

Urban Sensing Systems: Opportunistic or Participatory?[‡]

Nicholas D. Lane,^{*} Shane B. Eisenman,[†] Mirco Musolesi,^{*} Emiliano Miluzzo,^{*} Andrew T. Campbell ^{*}

^{*}Computer Science, Dartmouth College
Hanover, New Hampshire, USA

[†]Electrical Engineering, Columbia University
New York, New York, USA

Abstract

The development of sensing systems for urban deployments is still in its infancy. An interesting unresolved issue is the precise role assumed by people within such systems. This issue has significant implications as to where the complexity and the main challenges in building urban sensing systems will reside. This issue will also impact the scale and diversity of applications that are able to be supported. We contrast two end-points of the spectrum of conscious human involvement, namely participatory sensing, and opportunistic sensing. We develop an evaluation model and argue that opportunistic sensing more easily supports larger scale applications and broader diversity within such applications. In this paper, we provide preliminary analysis which supports this conjecture, and outline techniques we are developing in support of opportunistic sensing systems.

I. INTRODUCTION

Over the past decade, the focus of wireless sensor networking research has evolved from static networks of specialized devices deployed to sense the environment, to networks making use of robotic or other controlled mobility to adapt to the sensing conditions, to a people-centric approach relying on the mobility of people. An expanding research community is developing techniques to bring about very large scale urban sensing by leveraging the increasing sensing capabilities found in consumer devices such as cell phones (e.g., [1] [2] [3]). Data collected from these mobile sensors provide the bases for exciting people-centric applications (e.g., Google StreetView [4]).

A general purpose people-centric sensing system can be thought of in terms of the following stages: application query submission; device selection; sensor sampling; and data analysis, sharing and presentation. In the application query submission stage, a query is submitted to the system specifying at least one required sensor type (e.g., camera) and a set of conditions (i.e., sampling context) under which sampling of the required sensor should take place (e.g., location, time, physical orientation). In the device selection phase, a choice is made among available devices (e.g., via the cellular data channel for mobile phones, or via WiFi access points for WiFi-equipped devices) as to which will be tasked to meet a particular submitted application query. In the sensor sampling stage, the selected device decides when to sample the sensor required by the application query by comparing its own sampling context with that specified in the query. In the data analysis, sharing, and presentation stage, depending on the application requirements, the sensor samples are analyzed (e.g., filtering, classification, etc.) perhaps in combination with results from other queries. The results of the analysis, and possibly the raw samples, are then returned to the querying application and may also be shared with others, depending on issues such as connectivity and privacy.

Including consumer devices as a fundamental building block of the sensing system implies that the human owners of these devices play an important role in the resulting system architecture. In this paper, we consider the question of what roles people, as sensing device custodians, are willing to play in large scale urban sensing systems, particularly to what extent they should be conscious active participants in meeting application requirements. We examine the two end points on the spectrum of custodian awareness and involvement in the architecture, referring to one as *participatory* [3] and the other as *opportunistic* [2].

With participatory sensing the custodian consciously opts to meet an application request out of personal or financial interest. A participatory approach incorporates people into significant decision stages of the sensing system, such as deciding what data is shared and to what extent privacy mechanisms should be allowed to impact data fidelity. Consequently, a participatory system design focuses on tools and mechanisms that assist people to share, publish, search, interpret and verify information collected using custodian devices.

[‡]We raise the issue of the need for evaluation of sensing methodologies in a one page extended abstract included as an unpublished work-in-progress at the “Sensing on Everyday Mobile Phones in Support of Participatory Research” workshop, Nov 2007.

With opportunistic sensing, the custodian may not be aware of active applications. Instead a custodian's device (e.g., cell phone) is utilized whenever its state (e.g., geographic location, body location) matches the requirements of an application. This state is automatically detected; the custodian does not knowingly change the device state for the purpose of meeting the application request. To support symbiosis between the custodian and the system, sensor sampling occurs only if the privacy and transparency needs of the custodian are met. The main privacy concern is the potential leak of personally sensitive information indirectly when providing sensor data (i.e., the custodian's location). To maintain transparency, opportunistic use of a device should not noticeably impact the normal user experience of the custodian as he uses it for his own needs. Thus, the primary challenges in opportunistic sensing are determining when the state of the sensing device matches the requirements of applications, and sampling when application (device state) and custodian (privacy and transparency) requirements are met.

The characteristics of opportunistic and participatory sensing impact the applications that can be practically supported. Participatory sensing places demands on involved device custodians (e.g., prompting via their device GUI) that restrict the pool of willing participants. The tolerance of people to endure interruptions on behalf of applications limits the number (and query load) of concurrent applications that can likely be supported. Further, under the participatory approach, an application needs to have a critical mass of community appeal. These factors combine, we conjecture, to limit both an application's scale and the diversity of applications that are likely to be supported by a purely participatory people-centric network. Applications are best suited to the participatory model when they have a collection of interested custodians whose size is at least as large as number of sensors required to carry out the application. Thus, a strong motivation for opportunistic sensing is to increase the scale and scope/diversity of applications that may otherwise not be supported. Opportunistic sensing shifts the burden of supporting an application from the custodian to the sensing system, automatically determining when devices can be used to meet application requests. In this way, applications can leverage the sensing capabilities of all system users without requiring human intervention to actively and consciously participate in the application, lowering the bar for applications to run in people-centric networks.

In Section II, we propose a simple analytical model to quantitatively evaluate the characteristics of each of these extremes in terms of the probability a particular application query is successfully met. The model parameters capture at an abstract level the application query submission, device selection, and sensor sampling stages mentioned above, as they apply to a particular application scenario - sampling in support of Google StreetView [4]. Section III describes ongoing work in the MetroSense project [2] in support of opportunistic people-centric sensor networks. Related work is discussed in IV.

II. PARTICIPATORY OR OPPORTUNISTIC?: AN EVALUATION MODEL

In this section, we discuss a simple model to evaluate the choice of a participatory model against an opportunistic one. The model is used to explore quantitatively the most important dimensions of the problem.

Assumptions. In the sensing system, data is collected as the result of queries that are submitted to the mobile devices by remote users of the system. For example, a query may be communicated to an application installed by the custodian on his mobile phone via the cellular data channel. Each query specifies that a particular sensor type (e.g., camera) take samples under certain conditions, i.e., the sampling context (e.g., location, time, physical orientation). Meeting the sampling context specified in the query results in capturing data of the required fidelity with respect to the needs of the application.

Notation. The following notation is used in the development of our model: M is the total number of the devices composing the system that can be contacted to potentially meet the requirements of a given query; $P_{\text{accept query}}$ is the probability that a user, if given the option, will allow the query to run on their device (including the likelihoods that a user will accept the possible privacy leaks due to sampling and sharing of sensor data, and the resources that will be consumed by the device possibly without knowing this information in advance); $P_{\text{user adapt}}$ is the probability that a user will adapt their behaviour in support of the sensing activity if prompted (e.g., position their mobile device or modify their intended mobility path to meet the required sampling context, if required); $P_{\text{have sensor}}$ is the probability that the mobile sensing device has the sensor type required by the query; $P_{\text{context match}}$ is the probability that the mobile device has capacity to evaluate whether it matches the sampling context specified in the query, and does in fact match this required context without any user intervention; to shorten probability expressions, we use \bar{P} to mean $1 - P$.

The Model. The probability of success using a participatory strategy is given by

$$P_{success_{part}} = 1 - (1 - P_{accept \ query} \times P_{have \ sensor} \times (P_{context \ match} + \bar{P}_{context \ match} \times P_{user \ adapt}))^M.$$

The term $(P_{context \ match} + (\bar{P}_{context \ match} \times P_{user \ adapt}))$ indicates that either the sensing device has the appropriate context to sample, with respect to the query requirements, or that the human custodian is willing to participate in changing the context of his device to match the requirements of the query.

With the opportunistic strategy, the probability of success is given by

$$P_{success_{opp}} = 1 - (1 - P_{have \ sensor} \times P_{context \ match})^M$$

Combining these expressions and rearranging the terms, we group those terms that represent user interaction and those that reflect inherent properties of the system population. The opportunistic approach provides a higher success probability when the following condition holds:

$$\frac{P_{accept \ query} \times P_{user \ adapt}}{\bar{P}_{accept \ query}} / \frac{P_{context \ match}}{\bar{P}_{context \ match}} < 1. \quad (1)$$

In Equation 1, when equality holds the success probabilities of the participatory and opportunistic approaches are balanced; we call the ratio the *balance ratio* for convenience.

Quantitative Evaluation. Here we present a comparison of the participatory and opportunistic strategies in terms of the model given in Section II. To drive the selection of the model probabilities, we assume a Google StreetView [4] variant as an application submitting requests to the system. Google StreetView provides 360 degrees panoramic street-level views of several metropolitan centers across the United States. One can imagine user cell phones being tasked to opportunistically take street-level photos using the cell phone camera when the phone is out of the pocket. In the context of this application, we can derive the model probabilities as follows. $P_{have \ sensor}$ is the probability that the cell phone has a camera. More than eighty percent of existing cell phones are estimated to have camera phones [5]. We conservatively set $P_{have \ sensor} = 0.8$. In 2006, the average American spent 838 minutes per month on the phone [6]; we use this number to approximate the context matching probability (i.e., the phone is out of the pocket and can thus be used to take a picture). While an intelligent query process (e.g., do not attempt to map when most people are asleep) might be able to boost this probability, we balance this against the chance that the caller is using a hands-free device which would preclude useful photography, and use $P_{context \ match} = \frac{838}{365 \times 24 \times 60} = 0.0191$. Deriving realistic values for the query acceptance and user adaptation probabilities is more challenging given their multi-faceted and subjective nature. We consider statistics from a report on the feasibility of cell phone surveys [7] as related to those we might expect when anonymous queries are routed to a cell phone using the participatory model, with one major caveat. The study¹ reports contact rates for first call attempts, which maps conceptually to the notion of accepting a query, based on which we set $P_{query \ accept} = 0.3$. The same study reports survey completion rates from first call contacts, which given the fact that completing a survey requires some modification of planned behaviour conceptually relates to $P_{user \ adapt}$ and we set $P_{user \ adapt} = 0.2$ accordingly. However, we conjecture the actual inclination of a user to accept sensing queries and adapt their behavior accordingly will be substantially lower than the conceptually related values from [7], due to the much higher volume of queries involved. While the first contact and survey completion rates are based on a single survey instance, members of an urban sensing system need to answer multiple queries over time. If the system is supporting many applications simultaneously, a single user might be queried several times per day, which might quickly become tedious if the user has to manually participate each time. A natural solution to avoid the requirement of excessive manual intervention is the automatic application of a pre-configured user policy. However, it is not clear how policy-based approaches can be applied in practice. In aggregate, the authors of [8] [9] [10] argue that pre-configured policy is possible, citing studies where users were happy to build and use policies against classes of people (e.g., students, co-workers, etc.). On the other hand, studies by Palen and Dourish [11] and Lederer, et al [12] argue that a priori configuration of applications that disclose private information will usually not work given the dynamics of life. Further, it is difficult to imagine how such class-based policies can be defined in an urban sensing system where pseudo-anonymous system users

¹The study methodology [7] provides at least \$5 of compensation for each contact so that potential charges incurred against the participant's account do not bias the result.

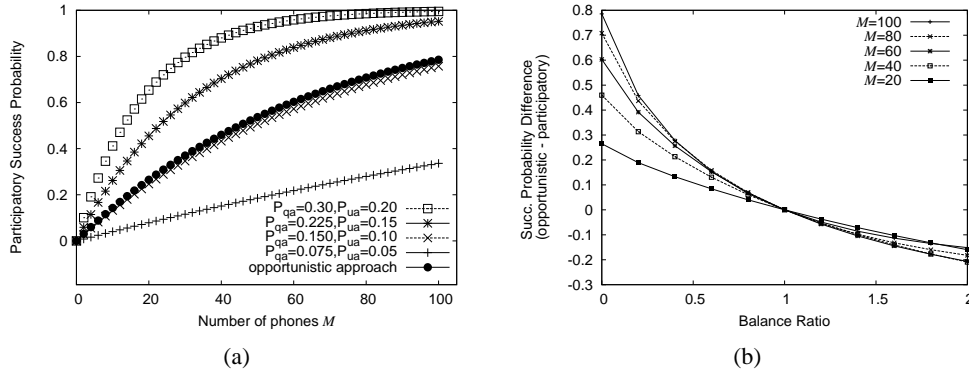


Fig. 1. In (a), the probability of query success when using the participatory strategy is plotted against the number of cell phones for different values of $P_{accept\ query}$ and $P_{user\ adapt}$. The query success probability when using the opportunistic approach is included for comparison. In (b), the difference in query success probability between the opportunistic and participatory approaches is plotted against the balance ratio (Equation 1) for different numbers of cell phones M .

submit queries to run on the mobile devices of others. Ultimately, given the lack of statistics on the matter, we use the values published in [7] but with the view that these represent (perhaps grossly) loose upper bounds.

Regardless of the absolute values, the relationship between the derived values for $P_{query\ accept}$ and $P_{user\ adapt}$ likely translates to urban sensing given that users are more likely to accept the resource and privacy hits implied by a query than to modify their behavior on its account. Therefore, in evaluating the model we keep the ratio of $3/2$ intact while scaling the values. In Figure 1(a), we plot the query success probability versus number of cell phones for both participatory and opportunistic approaches. The participatory curves are labeled with the values of $P_{query\ accept}$ (i.e., P_{qa}) and $P_{user\ adapt}$ (i.e., P_{ua}). Following intuition, with higher levels of human involvement, the participatory approach outperforms the opportunistic approach. However, as the probability of human cooperation falls, as is likely to happen as daily barrages of sensing queries begin to annoy the user, then the opportunistic approach becomes competitive and even outperforms the participatory one. Clearly, a model of how the probability of participation would drop in the face of an increasing query load would be helpful in determining the likely steady-state operating point of the participatory system, in terms of the values for $P_{user\ adapt}$ and $P_{query\ accept}$. We hope to learn such a model through planned large scale human-based experiments in the MetroSense project [13].

While our choices for the values of model parameters are well-reasoned, the numerical guidance given by Figure 1(a) is qualified by the lack of directly applicable statistical data. However, we can analyze the model itself to gain insight into the relative merits and sensitivities of the opportunistic and participatory approaches. In Figure 1(b), for different numbers of cell phones (M) we plot the simple difference between the opportunistic and participatory approaches in terms of query success probability against values of balance ratio (see Equation 1). We fix the $(3/2)$ ratio of $P_{query\ accept}$ to $P_{user\ adapt}$ and evaluate the success probabilities for the values of $P_{user\ adapt}$ and $P_{query\ accept}$ implied by the given values of the balance ratio). By definition, when the balance ratio is < 1 the opportunistic approach is advantageous; when the ratio is > 1 the participatory approach is advantageous. However, from Figure 1(b), moving left from the balance condition (ratio = 1) on a given curve (regardless of M) the opportunistic approach responds much faster to a favorable balance ratio than the participatory approach does when moving right away from the balance condition. We also observe that the comparative performance of the two approaches is most sensitive to M when the balance ratio is < 1 , reflecting the ability of the opportunistic approach to leverage the larger user population. When the ratio > 1 , the results are relatively tightly clustered. Thus, considering the risk and reward in terms of success probability, we observe that: (i) the opportunistic approach does well when the constituent model probabilities are in its favor, and its performance degrades more slowly when the probabilities are not in its favor; and (ii) increasing the number of available sensing devices only amplifies the positive aspect of the opportunistic approach. Therefore, when designing a large-scale urban sensing system, given the uncertainty in what the user population size will be and the uncertainty in user participation probabilities, supporting opportunistic sensing is well motivated.

III. TOWARDS OPPORTUNISM

There are significant challenges that must be overcome for opportunistic sensing to be feasible, including sensing coverage, sensor calibration, determining sampling context, and custodian privacy. Ongoing work in the MetroSense project, described in the following, is focused on developing techniques to overcome these challenges.

Data Collection. The opportunistic use of sensor custodian devices means that sensing is a background operation on the mobile device (e.g., cell phone). Consequently, the periods when the device is able to meet the sensing requirements defined by an application query may be short and intermittent. One can consider this as a sensing coverage problem. To increase the ability of an opportunistic sensing system to collect samples that meet a general set of application requirements (e.g., sensor type, location, physical orientation, time), we are developing an in situ sensor sharing mechanism. This approach allows application queries assigned to a particular device to borrow samples from the best-suited sensors of any device that is locally available at a given time, rather than just from the device to which the query is originally assigned. This ability to adapt in situ relieves the pressure on the sensor selection stage in picking the “perfect” mobile device with respect to a particular application query.

Calibration. While consumer devices are being increasingly augmented with sensors (e.g., accelerometer, camera), the usefulness of these new sensing resources will be limited until suitable calibration techniques are developed. We are developing an approach based on exploiting opportunistic rendezvous between uncalibrated sensors worn by custodians and calibrated infrastructure sensors treated as sources of ground truth. Our preliminary experiments indicate it is possible to estimate the sensor specific calibration curve by formulating the process as an average consensus problem. The algorithm is designed to incrementally converge to the calibration curve over a number of rendezvous by running the calibration algorithm when there are correlations in sensor context between the calibrated and the uncalibrated sensors.

Sensor Context. Sensor context is the meta-data that describes the conditions to which the sensing hardware is exposed that effects either its ability to perform the sensing operation or the sensor data itself. Knowledge of sensor context is required as an input in a number of operations of an opportunistic sensing system. During calibration it is used to determine when two sensors are likely to have sufficiently similar sensor readings. Within in situ sensor sharing, it is used to evaluate potential candidate sensor devices in terms of a given application query. During query execution it indicates when sampling should be initiated and ceased. We are currently working to develop a robust generic sensor context mechanism based on our experiences in building service specific mechanisms (i.e., in support of calibration and in situ sensor sharing).

Privacy. Opportunistic sensing faces barriers to wide scale adoption unless provable privacy guarantees can be provided. The potential exists for the leakage of sensitive personal information both from the data samples, and from the process by which the samples are collected. As an example of the latter, during in situ sensor sharing sensor data is solicited, in part, based upon the similarity of the sensor context requirements of the query and the current sensor context of the device (such as location). Sensor data (e.g., images and sound) may contain information that sensor custodians do not wish to expose. Ongoing work in the MetroSense project [14] provides custodians with anonymity and privacy [15] during the data sharing phase.

Incentives. Incentives play a fundamental role in an architecture using an opportunistic paradigm with resource-limited mobile devices. Existing models can be adapted from the mobile ad hoc networks literature (e.g., [16], [17]), or a trusted third party authority located in the back-end can be periodically accessed via WiFi or GPRS. These mechanisms do not require active user participation. In contrast, for a system based on a participatory paradigm, the user has to be actively involved (making case-by-case decisions) in the incentive-based mechanism. This is an open research area. The design of the system should include the definitions of graphical interfaces displaying the available credit in terms of sampling actions. In fact, psychological elements may also play a key role in the implementation of such mechanisms.

IV. RELATED WORK

While in this paper we model and compare the two extremes on the spectrum of user participation (i.e., opportunistic vs. participatory), in practice most systems represent a middle ground between these two. In the following, we outline a number of recent academic and industry projects exploring urban sensing applications, and indicate their bias towards the opportunistic or participatory extremes.

Both the work by Cartel [18] and Microsoft Research [19] lean towards opportunism. Cartel proposes mobile sensors placed in cars that service adhoc queries, delivered when connectivity is present via open WiFi

access points found in urban areas. Although the required sensors are present, the matching of query and sensor context is uncontrolled and dictates the sampling behavior. Work from [19] is based on ad hoc interactions with cell phones; the focus is on back end concerns such as the correct programmability abstractions, sensor selection and appropriate resource usage (i.e., reduce redundant sampling). The Urban Sensing project at CENS [3] is more closely tied to maintaining humans in the sensing loop and having them be active in the functioning of the system. For instance, DietSense [20] monitors the consumption of food by the participant by taking photographs of the food during meal situations to assist dietitians in monitoring patients. Manual intervention in the review and filtering of the image data is required. Nokia's Sensor Planet [21] focuses on the use of cell phones as mobile sensors. The game/experiment presented in [22] relies heavily on user involvement. The MetroSense Project [2] is focused on building sensing systems leveraging opportunistic whenever possible. Opportunistic sensing, while it raises its own set of challenges (see Section III), can support applications with limited mass appeal, and also a larger number of concurrent applications due to the smaller burden imposed on sensor custodians. We believe anonymous, transparent use of human-carried sensing devices is fundamental to enabling urban-scale sensor networks that provide broad-based application support.

V. CONCLUSION

We have contrasted two potential roles of sensor custodians in urban sensing systems. While we consider participatory and opportunistic solutions to be complementary, we believe that leveraging an opportunistic sensing design approach yields a system that more easily supports large scale deployments and application diversity and should be stressed. Our preliminary analysis supports this position. Urban systems are being built using both techniques and while the analysis shines some light on possible pros and cons the results from building out these systems will provide more conclusive evidence into what are the best architectural choices for urban sensing systems. In the MetroSense project [13], we are developing techniques in support of opportunistic sensing systems.

REFERENCES

- [1] T. Abdelzaher *et al.*, "Mobiscopes for human spaces," *IEEE Perv. Comp. - Mob. and Ubiqu. Sys.*, vol. 6, no. 2, April - June 2007.
- [2] A. T. Campbell *et al.*, "People-centric urban sensing," in *ACM/IEEE WICON 2006. Boston, MA, USA*.
- [3] M. Srivastava *et al.*, "Network system challenges in selective sharing and verification for personal, social, and urban-scale sensing applications," in *HotNets '06*, Irvine, California, USA, November 29 2006, pp. 37–42.
- [4] "Google StreetView," <http://maps.google.com/help/maps/streetview/>.
- [5] C. Noyes, "With new camera phones, you never know who's watching," *Quality Cities*, January/February 2005.
- [6] S. Singh, "Indians talk too much on mobile," http://economictimes.indiatimes.com/Indian_talk-time_461_minmth/articleshow/2118956.cms, June 13 2007.
- [7] J. M. Brick *et al.*, "Cell phone survey feasibility in the u.s.: Sampling and calling cell numbers versus landline numbers," *Public Opinion Quarterly*, pp. 1–17, March 7 2007.
- [8] J. Begole *et al.*, "Work rhythms: Analyzing visualizations of awareness histories of distributed groups," in *Proc. of CSCW '02*, 2002, pp. 334–343.
- [9] R. Hull *et al.*, "Enabling context-aware and privacy-conscious user data sharing," in *Proc. of the IEEE Int'l Conf. on Mobile Data Management (MDM'04)*, vol. 00, 2004, pp. 187–198.
- [10] J. S. Olson *et al.*, "Preferences for privacy sharing: Results and directions," in *CREW Technical Report*, 2004.
- [11] L. Palen and P. Dourish, "Unpacking privacy for a networked world," in *Proc. of CHI'03*, 2003, pp. 129–136.
- [12] S. Lederer *et al.*, "Personal privacy through understanding and action: Five pitfalls for designers," in *Personal and Ubiquitous Computing*, vol. 8, no. 6, November 2004, pp. 440–454.
- [13] "Metrosense project," <http://metrosense.cs.dartmouth.edu>.
- [14] S. B. Eisenman *et al.*, "Metrosense project: People-centric sensing at scale," in *World Sensor Web Workshop, ACM Sensys 2006, Boulder, Colorado*, 2006.
- [15] A. Kapadia *et al.*, "Virtual walls: Protecting digital privacy in pervasive environments," in *Proc. of the 5th Int'l Conference on Pervasive Computing (Pervasive)*, ser. LNCS, vol. 4480. Springer-Verlag, May 2007, pp. 162–179.
- [16] N. B. Salem *et al.*, "A charging and rewarding scheme for packet forwarding in multi-hop cellular networks," in *Proc. of MobiHoc '03*. New York, NY, USA: ACM Press, 2003, pp. 13–24.
- [17] J. Crowcroft *et al.*, "Modelling incentives for collaboration in mobile ad hoc networks," *Performance Evaluation*, vol. 57, no. 4, pp. 427–439, 2004.
- [18] "Cartel," <http://cartel.csail.mit.edu/>.
- [19] "Networked embedded computing group at microsoft research," <http://research.microsoft.com/nec/>.
- [20] S. Reddy *et al.*, "Image browsing, processing, and clustering for participatory sensing: lessons from a dietsense prototype," in *EmNets '07: Proc. of the 4th workshop on Embedded networked sensors*. New York, NY, USA: ACM Press, 2007, pp. 13–17.
- [21] "Sensorplanet," <http://www.sensorplanet.org>.
- [22] V. H. Tuulos, J. Scheible, and H. Nyholm, "Combining web, mobile phones and public displays in large-scale: Manhattan story mashup," in *Pervasive*, 2007, pp. 37–54.