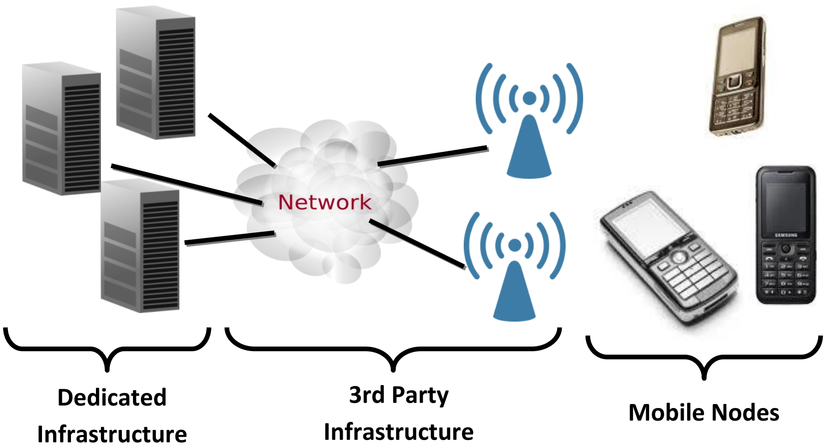

Figure 1: Hybrid network structure of our system model.

An overview of our hybrid network model is shown in Figure 1. We assume that mobile devices support two kinds of communication: First, all mobile nodes have the ability for ad-hoc communication (e.g. WLAN, Bluetooth) with different characteristics (range, quality, bandwidth) between nearby devices forming a MANET. Second, mobile nodes have infrastructure access (e.g. via local access points or based on the cellular infrastructure). The infrastructure is completely wired and connects mobile nodes with dedicated public sensing infrastructure nodes. Note that we assume that most local infrastructure components (e.g. access points, cellular base stations) belong to different administrative domains and we cannot run our algorithms on them or use their functionality. Deploying our own comprehensive dedicated infrastructure to achieve wide-area covering is too expensive.

Furthermore, there are consumer applications, running on any node within a public sensing system. Each of these applications can issue queries to the resource manager in order to plan and adapt their sensing tasks. These queries request views satisfying specific criteria (e.g. density of certain sensors), quality (e.g. maximum age) and area-of-interest (geographical area).

Finally, we assume that mobile nodes have the ability to determine their current position with some level of accuracy. There are numerous techniques for this, including GPS, cell-tower/access point triangulation, or beaconing [SCGL05].

3 Challenges and Requirements

Our grand general vision is to establish a demand-driven view on the situation of arbitrary areas of the world based on the sensors readily available in mobile devices. This raises some tough challenges, which we will discuss in the following. For each such challenge, we list a number of requirements that must be met in order to overcome them:

- **Challenge 1: Scalability** As we assume that there are a huge number of resources in our system, we need to gather and aggregate data without overloading resources such as wireless bandwidth, local resources of mobile devices, or infrastructure nodes in order to maintain scalability. Furthermore, the number of messages per node created due to resource management should be independent of the total number of participants in public sensing. Hence, interactions among nodes should be local and distributed. Also, in order to avoid overloading the infrastructure, it is imperative that devices filter or aggregate data streams early, i.e. before they enter the infrastructure. Furthermore, the resource manager should answer queries using existing (cached) data wherever possible.

- **Challenge 2: Minimize energy usage**: Some sort of application with access to system resources needs to be running on each individual device to enable public sensing. However, in most cases, this will not be the primary application. Instead, sensor and resource data gathering is likely to be a secondary task that must not disrupt normal device usage by draining the battery completely. This is very important in order to get users to participate in public sensing. Hence, it is important that the public sensing systems (and in particular, the resource manager) preserves as much energy as possible and gives the user control over the amount of energy that can be used for public sensing. Moreover, it has to be adaptive to the individual, possibly very heterogeneous energy availability among devices. Since ad-hoc communication (e.g. WiFi, Bluetooth) consumes much less energy per packet compared to using the cellular infrastructure (UMTS/GPRS) [BBV09], the former should be favored and the number of nodes communicating over the latter should be minimized. Furthermore, as we assume that queries specify a certain quality, the resource manager should provide answers only in the requested quality (e.g. accuracy). The specific quality requirements of a query can be used to minimize energy consumption. I.e. if an application only requires a medium accuracy, the system may stop collecting measurements as soon as this accuracy is reached. If a high accuracy is required, it may have to collect more measurements which also consumes more energy.

- **Challenge 3: Fairness**: The approach should be fair, especially in terms of energy usage, i.e. it is not acceptable that individual nodes are completely drained while others still have full capacity. This is not only important to keep participants satisfied, but also for the public sensing system to remain intact. However, we have to consider that individual nodes have very different amounts of energy which can be used for public sensing: Some devices might have only weak batteries (e.g. due to age etc.) or their user may drain much energy for working with the device intensively. Others devices, on the other hand, might have strong batteries and their user may carry them most of the time in his pocket without using them. Hence, the latter should be used to perform energy-intensive tasks in order to drain energy gracefully w.r.t. individual device capabilities in order to keep the public sensing system working and avoid draining individual devices. Furthermore, the user should be able to specify policies concerning the usage of his device, which the public sensing system should respect.

- **Challenge 4: Uncontrollable Mobility**: Since sensing tasks rely on devices carried by the general population, there is little control over their mobility. Hence, the number and density of sensors may be very heterogeneous among different regions. The configuration may change at any time (although it can be predicted on a large scale to some degree). Hence, an important requirement is being adaptive to mobility, density, as well as available resources. For example, in case of continuous queries, the set of nodes in a given area-of-interest may change while the query is running and thus, the set of queried nodes must be adapted to reflect these changes, e.g. in order to preserve the desired level of quality.

- **Challenge 5: Minimize Latency**: Available resources may change quite fast in a public sensing system. In order to allow applications to react to such changes promptly, a resource manager should provide answers to queries as fast as possible. Hence, communication paths should be limited (e.g. number of hops within a MANET) in terms of delay created by them. In order to enable retrieving cached data efficiently, it is sensible to assemble different query-results with the same or similar area of interest at the same or close-by infrastructure nodes. 'Close' does not necessarily mean geographical distance but is related to the latency of transmitting the data to the respective node. More distant infrastructure nodes, on the other hand, should only provide less detailed views to reduce the overall latency.

  Note that we also have to transmit the aggregated query result to the requesting application. Hence, it makes sense to execute the application code that requires the view as input at the node where the view is aggregated, e.g. by employing the concept of mobile code. This avoids the delay created by transmitting the query-result over long distances. Finally, infrastructure-based processing should also include a load-balancing mechanism in order to avoid high delays created due to overloaded nodes. The query load may vary greatly over time and space. For example, areas in a inner city are more interesting for most users than rural areas. Furthermore, at rush hour, we can expect frequent queries for traffic sensing while this is significantly less interesting at night. Hence, infrastructure nodes should be assigned adaptively to geographical areas based on their respective query-load. To maintain scalability (Challenge 1) this should happen in a distributed manner.

- **Challenge 6: Prediction**: In order to reduce delays and inaccuracies for adaptation of public sensing applications, the resource manager should be able to predict future resource states. Existing prediction mechanisms are involved with extrapolating time series of past resource usage on the level of individual devices [FK06]. However, in public sensing we are dealing with large masses of sensor-carrying entities. Furthermore, the prediction of fine-grained mobility of individual users is very hard. For example, approaches that predict the next location based on the history of locations yield only an accuracy of 65-72% [SKJH03]. Predicting the time when the user enters the new location becomes even worse [SDK+06]. Hence, individual user prediction may be only feasible with reasonable accuracy for short time frames. Therefore, we need to integrate prediction models that focus on general mobility patterns of larger number of devices, e.g. utilizing research in the area of sociology [HFMV02, HJT05, MHG+09] that relates to the behaviour of groups rather than individual entities.

4 Meeting the Challenges

Although many current mobile devices can take advantage of cellular infrastructure connectivity (e.g. GPRS, UMTS), communicating over such infrastructures is expensive in terms of energy

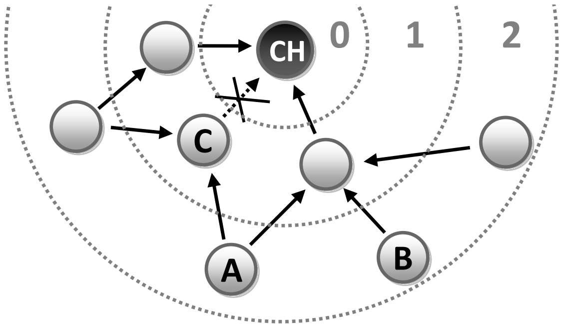

Figure 2: Aggregation in the ring topology.

consumed on the mobile device compared to using ad-hoc communication (WiFi, Bluetooth...) and there is a high per-packet overhead [BBV09]. We have conducted simulation studies with representative energy models [BBV09, XSK+10] that show that this holds even if a multi-hop network is maintained in the ad hoc network. Hence, we propose a cluster-based approach where few special nodes (known as cluster heads) collect the local resource state sensed by near-by nodes using ad-hoc communication before sending the *aggregated* data to a query manager component which runs in the dedicated infrastructure. We plan to use a robust *ring-based topology* aggregation scheme within clusters, originally proposed by Nath et al. [NGSA04] for multi-hop in-network aggregation in static sensor networks. As shown in Figure 2, nodes always aggregate data which they received from the next outer ring with their own data before forwarding the aggregate to the next inner ring. Thus, sensor readings are broadcast by each node within a ring, starting with the outmost ring and continuing towards the center node which is the data sink. As individual sensor readings travel over multiple paths, this aggregation scheme is very robust against collisions or transmission errors. Based on this redundancy, we argue that the approach is also suited for extension into mobile environments (Challenge 4) by building the ring-topology in the course of broadcasting the query. During this broadcast, nodes can also infer their successors in the topology. Hence, nodes know when they have received all data from their successors and when to forward the aggregated data. This drops the assumption of time synchronization made by Nath et al. and thus, makes the query response aggregation much faster (Challenge 5). The general idea is to propose a clustering mechanism that, contrary to existing approaches, exploits the existing query/response communication in order to piggyback information which can be used for cluster maintenance with a very low overhead. This is important for coping with Challenge 1. Furthermore, the ring-based aggregation scheme allows for early aggregation/filtering of data within the local ad-hoc network and thus, strongly reduces the load that reaches the infrastructure. Furthermore, the energy is reduced (Challenge 2) compared to a purely infrastructure-based approach since only few nodes (i.e. cluster heads) use energy-expensive cellular communication and they only send data which was already aggregated/filtered. By selecting aggregation nodes based on the energy available for public sensing, we exploit energy-rich devices and thus drain the energy available for public sensing in an adaptive fashion (Challenge 3). Furthermore, the user can define which fraction of the overall energy is made available to the public sensing system. This prevents the batteries from being drained by public sensing and preserves enough energy for other applications.

For processing queries further within the infrastructure, we propose the following approach: Infrastructure nodes are mapped to a so called latency-space, in which nodes are placed in an n-dimensional space such that the euclidean distance between two nodes approximates their pairwise latency. Several algorithms exist that maintain such a latency space in a distributed manner [PLS+06]. We propose to map a 2-dimensional latency-space of public sensing infrastructure nodes to the geographical area covered by public sensing. For a given query Q with area-of-interest A, Q gets assembled (and its data cached) at one or more nodes in the latency space that overlap with A. Hence, queries with overlapping area-of-interest are always assembled at nodes which are close to each other in terms of latency which allows retrieving cached data fast (Challenge 4). This mapping between the latency-space and the geographical area covered by public sensing should be performed in an adaptive way according to the current query-load: The number of devices assigned to a given area A should be proportional to the query-load of A. Conceptually, during the mapping, we 'stretch' the latency-space in areas with low load, and 'tighten' it in areas of high load.

Another measure for meeting the challenges is to exploit that queries to the resource manager will often have similar aspects as well as an overlapping area-of-interest. In this case, we can decrease the energy consumption (Challenge 2) and the overall load within the system (Challenge 1) by using two simple techniques: First, we cache data at different levels within the resource manager, e.g. at the level of cluster members and cluster heads, as well as within the infrastructure. Hence, queries can be answered by asking only a fraction of mobile devices or even none at all. Second, queries can be deferred and answered in batch-mode (i.e. answer several queries within one interaction) based on their specified maximum age.

In order to make the resource manager predictive (Challenge 6), we mainly draw on the research in crowd behavior [HFMV02, HJT05, MHG+09] to derive estimations of how a group of sensing devices (carried by mobile users) will move. This will provide the basis for extrapolating resource information. Models of the behavior of large groups allow for relatively accurate predictions of aspects like device density in a given area. This is a very promising approach for rendering the resource manager proactive and thus, for reducting query latency.

5 Discussion and Conclusions

Public sensing uses a huge mass of uncontrollable mobile sensor devices in order to perform large-scale sensing tasks. Compared a static, specifically placed static sensor networks, public sensing provides several advantages: First, it is much cheaper since a huge number of mobile sensors are already in widespread use among the general population. Second, mobile sensors can capture places that are out of reach of any static sensor. Third, the great number and density of mobile devices may increase the quality of the sensed data by complementing static sensors. Finally, maintenance and administrative tasks (in particular recharging the battery) are carried out by the users of the devices.

In this paper, we investigated the topic of resource management in public sensing. We pointed out a number of challenges such as scalability, minimizing energy usage, coping with uncontrollable mobility as well as minimizing the latency of query execution and predicting query results.

We have proposed concepts in order to cope with those challenges, both within the mobile ad-hoc network as well as within the infrastructure. In the former, we aggregate data among nearby resources in order to reduce both, the energy usage as well as the load in the infrastructure. In the latter, we propose a latency-space based processing that provides access to cached data with low delay and adapts to different query-loads. In order to predict future resource states, we advocate to use existing models of human group behavior.

The resulting resource management system for public sensing will be a crucial enabling technology that allows for large-scale and highly flexible sensing applications on top of a generic public sensing substrate. This is an important step towards realizing the general vision of public sensing from today's state of the art. We are currently working on a concrete implementation and evaluation of the proposed concepts with respect to the discussed challenges.

Bibliography

[AAB+07] T. Abdelzaher, Y. Anokwa, P. Boda, J. Burke, D. Estrin, L. Guibas, A. Kansal, S. Madden, J. Reich. Mobiscopes for Human Spaces. *IEEE Pervasive Computing* 6(2):20–29, 2007.

[BBV09] N. Balasubramanian, A. Balasubramanian, A. Venkataramani. Energy consumption in mobile phones: a measurement study and implications for network applications. In *IMC '09: Proceedings of the 9th ACM SIGCOMM conference on Internet measurement conference*. Pp. 280–293. ACM, New York, NY, USA, 2009.

[CEL+06] A. T. Campbell, S. B. Eisenman, N. D. Lane, E. Miluzzo, R. A. Peterson. People-centric urban sensing. In *WICON '06: Proceedings of the 2nd annual international workshop on Wireless internet*. P. 18. ACM, New York, NY, USA, 2006.

[CEL+08] A. T. Campbell, S. B. Eisenman, N. D. Lane, E. Miluzzo, R. A. Peterson, H. Lu, X. Zheng, M. Musolesi, K. Fodor, G.-S. Ahn. The Rise of People-Centric Sensing. *IEEE Internet Computing* 12(4):12–21, 2008.

[CGM06] M. Caliskan, D. Graupner, M. Mauve. Decentralized discovery of free parking places. In *VANET '06: Proceedings of the 3rd international workshop on Vehicular ad hoc networks*. Pp. 30–39. ACM, New York, NY, USA, 2006.

[CHK08] D. Cuff, M. Hansen, J. Kang. Urban sensing: out of the woods. *Commun. ACM* 51(3):24–33, 2008.

[CKK+08] C. Cornelius, A. Kapadia, D. Kotz, D. Peebles, M. Shin, N. Triandopoulos. Anonysense: privacy-aware people-centric sensing. In *Proceeding of the 6th international conference on Mobile systems, applications, and services*. Pp. 211–224. ACM, New York, NY, USA, 2008.

[FK06] U. Farooq, W. Khalil. A Generic Mobility Model for Resource Prediction in Mobile Grids. In *Proceedings of the International Symposium on Collaborative Technologies and Systems*. Pp. 189–193. IEEE Computer Society, Washington, DC, USA, 2006.

[HBZ+06] B. Hull, V. Bychkovsky, Y. Zhang, K. Chen, M. Goraczko, A. Miu, E. Shih, H. Balakrishnan, S. Madden. CarTel: a distributed mobile sensor computing system. In *SenSys '06: Proceedings of the 4th international conference on Embedded networked sensor systems*. Pp. 125–138. ACM, New York, NY, USA, 2006.

[HFMV02] D. Helbing, I. Farkas, P. Molnar, T. Vicsek. Simulation of pedestrian crowds in normal and evacuation situations. *Pedestrian and evacuation dynamics*, pp. 21–58, 2002.

[HJT05] D. Helbing, R. Jiang, M. Treiber. Analytical investigation of oscillations in intersecting flows of pedestrian and vehicle traffic. *Phys. Rev. E* 72(4):046130, Oct 2005.

[Int09] International Telecommunication Union (ITU). The World in 2009: ICT facts and figures. 2009.

[KBRL09] E. Kanjo, J. Bacon, D. Roberts, P. Landshoff. MobSens: Making Smart Phones Smarter. *Pervasive Computing, IEEE* 8(4):50–57, Oct.-Dec. 2009.

[LEM+08] N. D. Lane, S. B. Eisenman, M. Musolesi, E. Miluzzo, A. T. Campbell. Urban sensing systems: opportunistic or participatory? In *HotMobile '08: Proceedings of the 9th workshop on Mobile computing systems and applications*. Pp. 11–16. ACM, New York, NY, USA, 2008.

[LLEC09] H. Lu, N. D. Lane, S. B. Eisenman, A. T. Campbell. Bubble-sensing: Binding sensing tasks to the physical world. *Pervasive and Mobile Computing* In Press, Corrected Proof:–, 2009.

[LPL+09] H. Lu, W. Pan, N. D. Lane, T. Choudhury, A. T. Campbell. SoundSense: scalable sound sensing for people-centric applications on mobile phones. In *MobiSys '09: Proceedings of the 7th international conference on Mobile systems, applications, and services*. Pp. 165–178. ACM, New York, NY, USA, 2009.

[LSM10] C. Lochert, B. Scheuermann, M. Mauve. A probabilistic method for cooperative hierarchical aggregation of data in VANETs. *Ad Hoc Networks* In Press, Uncorrected Proof:–, 2010.

[MHG+09] M. Moussa
"ıd, D. Helbing, S. Garnier, A. Johansson, M. Combe, G. Theraulaz. Experimental study of the behavioural mechanisms underlying self-organization in human crowds. *Proceedings of the Royal Society B: Biological Sciences* 276(1668):2755–2762, 2009.

[MSN+09] N. Maisonneuve, M. Stevens, M. E. Niessen, P. Hanappe, L. Steels. Citizen noise pollution monitoring. In *dg.o '09: Proceedings of the 10th Annual International Conference on Digital Government Research*. Pp. 96–103. Digital Government Society of North America, 2009.

[NGSA04] S. Nath, P. B. Gibbons, S. Seshan, Z. R. Anderson. Synopsis diffusion for robust aggregation in sensor networks. In *SenSys '04: Proceedings of the 2nd international conference on Embedded networked sensor systems*. Pp. 250–262. ACM, New York, NY, USA, 2004.
doi:http://doi.acm.org/10.1145/1031495.1031525

[PLS+06] P. Pietzuch, J. Ledlie, J. Shneidman, M. Roussopoulos, M. Welsh, M. Seltzer. Network-Aware Operator Placement for Stream-Processing Systems. In *ICDE '06: Proceedings of the 22nd International Conference on Data Engineering*. P. 49. IEEE Computer Society, Washington, DC, USA, 2006.

[PMM07] *Opportunistic mobile sensor data collection with scar*. 2007.

[PRS+06] A. Parker, S. Reddy, T. Schmid, K. Chang, G. Saurabh, M. Srivastava, M. Hansen, J. Burke, D. Estrin, M. Allman, V. Paxson. Network System Challenges in Selective Sharing and Verification for Personal Social, and urban-Scale Sensing Applications. In *In HotNets*. 2006.

[SCGL05] G. Sun, J. Chen, W. Guo, K. Liu. Signal processing techniques in network-aided positioning: a survey of state-of-the-art positioning designs. *Signal Processing Magazine, IEEE* 22(4):12 – 23, july 2005.

[SDK+06] L. Song, U. Deshpande, U. C. Kozat, D. Kotz, R. Jain. Predictability of WLAN Mobility and Its Effects on Bandwidth Provisioning. In *INFOCOM 2006. 25th IEEE International Conference on Computer Communications. Proceedings*. Pp. 1 –13. april 2006.

[SKJH03] L. Song, D. Kotz, R. Jain, X. He. Evaluating location predictors with extensive Wi-Fi mobility data. *SIGMOBILE Mob. Comput. Commun. Rev.* 7(4):64–65, 2003.

[XSK+10] Y. Xiao, P. Savolainen, A. Karppanen, M. Siekkinen, A. Ylä-Jääski. Practical power modeling of data transmission over 802.11g for wireless applications. In *Proceedings of the 1st International Conference on Energy-Efficient Computing and Networking*. 2010.