Participatory Privacy: Enabling Privacy in Participatory Sensing

Emiliano De Cristofaro, Palo Alto Research Center (PARC)
Claudio Soriente, ETH Zurich, Switzerland

Abstract

Participatory sensing is an emerging computing paradigm that enables the distributed collection of data by self-selected participants. It allows the increasing number of mobile phone users to share local knowledge acquired by their sensor-equipped devices (e.g., to monitor temperature, pollution level, or consumer pricing information). While research initiatives and prototypes proliferate, their real-world impact is often bounded to comprehensive user participation. If users have no incentive, or feel that their privacy might be endangered, it is likely that they will not participate. In this article, we focus on privacy protection in participatory sensing and introduce a suitable privacy-enhanced infrastructure. First, we provide a set of definitions of privacy requirements for both data producers (i.e., users providing sensed information) and consumers (i.e., applications accessing the data). Then we propose an efficient solution designed for mobile phone users, which incurs very low overhead. Finally, we discuss a number of open problems and possible research directions.

In the last decade, researchers have envisioned an outbreak of wireless sensor networks (WSNs) and predicted the widespread installation of sensor (e.g., in infrastructures, buildings, woods, rivers, or even the atmosphere). This has triggered a lot of interest in many different WSN topics, including identifying and addressing security issues, such as data integrity, node capture, and secure routing. On the contrary, privacy has not really been a concern in WSNs, as sensors are usually owned, operated, and queried by the same entity. (For instance, the National Department of Transportation deploys sensors and collects traffic information related to national highways.)

On the other hand, the proliferation of mobile phones, along with their pervasive connectivity, has propelled the amount of digital data produced and processed every day. This has driven researchers and IT professionals to discuss and develop a novel sensing paradigm, where sensors are not deployed in specific locations, but are carried around by people. Today, many different sensors are already deployed in our mobile phones, and soon all our gadgets (e.g., even our clothes or cars) will embed a multitude of sensors (GPS, digital imagers, accelerometers, etc.). As a result, data collected by sensor-equipped devices becomes of extreme interest to other users and applications. For instance, mobile phones may report (in real time) temperature or noise level; similarly, cars may inform on traffic conditions.

This paradigm is called participatory sensing (PS) — sometimes also referred to as *opportunistic* or *urban* sensing [3]. It combines the ubiquity of personal devices with sensing capabilities typical of WSNs. As the number of mobile phone subscriptions exceeds 5 billion, PS becomes a cutting-edge and effective distributed computing (as well as business) model. We argue that PS appreciably expands the capabilities of WSN applications by, for example, allowing effective monitoring in scenarios where the setup of a WSN is either not economical or infeasible.

However, its success is strongly related to the number of users actually willing to commit personal device resources to sensing applications, and thus to associated privacy concerns. Observe that sensing devices are no longer "dull" gadgets owned by the entity querying them. They are personal devices that follow users at all times, and their reports often expose personal and sensitive information. Consider, for instance, a PS application like http://www.gasbuddy.com/ where gas prices are monitored via user reports, and information announced by participants inevitably exposes their current and past locations, and hence their movements. If users have no incentive to contribute sensed data or feel that their privacy might be violated, they will (most likely) refuse to participate. Thus, not only traditional security but also privacy issues must be taken into account. In this article, we focus on privacy protection in PS. We define privacy in this new context, present a privacy-enhanced PS infrastructure, and elaborate on a number of desirable features that constitute challenging research problems. The proposed privacy-protecting layer can easily be adopted by available PS applications to enforce privacy and enhance user participation.

Participatory Sensing

What Is Participatory Sensing?

PS is an emerging paradigm that focuses on the seamless collection of information from a large number of connected, always on, always carried devices, such as mobile phones. PS leverages the wide proliferation of commodity sensor-equipped devices and the ubiquity of broadband network infrastructure to provide sensing applications where deployment of a WSN infrastructure is not economical or infeasible.

0890-8044/13/$25.00 © 2013 IEEE

PS provides fine-grained monitoring of environmental trends without the need to set up a sensing infrastructure. Our mobile phones *are* the sensing infrastructure, and the number and variety of applications are potentially unlimited. Users can monitor gas prices (http://www.gasbuddy.com/), traffic information (http://www.waze.com/), and available parking spots (http://spotswitch.com/), just to cite a few. We refer readers to [4] for an updated list of papers and projects related to PS.

What Is Not Participatory Sensing?

PS is not a mere evolution of WSNs, where motes are replaced by mobile phones. Sensors are now relatively powerful devices, such as mobile phones, with much greater resources than WSN motes. Their batteries can easily be recharged and production cost constraints are not as tight. They are extremely *mobile*, as they leverage the ambulation of their carriers. Moreover, in traditional WSNs, the network operator is always assumed to manage and own the sensors. On the contrary, this assumption does not fit most PS scenarios, where mobile devices are *tasked* to participate in gathering and sharing local knowledge. Hence, a sensor (or its owner) might choose whether to participate or not. As a result, in PS applications, different entities coexist and might not trust each other.

Participatory Sensing Components

A typical PS infrastructure involves (at least) the following parties:

- **Mobile nodes** are the union of a carrier (i.e., a user) with a sensor installed on a mobile phone or other portable wireless-enabled device. They provide reports and form the basis of any PS application.
- **Queriers** subscribe to information collected in a PS application (e.g., "temperature in Irvine, CA") and obtain corresponding reports.
- **Network operators** manage the network used to collect and deliver sensor measurements (e.g., they maintain GSM and/or third/fourth generation, 3G/4G, networks).
- **Service providers** act as intermediaries between queriers and mobile nodes, in order to deliver reports of interest to queriers.

Queriers can subscribe to the appropriate service provider for one or more types of measurements. For example, assume that Alice subscribes to "available parking spots on W 16th Street, New York," or Bob is interested in the "temperature in Central Park, New York." In turn, mobile nodes share local knowledge — either voluntarily or in return for some profit — with one or more service providers, which make information available to queriers. For example, assume Carol's mobile phone sends report "3 available parking spots on E 56th, New York," while John's device sends "74°F in Central Park, New York."

As mobile nodes and queriers have no direct communication or mutual knowledge, service providers route reports matching specific subscriptions to their original queriers. In fact, mobile nodes ignore which queriers (if any) are interested in their reports. For example, the service provider forwards John's temperature report to Bob; Carol's parking report is not sent to Alice as it refers to a different location.

Privacy Concerns

PS provides an effective solution to a wide range of applications; however, it prompts several security and privacy concerns that need to be carefully addressed.

On one hand, issues such as confidentiality or integrity can be mitigated using state-of-the-art techniques. For instance, all parties can be protected from external eavesdroppers using Secure Socket Layer (SSL)/Transport Layer Security (TLS). The latter provides a secure channel between any two parties, so communications between mobile nodes and service providers or between service providers and queriers are kept confidential.

On the other hand, the need for privacy protection stems from the potential leakage of personal information to *internal adversaries*. Indeed, as the service provider collects all data (i.e., reports and queries), it might learn a considerable amount of sensitive information about both mobile nodes and queriers, and violate the privacy of their movements, interests, habits, and more. For instance, the service provider learns that both Bob and John are located in Central Park, New York. It also learns that Alice is driving on West 16th Street, looking for parking. The continuous collection of information over long periods allows the service provider to meticulously profile users.

Furthermore, as data collected through PS applications becomes available to external entities and organizations (i.e, the queriers), query interests also become sensitive and need to be hidden. For instance, service providers should not learn which interests are "hot."

Finally, there is a tension between privacy and accountability as PS business models may require, at the very least, that reports are available only to entitled (e.g., authorized or paying) members.

However, we claim there is one main reason to protect privacy. If users feel that their privacy is endangered, they will deny sharing their reports. Specifically, it is required that the service provider performs report/query matching but learns no information about query interests. Also, data reports should not reveal to the service provider, the network operator, or unauthorized queriers any information about a mobile node's identity, its location, the type of measurement (e.g., temperature), or the quantitative information (e.g., 74°F).

A Novel Privacy-Enhanced Participatory Sensing Infrastructure

We now present our innovative solution for a Privacy-Enhanced Participatory Sensing Infrastructure (PEPSI). We describe its architecture and privacy desiderata, and overview our instantiation. Finally, we discuss efficiency costs introduced by the privacy-protecting layer.

PEPSI Architecture

PEPSI protects privacy using efficient cryptographic tools. Similar to other cryptographic solutions, it introduces an additional (offline) entity, the registration authority. It sets up system parameters and manages mobile nodes or queriers registration. However, the registration authority is not involved in real-time operations (e.g., query/report matching); nor is it trusted to intervene for protecting participants' privacy.

Figure 1 illustrates the PEPSI architecture. The registration authority can be instantiated by any entity in charge of managing participants registration (e.g., a phone manufacturer). A service provider offers PS applications (used, e.g., to report and access pollution data) and acts as an intermediary between queriers and mobile nodes. Finally, mobile nodes send measurements acquired via their sensors using the network infrastructure, and queriers are users or organizations (e.g., bikers) interested in obtaining reports (e.g., pollution levels).

PEPSI allows the service provider to perform report/query matching while guaranteeing the privacy of both mobile nodes and queriers. It aims at providing (provable) privacy by design, and starts off with defining a clear set of privacy properties.

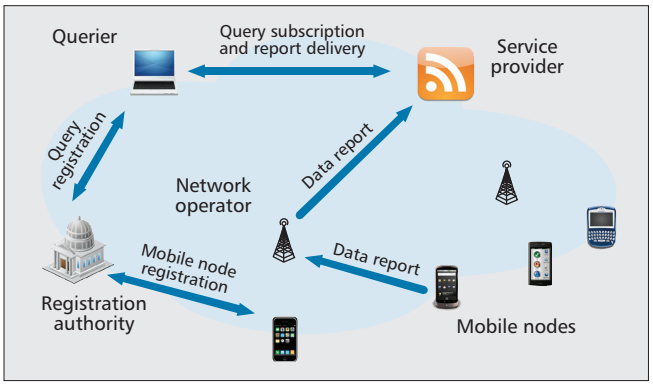

Figure 1. *Privacy-enhanced participatory sensing infrastructure.*

Privacy Desiderata

The *privacy desiderata* of PS applications can be formalized as follows:

- **Soundness:** Upon subscribing to a query, queriers in possession of the appropriate authorization always obtain the desired query results.
- **Node privacy:** Neither the network operator, the service provider, nor any unauthorized querier learn any information about the type of measurement or the data reported by a mobile node. Also, mobile nodes should not learn any information about other nodes' reports. Only queriers in possession of the corresponding authorization obtain reported measurements.
- **Query privacy:** Neither the network operator, the service provider, nor any mobile node or any other querier learn any information about queriers' subscriptions.
- **Report unlinkability:** No entity can successfully link two or more reports as originating from the same mobile node. However, as we discuss below, we do not pursue report unlinkability with respect to the network operator.
- **Location privacy:** No entity can learn the current location of a mobile node (again, excluding the network operator).

In realistic scenarios, it appears unlikely — if not impossible — to guarantee report unlinkability and location privacy with respect to the network operator. In fact, PS strongly relies on the increasing use of broadband 3G/4G connectivity. In these networks, current technology does not allow providing user anonymity with respect to the network operator. Mobile nodes are identified through their International Mobile Subscriber Identity, and any technique for identifier obfuscation would lead to service disruption (e.g., the device would not receive incoming calls). Furthermore, the regular usage of cellular networks (e.g., incoming/outgoing phone calls), as well as heartbeat messages exchanged with the network infrastructure, irremediably reveal a device's location. To provide report unlinkability/location privacy with respect to other parties, we need to trust the network operator (who routes mobile nodes' reports to service providers) not to forward any information identifying the mobile nodes (the identifier, the cell from which the report was originated, etc.).

PEPSI Construction

One of the main goals of PEPSI is to hide reports and queries from unintended parties. Thus, those cannot be transmitted *in the clear*, but must be encrypted. In this section, we discuss how to achieve, at the same time:

- Secure encryption of reports and queries
- Efficient and oblivious matching by the service provider

Due to space limitations and to ease presentation, we only provide an overview of our construction (with no technical details). We refer interested readers to [5] for a complete description of our techniques, as well as formal cryptographic proofs.

A Naive Solution — Traditional confidentiality means are not suited for PS applications. Recall that in our context, mobile nodes and queriers have no mutual knowledge or common history; that is, mobile nodes provide reports oblivious of (any) potential receiver, while queriers subscribe to data reports not knowing who (if anyone) will provide measurements of interest. Hence, we cannot assume that each mobile node shares a unique pairwise secret key with each querier or that reports are encrypted under that key via a symmetric key cipher (e.g., Advanced Encryption Standard, AES). Even if we were to allow interactions between mobile nodes and queriers, we would still need the former to encrypt reports under each key shared with queriers. This would generate a number of ciphertexts quadratic in the number of measurements. Alternatively, we could use a public key encryption scheme and provide mobile nodes with the public keys of queriers. Still, scalability would be an issue as each report would be encrypted under the public key of each querier. In general, because of scalability and loose coupling between data producers and consumers, mobile nodes cannot provide measurements intended for a specific querier, and the latter cannot ask for data from a given mobile node.

Our main building block is Identity-Based Encryption (IBE) — a cryptographic primitive, based on bilinear map pairings, that enables asymmetric encryption using any string ("identity") as a public key. In IBE, anyone can derive public keys from some unique information about the recipient's identity. Private decryption keys are generated by a third party, called the private key generator (PKG). *Our intuition is to use a tagging mechanism on top of IBE.*

Report Encryption — We assume that each report or subscription is identified by a set of labels, or keywords. These are used as "identities" in an IBE scheme. For example, labels "Temperature" and "Central Park, NY" can be used to derive a unique public encryption key, associated to a secret decryption key. Thus, mobile nodes can encrypt sensed data using a report's labels as the (public) encryption key. Queriers should then obtain the private decryption keys corresponding to the labels of interest. Those are obtained, upon query registration, from the registration authority, which, in practice, acts like a PKG.

Efficient Matching using Cryptographic Tags — After enabling encryption/decryption of reports, we need to allow the service provider to efficiently match them against queries. In fact, the application of IBE to PS settings is not trivial: with straightforward use of IBE, oblivious matching of queries and reports would be impossible. In other words, the service provider would forward *all* (encrypted) reports to all queriers; each of them will only be able to decrypt reports of interests (i.e., the ones for which they hold the decryption keys). However, given the large amount of reports produced by mobile nodes, this would incur a considerable overhead for the querier, which must try to decrypt all reports using each of her decryption keys. To address this problem, we propose an efficient tagging mechanism: mobile nodes tag each report with a cryptographic token that identifies the nature of the report only to authorized queriers, but does not leak any information about the report itself. Tags are computed using the same labels used to derive encryption keys. Similarly, queriers compute tags for the labels defining their interests (using the corresponding decryption keys) and provide them to the service provider at query subscription.

Our main contribution, in this context, is to exploit the mathematical properties of bilinear mapping: we ensure that whenever a report matches a query, corresponding tags also match. In other words, a tag computed by John using the encryption key derived from label "temperature in Central Park, New York," is equal to the tag computed by Bob using the decryption key computed over the same label. Specifically, mobile nodes upload reports along with the respective tags, while queriers define their subscriptions uploading the tags they compute at the service provider. The latter can find matches (i.e., a tag related to a report equals the tag related to a subscription) without learning any information about underlying queries/reports.

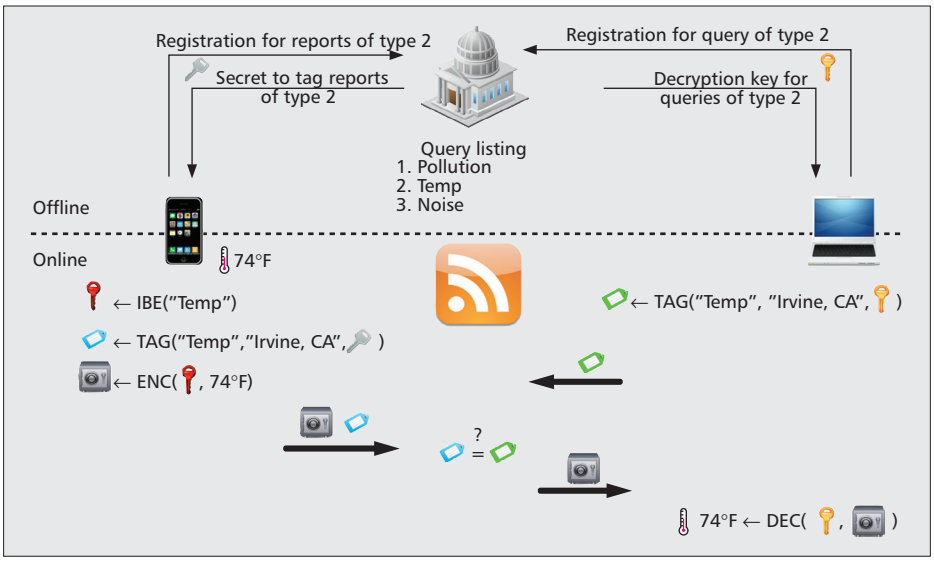

Figure 2. *PEPSI operations.*

PEPSI Operations

Figure 2 shows how PEPSI works. The upper part of the figure depicts the offline operations where the registration authority is involved to register both mobile nodes and queriers.

Querier Registration — In the example, querier $\mathcal{Q}$ (the laptop on the right side) picks "Temp" among the list of available queries and obtains the corresponding decryption key (yellow key).

Mobile Node Registration — Similarly, mobile node $\mathcal{M}$ (the mobile phone on the left side) decides to report the temperature in its location and obtains the corresponding secret used for tagging (grey key).

The bottom part of Fig. 2 shows the online operations where the service provider is involved.

Querier Subscription — $\mathcal{Q}$ subscribes to queries of type "Temp" in "Irvine, CA" using these keywords and the decryption key acquired offline to compute a (green) tag; the algorithm is referred to as ← TAG(). The tag leaks no information about $\mathcal{Q}$'s interest and is uploaded at the service provider.

Data Report — Any time $\mathcal{M}$ wants to report on temperature, it derives the public decryption key (red key) for reports of type "Temp" (via the ← IBE() algorithm) and encrypts the measurement; encrypted data is pictured as a vault. $\mathcal{M}$ also tags the report using the secret key acquired offline and a list of keywords characterizing the report; in the example, $\mathcal{M}$ uses keywords "Temp" and "Irvine, CA." Our tagging mechanism leverages the properties of bilinear maps to make sure that if $\mathcal{M}$ and $\mathcal{Q}$ use the same keywords, they will compute the same tag, despite each of them using a different secret key ($\mathcal{M}$ is using the grey key while $\mathcal{Q}$ is using the yellow one). As before, the tag and the encrypted report leak no information about the nature of the report or the nominal value of the measurement. Both the tag and encrypted data are forwarded to the service provider.

Report Delivery — The service provider only needs to match tags sent by mobile nodes with the ones uploaded by queriers. If the tags match, the corresponding encrypted report is forwarded to the querier. In the example of Fig. 2 the green tag matches the blue one, so the encrypted report (the vault) is forwarded to $\mathcal{Q}$. Finally, $\mathcal{Q}$ can decrypt the report using the decryption key and recover the temperature measurement.

PEPSI Overhead

Resources in PS are not as constrained as in WSNs; nonetheless, overhead incurred at mobile nodes should still be minimized. To foster the adoption of our solution in current PS applications we provide an experimental evaluation of the cost of cryptographic operations used to achieve intended privacy features. We implemented protocol operations executed by mobile nodes on a Nokia N900 (equipped with a 600 MHz ARM processor and 256 Mbytes RAM). Computation overhead, for every report, is due to the computation of the tag and the encryption of the measurement. In our experiments, we experience an average time (over 100 trials) of 93.47 ms to perform these operations.

Communication overhead is merely due to the transmission of the tag, which is the output of a hash function (e.g., SHA-1); thus, it is relative small (160-bit). The encryption of the measurement generates almost no overhead, since, using state-of-the-art symmetric-key ciphers (e.g., AES), the ciphertext's length is almost the same as plaintext's.

Tag computation by queriers is performed only once, during query subscription. Upon reception of measurement of interests, queriers perform symmetric-key decryption, which incurs a negligible overhead.

Finally, note that the service provider incurs no communication or computational overhead; its task is limited to comparing output of hash functions (i.e., tags) and forwarding reports. From a functional point of view, the work of the service provider is no different from that in a non privacy-preserving solution. Thus, privacy protection incurs no overhead at the service provider and enjoys scalability to large-scale scenarios. We conclude that our architecture is practical enough, today, to be deployed for real-world PS applications.

Related Work

Participatory Sensing Projects

In the last few years, PS initiatives have multiplied, ranging from research prototypes to deployed systems. Due to space limitations we briefly review some PS application that apparently expose participant privacy (location, habits, etc.). Each of them can easily be enhanced with our privacy-protecting layer. Interested readers can find a larger list of PS applications at [4]. Quake-Catcher [1] aims at building the world's largest low-cost strong-motion seismic network by utilizing accelerometers embedded in any Internet-connected device.

Kim *et al.* [10] use the power of PS for meaningful places (e.g., home, office) discovery. PS has been shown to be an effective means to monitor levels of air pollution [14], noise pollution [13], and water quality [11]. PS to aid health care providers in patient monitoring has been investigated in [12].

Privacy

Only a little attention has been paid to arising privacy issues in PS [16]. The authors of [2] study privacy in participatory sensing relying on weak assumptions: they attempted to protect *anonymity* of mobile nodes through the use of Mix Networks. (A Mix Network is a statistical-based anonymizing infrastructure that provides k-anonymity; i.e., an adversary cannot tell a user from a set of k.) However, Mix Networks are unsuitable for many PS settings. They do not attain provable privacy guarantees and assume the presence of a ubiquitous WiFi infrastructure used by mobile nodes, whereas PS applications do leverage the increasing use of broadband 3G/4G connectivity. In fact, a ubiquitous presence of open WiFi networks is not realistic today or anticipated in the near future. By contrast, our work aims at identifying a minimal set of realistic assumptions, and defining system properties and clear privacy guarantees to be achieved with provable security.

The work in [15] studies privacy-preserving data aggregation (computation of sum, average, variance, etc.). Similarly, [7] presents a solution for community statistics on time-series data, while protecting anonymity (using data perturbation in a closed community with a known empirical data distribution). Finally, [8] aims at guaranteeing integrity and authenticity of user-generated contents, by employing trusted platform modules (TPMs).

The main technical challenge in providing provable privacy in participatory sensing infrastructure stems from the simultaneous presence of several mutually untrusted (and potentially unknown) entities, including data producers, data consumers, and service providers. A similar scenario arises in the context of publish-subscribe networks [6], which face similar privacy concerns. However, state-of-the-art solutions (e.g., [9]) assume an a priori knowledge (and key exchange) between publishers and subscribers, while PS application requires loose coupling between mobile nodes and queriers. This makes it impossible to apply them to the PS scenario, where data producers and consumers may not know each other. Our solution protects their privacy while requiring no direct interaction between the two parties.

Conclusion and Open Problems

Participatory sensing is a novel computing paradigm that bears great potential. If users are incentivized to contribute personal device resources, a number of novel applications and business models will arise. In this article we discuss the problem of protecting privacy in participatory sensing. We claim that user participation cannot be afforded without protecting the privacy of both data consumers and data producers. We also propose the architecture of a privacy-preserving participatory sensing infrastructure and introduce an efficient cryptographic solution that achieves privacy with provable security. Our solution can be adopted by current participatory sensing applications to enforce privacy and enhance user participation, with little overhead.

This work represents an initial foray into robust privacy guarantees in PS; thus, much remains to be done. Items for future work include (but are not limited to):

- Protecting query privacy with respect to the registration authority. Recall, in fact, that querier Alice needs to obtain the IBE decryption keys from the registration authority, which would then learn Alice's query interests.
- Protecting node privacy with respect to the network operator. Current technology does not allow users' locations and identities to be hidden from the network operator. Hence, it is an interesting challenge to guarantee node anonymity in broadband networks.
- Addressing collusion attacks, where multiple entities might collaborate in order to violate the privacy of mobile nodes or queriers.
- Improving the syntax of supported query types. In fact, PEPSI so far allows query/report matching based on the tags provided by both mobile nodes and queriers. However, PS applications might require more complex queries where queriers are interested in an aggregate of the reports (e.g., average or sum), or even complex query predicates (e.g., comparisons). While simple aggregate function evaluation over encrypted data is viable with available cryptographic techniques (e.g., homomorphic encryption), enabling *efficient* evaluation of complex predicates remains an open challenge.

References

[1] E. S. Cochran *et al.*, "The QuakeCatcher Network: Citizen Science Expanding Seismic Horizons," *Seismological Research Letters*, vol. 80, 2009, pp. 26–30.
[2] C. Cornelius *et al.*, "AnonySense: Privacy-Aware People-Centric Sensing," *6th Int'l. Conf. Mobile Systems, Applications, and Services*, 2008, pp. 211–24.
[3] D Cuff, M. H. Hansen, and J. Kang, "Urban Sensing: Out of the Woods," *Commun.* ACM, vol. 51, no. 3, 2008, pp. 24–33.
[4] E. De Cristofaro and C. Soriente, "Privacy-Preserving Participatory Sensing Infrastructure," http://sprout.ics.uci.edu/PEPSI/index.php?page=projects.php.
[5] E. De Cristofaro and C. Soriente, "Privacy-Enhanced Participatory Sensing Infrastructure," http:// sprout.ics.uci.edu/PEPSI/TR-2011-01.pdf.
[6] P.T. Eugster *et al.*, "The Many Faces of Publish/Subscribe," ACM *Computing Surveys*, vol. 35, no. 2, 2003, pp. 114–31.
[7] R. K. Ganti *et al.*, "PoolView: Stream Privacy for Grassroots Participatory Sensing," *6th Int'l. Conf. Embedded Networked Sensor Systems*, 2008, pp. 281–94.
[8] P. Gilbert *et al.*, "Toward Trustworthy Mobile Sensing," *11 Wksp. Mobile Computing Systems and Applications*, 2010, pp. 31–36.
[9] M. Ion, G. Russello, and B. Crispo, "Supporting Publication and Subscription Confidentiality in Pub/Sub Networks," *6th Int'l. ICST Conf. Security and Privacy in Communication Networks*, 2010, pp. 272–89.
[10] D. H. Kim *et al.*, "Discovering Semantically Meaningful Places from Pervasive RF-Beacons," *11th Int'l. Conf. Ubiquitous Computing*, 2009, pp. 21–30.
[11] S. Kuznetsov and E. Paulos, "Participatory Sensing in Public Spaces: Activating Urban Surfaces with Sensor Probes," ACM *Conf. Designing Interactive Systems*, 2010, pp. 21–30.
[12] B. Longstaff, S. Reddy, and D. Estrin, "Improving Activity Classification for Health Applications on Mobile Devices Using Active and Semi-Supervised Learning," *4th Int'l. Conf. Pervasive Computing Technologies for Healthcare*, 2010, pp. 1–7.
[13] N. Maisonneuve *et al.*, "NoiseTube: Measuring and Mapping Noise Pollution with Mobile Phones," *4th Int'l. ICSC Symp. Information Technologies in Environmental Engineering*, 2009, pp. 215–28.
[14] E. Paulos, R. J. Honicky, and E. Goodman, "Sensing Atmosphere," *Sensing on Everyday Mobile Phones in Support of Participatory Research*, 2007, pp. 1–3.
[15] J. Shi *et al.*, "PriSense: Privacy-Preserving Data Aggregation in People-Centric Urban Sensing Systems," *29th IEEE INFOCOM*, 2010, pp. 758–66.
[16] K. Shilton, "Four Billion Little Brothers?: Privacy, Mobile Phones, and Ubiquitous Data Collection," *Commun.* ACM, vol. 52, no. 11, 2009, pp 48–53.

Biographies

EMILIANO DE CRISTOFARO (edc@parc.com) received his B.Sc. degree in computer science from the University of Salerno, Italy, and his Ph.D. degree in networked systems from the University of California, Irvine. Currently, he is a research scientist in the Security and Privacy group at the Palo Alto Research Center (PARC, a Xerox Company). His research interests include security, privacy, and applied cryptography. More information can be found at http://www.emilianodc.com.

CLAUDIO SORIENTE (claudio.soriente@inf.ethz.ch) received a Ph.D. in networked systems from the University of California, Irvine. He is currently a post-doctoral researcher at ETH Zurich, Switzerland. His research interests include privacy and distributed system security.