

Short Paper: PEPSI: Privacy-Enhanced Participatory Sensing Infrastructure

Emiliano De Cristofaro
Computer Science Department
University of California, Irvine
Irvine, CA, 92617, U.S.A.
edecrist@uci.edu

Claudio Soriente
Distributed Systems Laboratory
Universidad Politécnica de Madrid
28660 Boadilla del Monte (Madrid), Spain
csoriente@fi.upm.es

ABSTRACT

Participatory Sensing combines the ubiquity of mobile phones with the sensing capabilities of Wireless Sensor Networks. It targets the pervasive collection of information, e.g., temperature, traffic conditions, or medical data. Users produce measurements from their mobile devices, thus, a number of privacy concerns – due to the personal information conveyed by reports – may hinder the large-scale deployment of participatory sensing applications. Prior work has attempted to protect privacy in participatory sensing, but it relied on unrealistic assumptions and achieved no provably-secure guarantees. In this paper, we introduce PEPSI: Privacy-Enhanced Participatory Sensing Infrastructure. We explore realistic architectural assumptions and a minimal set of formal requirements aiming at protecting privacy of both data producers and consumers. We also present an instantiation that attains privacy guarantees with provable security at very low additional computational cost and almost no extra communication overhead. Finally, we highlight some problems that call for further research in this developing area.

Categories and Subject Descriptors

C.2.0 [Computer-Communication Networks]: General—Security and Protection; K.4.1 [Computers and Society]: Public Policy Issues—Privacy

General Terms

Algorithms, Security

Keywords

Participatory Sensing, Privacy, Security

1. PRELUDE

Participatory sensing is an emerging paradigm that targets the seamless collection of data from a large number of user-carried devices. By embedding a sensor to a mobile phone, participatory sensing (also called *opportunistic* or *urban* sensing) enables harvesting dynamic information about

environmental trends, such as ambient air quality [21], urban traffic patterns [20], health-related information [15], parking availabilities [19], sound events [16], etc. To allow large-scale deployment, researchers are proposing platforms for application developers [7] and devising innovative business models, based on incentive mechanisms for the capitalization on sensed data [14].

Participatory sensing combines the ubiquity of mobile phones with sensing capabilities typical of Wireless Sensor Networks (WSNs). However, it differs in several aspects. *Sensors* are high-end devices, such as smartphones, with much greater resources than traditional WSN sensors. Their batteries can be easily 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 assumed to own and query all sensors, while this assumption does not apply to most participatory sensing scenarios. Indeed, mobile devices are *tasked* to participate into gathering and sharing local knowledge; thus, different entities co-exist and might not trust each other.

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

- **Sensors:** Installed on smartphones or other wireless-enabled devices, they emit data reports and form the basis of the participatory sensing infrastructure.
- **Carriers:** Usually envisioned as the people carrying their smartphones, they could also be vehicles, animals or any other entity carrying the mobile sensing device. In the rest of the paper, we refer to a sensor and its carrier as a *Mobile Node*.
- **Network Operators:** They manage the network used to collect and deliver reports, e.g., maintaining WiFi, GSM, or 3G network infrastructure.
- **Queriers:** They subscribe to specific information collected in a participatory sensing application (e.g., “*temperature readings in Irvine, CA*”) and obtain corresponding data reports.

Motivation. The number and the heterogeneity of entities involved in participatory sensing prompts a range of new challenges. Unlike in WSNs, sensing devices are not “dull” gadgets, owned by the network operator; they are personal devices that follow users at all time, and their reports often expose personal information. Thus, not only traditional security but also privacy issues must be taken into account, as concerns on personal information disclosure may constitute a fundamental obstacle to large-scale deployment.

Contributions. Prior work has focused on privacy issues in participatory sensing and proposed a few solutions that, however, introduce unrealistic assumptions and provide no *provable* guarantees. On the contrary, we aim at a *cryptographic treatment* of privacy protection in participatory sensing.

We investigate realistic architectural assumptions and a minimal set of formal privacy definitions, intended to protect privacy of both data producers (i.e., mobile nodes) and data consumers (i.e., queriers). Finally, we provide an instantiation that attains privacy guarantees with provable security, at very low additional computational cost and almost no extra communication overhead.

Organization. The next section reviews previous privacy-enhancing solutions and highlights their limitations. Section 3 presents the PEPSI infrastructure and its privacy requirements, while Section 4 yields an efficient instantiation with provable security. Section 5 concludes the paper with a list of open problems.

2. RELATED WORK

In the last years, research interest in participatory sensing has ramped-up. Many researchers have highlighted security and privacy challenges [24], [13], [5], but without proposing actual solutions. Recent proposals in [6] and [12] are—to the best of our knowledge—the only results to address privacy-related problems, hence, they are most related to our work. They aim at protecting anonymity of users, using Mix Network techniques [4], and provide either k -anonymity [25] or l -diversity [18]. They rely on statistical methods to protect privacy and do not achieve provably-secure guarantees. Both proposals only provide limited confidentiality, as reports are encrypted under the public key of a *Report Service* (RS), a trusted party responsible for collecting reports and distributing them to queriers. That is, the RS learns both sensors’ reports and queriers’ interests.

Additional research work focuses on somewhat related problems. [3] argues that privacy issues can be addressed if each user has access to a private server and uses it as a proxy between her *sensors* and the application requesting her data. However, given the number of contributors in a participatory sensing application, the requirement of *per-user* proxies would severely limit the feasibility of this approach. [23] studies privacy-preserving data aggregation, e.g., computation of sum, average, variance, etc. Similarly, [10] 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). Other proposals, such as [9] and [11], aim at guaranteeing integrity and authenticity of user-generated contents by employing Trusted Platform Modules (TPMs).

2.1 Limitations of prior work

We now discuss in detail limitations and open problems of prior work on security and privacy in participatory sensing.

Assuming an ubiquitous WiFi infrastructure. One common feature of existing proposals is the assumption of an ubiquitous WiFi infrastructure used to collect and deliver reports [6, 12, 23]. In particular, [6, 12] use standard MAC-IP address recycling techniques to guarantee user unlinkability between reports with respect to WiFi access points. Such an assumption imposes severe limitations on the scope of par-

ticipatory sensing applications, as an ubiquitous presence of open WiFi networks is not realistic today nor anticipated in the next future.

Actually, the majority of existing participatory sensing applications operate from smartphones and use the cellular network to upload reports [22, 21, 19]. Thus, one cannot use WiFi-based anonymization techniques and, in particular, cannot leverage MAC-IP address recycling to guarantee unlinkability. In cellular networks, devices are identified through their International Mobile Subscriber Identity (IMSI), and ID recycling—besides being impossible with current technologies—would lead to denial of service (e.g., the device would not receive incoming calls for its original ID). Thus, it seems not possible to protect privacy of user locations with respect to the network operator. In fact, the regular usage of cellular networks (e.g., including incoming/outgoing phone calls), as well as heartbeat messages exchanged with the network infrastructure, irremediably reveal device’s location to the operator.

Using Mix Networks. Another limitation of prior work, such as [6], concerns the use of Mix Networks [4] – anonymizing channels used to de-link reports submitted by sensors before they reach the applications. In other words, Mix Networks act as proxies to forward user reports only when some system-defined criteria are met. Several metrics, such as k -anonymity [25] or l -diversity [18] have been defined to characterize privacy through Mix Networks. Observe that a Mix Network may wait to receive k reports before forwarding them to the application, e.g., to guarantee k -anonymity. However, the anonymity level directly depends on the number of reports received and “mixed” by the Mix Network. They rely on statistical methods to protect privacy and do not guarantee provably-secure privacy. Moreover, there could be scenarios where a relatively long time could pass before the desired level of anonymity is reached (when “enough” reports have been collected). As a result, Mix Networks may remarkably decrease system throughput and cannot be used in settings where *timely* reports are required.

Multiple Semi-Trusted Parties. Available techniques to protect privacy in participatory sensing often involve many semi-trusted independent parties, that are always assumed not to collude. The solutions in [6, 12], besides Mobile Nodes, Registration Authority, and WiFi Access Points, also assumes the presence and the non-collusion of a Task Service (used to distribute tasks to users), a Report Service (to receive reports from sensors), and several Mix Network nodes (i.e., a trusted anonymizing infrastructure). It is not clear how to deploy the Task and the Report services as two separate entities having no incentive to collude. Whereas, we aim at minimizing the number of needed semi-trusted parties (and, in general, the number of involved entities), and propose a participatory sensing infrastructure that can be deployed with formal privacy guarantees.

3. PRELIMINARIES

In this section, we formalize: (i) the entities involved in a privacy-enhanced participatory sensing infrastructure, (ii) involved operations, and (iii) privacy requirements.

3.1 Infrastructure

We envision a participatory sensing infrastructure composed by the following entities:

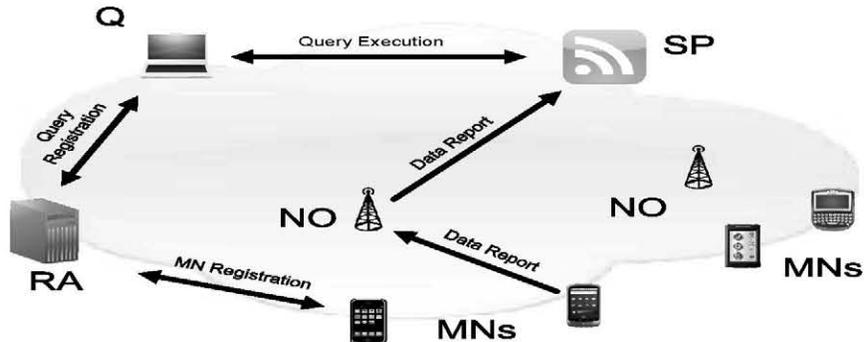


Figure 1: Privacy-Enhanced Participatory Sensing Infrastructure: Mobile Nodes (MNs) register to the Registration Authority (RA) and, subsequently, report sensed data to a Service Provider (SP). Queriers, after registering to RA, subscribe to queries offered by the SP and receive corresponding reports.

Mobile Nodes (MNs). They are computing devices with sensing capabilities (i.e., equipped with one or more *sensors*) and with access to a cellular network. They are carried by people or attached to mobile entities. We assume that MNs run on smartphones and that users voluntarily engage into participatory sensing. We denote with \mathcal{N} a generic mobile node of a participatory sensing application.

Queriers. Queriers are end-users interested in receiving sensor reports in a given participatory sensing application. A generic querier is denoted with Q .

Network Operator (NO). The Network Operator is responsible for the communication infrastructure. We assume that the NO maintains, and provides access to, a cellular network infrastructure (e.g., GSM or 3G).

Registration Authority (RA). The Registration Authority handles the application setup, as well as the registration of participating parties. In our solutions, the RA also contributes to privacy protection, by generating cryptographic public parameters, handling the registration of MNs, and managing queriers’ subscription.

Service Providers (SP). The Service Provider acts as an intermediary between the nodes reporting readings and queriers subscribed to them. We envision one or more SPs running participatory sensing applications that offer different query types. (For example, a national service provider might run a pollution monitoring application and define queries to retrieve reports of pollution levels in different cities). Service provider’s duties may include listing available sensing services, micropayment, data collection, and notification to queriers.

3.2 Operations

We now describe the common operations performed within participatory sensing applications.

Setup. In this phase, the RA generates all public parameters and its own secret key.

MN Registration. Users register their sensor-equipped device to the RA and install participatory sensing software.

Query Registration. Queriers approach the appropriate RA and request an *authorization* to query the participatory sensing application to obtain a specific type of data

reports, e.g., “Pollution level in Madrid, Spain”. (A public list of available sensing services and query syntax may be available from the SP or the RA). Next, they may subscribe to one or more queries, by submitting a request to the SP and awaiting for the responses containing the desired readings. Ideally, only queriers authorized by the RA should receive the desired reports. Also, no information about query interests should be revealed to the SP.

Data Report. MNs report to the SP their readings, using the network access provided by the NO. Ideally, this operation should not reveal to the SP, the NO, or unauthorized queriers any information about reported data, such as type of reading (e.g., pollution) or quantitative information (e.g., $35\text{mg}/\text{m}^3$ carbon oxide). Also, the SP and any querier should not learn the identity of the source MN.

Query Execution. With this operation, the SP matches incoming data reports with query subscriptions. Ideally, this should be done *blindly*, i.e., the SP should learn nothing beyond the occurrence of an (unspecified) match, if any.

In Figure 1, we illustrate our participatory sensing infrastructure. In the depicted scenario, one may envision that a phone manufacturer (e.g., Nokia) acts as the RA and embeds a given type of sensor (e.g., air pollution meter) in one or more of its phone models, operated by smartphone users, i.e., the MNs. A service provider (such as Google, Microsoft, Yahoo, or a non-profit/academic organization) offers participatory sensing applications (used, for instance, to report and access pollution data), and acts as an intermediary between queriers and mobile nodes. Finally, queriers are users or organizations (e.g., bikers) interested in obtaining readings (e.g., pollution levels).

Observe that—similar to related work—we do not address the problem of encouraging mobile phone users to run participatory sensing applications, nor we focus on business incentives for phone manufacturers or for service providers. Nonetheless, it is reasonable to envision that queriers are willing to pay small fees (or receive advertisement) in return to obtaining measurements of interest.

3.3 Privacy Requirements

Before entering the details of our privacy requirements, observe that the main purpose of a participatory sensing ap-

plication is to allow queriers to obtain MNs reports. While our main goal is to protect the privacy of both data producers and consumers, entities registered as a querier should still be able to subscribe to a query and receive desired measurements, thus, techniques to identify legitimate parties before registering them are beyond the scope of our work.

We now define the requirements of a Privacy-Enhanced Participatory Sensing Infrastructure (PEPSI). Our definitions are concise due to space limitation, whereas, adversarial games can be found in paper’s full version [8].

Soundness. We say that PEPSI is *sound* if, upon subscribing to a query, a querier in possession of the appropriate authorization obtains the desired readings (if any).

Node Privacy. We say that PEPSI is *node-private* if neither the NO, the SP, nor any unauthorized querier, learn any information about the type of reading or the data reported by the MN. Also, other mobile nodes should not learn any information about a given node’s reports. In other words, only queriers in possession of the corresponding authorization obtain MN’s readings.

Query Privacy. We say that a PEPSI is *query-private* if neither the NO, the SP, nor any mobile node or any other querier, learn any information about the query subscribed by a querier. We leave as part of future work how to guarantee query privacy also w.r.t. the RA.

Report Unlinkability. Report unlinkability prevents any party from linking two or more reports as originating from a given mobile node. PEPSI provides report unlinkability if different reports originating by the same MN cannot be linked to the source by the RA, the SP, the queriers or any other MN¹.

As discussed earlier, it seems unlikely to provide node location privacy w.r.t. NO, as the NO knows phone’s position at any time. We still consider node location privacy as for the RA, the SP, or unauthorized queriers.

4. PEPSI

We now present our PEPSI instantiation, in accordance with the architectural design in Figure 1 and that complies with privacy definitions of Section 3.3. We also discuss arguments of its provable security and estimate its performance.

4.1 PEPSI Instantiation

In PEPSI, data reports are always labeled using keywords that identify the nature of the information announced by mobile nodes. Similarly, queriers subscribe to given queries by specifying the corresponding keywords. In the rest of the paper, we use the term *identifier*, and the notation ID (or ID^*) to identify the data report/query type. Examples of such identifiers include: “*Temperature in Irvine, CA*” or “*Pollution in Madrid, Spain*”. The list of identifiers – depending on the application – can be obtained either from the SP or the RA. In particular, the RA defines which services (i.e., queries) will be available for mobile nodes to contribute and for users to query. However, as these identifiers can be public, they can be downloaded from the SP or any bulletin board. For ease of presentation, in the rest of the paper, we assume that query identifiers are available at the RA.

¹Observe that we cannot guarantee user anonymity and report unlinkability with respect to the NO, thus, we do not consider the NO in our definition.

4.1.1 High Level Description

PEPSI works as follows. The RA runs the *Setup* algorithm to generate public parameters and secret keys. In order to pose a query, e.g., identified by ID , queriers first need to register to the RA and obtain the corresponding authorization (*Query Authorization*). Then, they subscribe their queries to the SP (*Query Subscription*): in PEPSI, this process reveals nothing about queriers’ interests.

At the same time, before starting reporting data, MNs need to authenticate to the RA, and obtain: (i) the identifier ID corresponding to the type of their reports, and (ii) a token that allows them to announce data (*MN Registration*).

The *on-line* part of PEPSI includes two operations: *Data Report* and *Query Execution*. With the former, MNs upload encrypted reports to the SP. In the latter, the SP *blindly* matches received reports with queries and forwards (matching) reports to all subscribed queriers. Only authorized queriers obtain query responses, can decrypt data reports, and retrieve original measurements. Finally, we let the RA run a *Nonce Renewal* procedure to evict malicious MNs from the participatory sensing application. This procedure is run periodically and the new nonce is securely delivered to honest MNs using broadcast encryption [2].

Our main cryptographic building block is Identity-Based Encryption (IBE), specifically, the construction in [1]. The main advantage in using IBE, as opposed to standard public-key cryptography, is to enable non-interactivity in our query protocol design. This is crucial in participatory sensing scenarios, where MNs and queriers have no direct communication nor mutual knowledge.

4.1.2 Algorithms Specification

Setup. The Registration Authority, given a security parameter λ , generates a prime q , two groups $\mathbb{G}_1, \mathbb{G}_2$ of order q , a bilinear map $e : \mathbb{G}_1 \times \mathbb{G}_1 \rightarrow \mathbb{G}_2$.² Next, a random $s \in \mathbb{Z}_q^*$ and a random generator $P \in \mathbb{G}_1$, are chosen; Q is set such that $Q = P^s$. (P, Q) are public parameters. s is RA’s private master key. Also, a nonce $z \in \mathbb{Z}_q^*$ is selected and R is set such that $R = P^z$. Finally, three cryptographic hash function, $H_1 : \{0, 1\}^* \rightarrow \mathbb{G}_1$, $H_2 : \{0, 1\}^{\mathbb{G}_2} \rightarrow \{0, 1\}^\lambda$, and $H_3 : \{0, 1\}^{\mathbb{G}_2} \rightarrow \{0, 1\}^\lambda$ are chosen.

MN Registration. The MN registration is run between MN and RA. We assume that, after being identified by RA, the Mobile Node \mathcal{N} obtains the pair (z, ID) , where z is the nonce generated by RA during setup, and ID identifies the nature of the readings for which \mathcal{N} provides reports.

Query Registration. The query registration is as follows:

1. **Query Authorization:** Querier \mathcal{Q} registers to RA to get an authorization to receive readings for a given query, identified by ID^* . \mathcal{Q} obtains:

$$sig = H_1(ID^*)^s.$$

2. **Query Subscription:** Querier \mathcal{Q} —subscribing to readings identified by ID^* —sends to the SP:

$$T^* = H_2[e(R, sig)].$$

Data Report. Mobile node \mathcal{N} periodically submits data reports to SP, using NO’s infrastructure. In our protocol,

²Recall that the map e is *bilinear* if $e(U^a, V^b) = e(U, V)^{ab}$.

to upload a reading, \mathcal{D} , related to query ID , \mathcal{N} sends SP (using NO’s infrastructure) the pair $\langle T, CT \rangle$, such that:

$$T = H_2[e(Q, H_1(ID)^z)], \quad CT = ENC_k(\mathcal{D})$$

for $k = H_3[e(Q, H_1(ID)^z)]$.

Query Execution. The query execution is as follows:

1. **Blind Matching:** The Service Provider matches T with stored T^* ’s and returns \mathcal{Q} all matching T^* ’s, alongside associated CT^* ’s.
2. **Notification:** On receiving $\langle T, CT \rangle$, \mathcal{Q} computes $k^* = H_3[e(R, sig)]$ and obtains:

$$\mathcal{D} = Dec_{k^*}(CT)$$

Nonce renewal. We assume a dynamic set of subscribed MNs where new sensors can register and malicious ones are evicted. In order to ban misbehaving sensors, the RA periodically generates and distributes a fresh z to sensors and $R = P^z$ to queriers. The former can be securely distributed to honest sensors using broadcast encryption [2].

4.2 Privacy Analysis

We now consider privacy properties of PEPSI. Due to space limitation, we only provide concise proofs at this stage and defer detailed proofs to [8]. We assume that the system is immune to eavesdropping. In fact, 3G networks encrypt communication between mobile phones and the network operator. Communication between other parties (i.e., $RA \rightleftharpoons MN$, $Querier \rightleftharpoons SP$, etc.) are encrypted using standard techniques, e.g., using SSL.

Soundness. Our PEPSI solution is *sound*, since: for any $\langle ID^*, sig \rangle$ held by a querier \mathcal{Q} , and ID held by a node \mathcal{N} , if: (1) $sig = H_1(ID^*)^s$, where s is RA’s secret key, and (2) $ID^* = ID$, we obtain :

$$\begin{aligned} T &= H_2[e(Q, H_1(ID)^z)] = H_2[e(P^z, H_1(ID^*)^s)] \\ &= H_2[e(R, sig)] = T^* \end{aligned}$$

and, similarly, also $k = k^*$. Therefore, (1) SP correctly matches \mathcal{Q} ’s (authorized) request T^* with the appropriate sensor report $\langle T, CT \rangle$, and (2) \mathcal{Q} can correctly decrypt CT and recover \mathcal{D} .

Node Privacy. Our PEPSI solution is *node-private* since only authorized queriers in possession of a valid sig can learn any information about the report $\langle T, CT \rangle$. Privacy w.r.t. the NO, the SP, other MNs and non-authorized queriers, stems from the security of the underlying Identity-based Encryption scheme [1]. The main intuition is that this kind of adversary could obtain information about a node’s report only if she was in possession of the appropriate sig or—assuming that ID ’s have low entropy—the nonce z . Assuming a CPA-secure and key-private IBE system, the resulting PEPSI scheme is trivially *node-private* w.r.t. the NO, the SP, and non-authorized queriers. We sketch a prove of this claim by contradiction. Assuming our claim is not true, then there exists a polynomial-bounded adversary \mathcal{A} that violates node privacy of PEPSI. \mathcal{A} is given ID and the IBE-encryption of \mathcal{D} under the key ID but not the corresponding $sig = H_1(ID)^s$. If \mathcal{A} decrypts \mathcal{D} with non-negligible probability, then we can construct a polynomial-bounded adversary \mathcal{B} which uses \mathcal{A} to break the CPA-security of IBE. This contradicts our assumption.

Note that the RA could use its secret key to “test” an arbitrary ID^* against an eavesdropped report $\langle T, CT \rangle$. That

is, the RA could learn whether $ID^* = ID$ and violate node privacy. However, assuming that reports $\langle T, CT \rangle$ are super-encrypted under SP’s public key, the RA can access nodes’ reports only if it colludes with the SP.

Query Privacy. PEPSI is *query-private* since neither the NO, the SP, other queriers, nor any mobile node, learn any information about query interests of a querier \mathcal{Q} . Query privacy stems from the security of the underlying Identity-based Encryption scheme [1]. Arguments behind this claim mirror those outlined above for node privacy, thus, we do not repeat them here.

Observe that the privacy of a querier \mathcal{Q} , subscribed to ID^* (i.e., in possession of $sig = H_1(ID^*)^s$), could be violated by a malicious party, subscribed to ID' (i.e., in possession of $sig' = H_1(ID')^s$), only if: (1) she obtains $T = H_2[e(R, sig)]$ sent from \mathcal{Q} to the SP during *Query Subscription* and (2) $ID^* = ID'$. Since the communication between \mathcal{Q} and the SP is encrypted, (1) happens only if such a malicious party colludes with the SP.

Report Unlinkability and Location Privacy. As argued above, it is not possible to guarantee report unlinkability with respect to the network operator. However, one could trust the NO to remove privacy-sensitive metadata from each report (such as mobile nodes’ identifiers, the cell from which the report was originated, etc.), before forwarding it to the SP. Nonetheless, this would not require the NO to act as a Mix Network. Also, the NO never delays message forwarding, e.g., until “enough” reports to protect privacy are collected, but forwards “the payload” of each report (i.e., $\langle T, CT \rangle$) as soon as it is received.

Trust Assumptions. The security of PEPSI only relies on the assumption that the SP is not colluding with either the RA or queriers — on the contrary, prior work assumed the presence of *several* non-colluding and/or fully-trusted parties [6, 12, 9, 11]. Specifically, if the RA and the SP colluded, they could violate node privacy using RA’s secret key, s , and pairs $\langle T, CT \rangle$ received by the SP. Also, recall that any party registered as a querier could potentially collude with the SP and try violating query privacy: it could test a given sig (obtained during the *Query Authorization*) against messages sent by the victim querier to the SP (during *Query Subscription*). We argue that assuming a non-colluding SP is realistic since in participatory sensing, SPs often capitalize on the services they provide, thus, they have no incentive to deviate from an honest-but-curious behavior.

Furthermore, PEPSI needs to trust the NO to remove sensitive MN information from reports before forwarding to the SP. Recall that this assumption is essential since anonymity w.r.t the NO is not achievable in 3G networks.

4.3 Performance Evaluation

Even if resources in participatory sensing are not as constrained as in WSNs, we aim at minimizing the overhead incurred at mobile nodes. This section provides preliminary figures on the cost of cryptographic operations used to achieve intended privacy features.

We implemented protocol operations executed by MNs on a Nokia N900 (equipped with a 600 MHz ARM processor and 256 MB RAM) running the `libpbc` cryptographic library [17]. We selected Type-A pairings and 160-bit prime q . Computation overhead is due to the computation of T , the encryption key k , and the encrypted report CT . Note

that the first two values can be computed off-line, independently of the sensed data. Communication overhead is merely due to the transmission of T , which is the output of a hash function (e.g., SHA-1), and can be as small as 160-bit. Indeed, using available symmetric-key cryptosystems (e.g., AES), the length of CT is almost the same as a reading D .

Without leveraging off-line precomputation, we measured the time to compute and transmit $\langle T, CT \rangle$, using integers as data reports. Over 100 experiments, we experienced an average time of 93.47ms to compute $\langle T, CT \rangle$ and around 80ms for transmission over the 3G network. Note that a naïve (non-private) solution would save in computation (since data would not be encrypted) but would spend roughly the same transmission time to send the report. Finally, remark that the SP incurs no communication nor computational overhead: its task is limited to hash comparison and forwarding. Similarly, the only additional operation that queriers perform during query execution is the symmetric decryption of received readings, which incurs a negligible overhead.

5. CONCLUSION

The participatory sensing paradigm bears an irrefutably great potential. However, its success depends on the number of users willing to report measurements from their mobile devices. Clearly, a wide-scale user participation is bound to effective protocols that preserve privacy of both data producers (i.e., mobile nodes) and data consumers (i.e., queriers). In this paper, we highlighted shortcomings of previous solutions and we embarked towards a cryptographic treatment of privacy in participatory sensing. To this aim, we analyzed which are the privacy features that can be guaranteed with provable security and introduced a participatory sensing protocol that attains them. Finally, we provided figures of the incurred overhead at mobile nodes.

As often happens, deploying actual solutions based on our proposal requires addressing additional (potential) security issues, such as authentication, data integrity, DoS prevention, active attacks, Sybil attacks, etc. Our next step is to deploy testing applications using the PEPSI infrastructure, as well as to devise a large-scale evaluation of its global overhead. Our future work also includes extending the protocols to efficiently support query privacy w.r.t. the RA (i.e., queriers can register without the RA learning their interests). Interesting open challenges remain in how to provide location privacy with respect to cellular network operator, addressing potential collusion between different parties, and supporting more complex queries (e.g., aggregate and conjunctive queries).

Acknowledgments. This research has been partially funded by US Intelligence Advanced Research Projects Activity (IARPA) under grant FA8750-09-2-0071, the Madrid Regional Council – CAM under project CLOUDS (S2009TIC-1692), the Spanish Research Agency – MICINN under project CloudStorm (TIN2010-19077), and the European Commission under projects MASSIF (FP7-257475) and STREAM (FP7-216181). We are also grateful to Nokia for the devices used in our experiments.

6. REFERENCES

- [1] D. Boneh and M. Franklin. Identity-based encryption from the Weil pairing. In *Crypto*, 2001.
- [2] D. Boneh, C. Gentry, and B. Waters. Collusion resistant broadcast encryption with short ciphertexts and private keys. In *Crypto*, 2005.
- [3] R. Cáceres, L. P. Cox, H. Lim, A. Shakimov, and A. Varshavsky. Virtual individual servers as privacy-preserving proxies for mobile devices. In *MobiHeld Workshop*, 2009.
- [4] D. L. Chaum. Untraceable electronic mail, return addresses, and digital pseudonyms. *Communications of ACM*, 24(2), 1981.
- [5] D. Christin, M. Hollick, and M. Manulis. Security and Privacy Objectives for Sensing Applications in Wireless Community Networks. In *IEEE ICCCN*, 2010.
- [6] C. Cornelius, A. Kapadia, D. Kotz, D. Peebles, M. Shin, and N. Triandopoulos. AnonySense: Privacy-aware people-centric sensing. In *Mobisys*, 2008.
- [7] T. Das, P. Mohan, V. Padmanabhan, R. Ramjee, and A. Sharma. PRISM: Platform for Remote Sensing using Smartphones. In *Mobisys*, 2010.
- [8] E. De Cristofaro and C. Soriente. PEPSI: Privacy-Enhanced Participatory Sensing Infrastructure. <http://sprout.ics.uci.edu/PEPSI>, 2011.
- [9] A. Dua, N. Bulusu, W. Feng, and W. Hu. Towards trustworthy participatory sensing. In *HotSec*, 2009.
- [10] R. Ganti, N. Pham, Y. Tsai, and T. Abdelzaher. PoolView: stream privacy for grassroots participatory sensing. In *Sensys*, 2008.
- [11] P. Gilbert, L. Cox, J. Jung, and D. Wetherall. Toward trustworthy mobile sensing. In *HotMobile*, 2010.
- [12] K. Huang, S. Kanhere, and W. Hu. Preserving privacy in participatory sensing systems. *Computer Communications*, 33(11), 2010.
- [13] A. Kapadia, D. Kotz, and N. Triandopoulos. Opportunistic sensing: Security challenges for the new paradigm. In *COMNETS*, 2009.
- [14] J. Lee and B. Hoh. Sell Your Experiences: A Market Mechanism based Incentive for Participatory Sensing. In *PerCom*, 2010.
- [15] B. Longstaff, S. Reddy, and D. Estrin. Improving activity classification for health applications on mobile devices using active and semi-supervised learning. In *PervasiveHealth*, 2010.
- [16] H. Lu, W. Pan, N. Lane, T. Choudhury, and A. Campbell. SoundSense: scalable sound sensing for people-centric applications on mobile phones. In *Mobisys*, 2009.
- [17] B. Lynn. PBC: The Pairing-Based Cryptography Library. <http://crypto.stanford.edu/pbc/>, Last Accessed, 2011.
- [18] A. Machanavajjhala, D. Kifer, J. Gehrke, and M. Venkatasubramanian. l-diversity: Privacy beyond k-anonymity. *ACM Transactions on Knowledge Discovery from Data (TKDD)*, 1(1), 2007.
- [19] S. Mathur, T. Jin, N. Kasturirangan, J. Chandrasekaran, W. Xue, M. Gruteser, and W. Trappe. ParkNet: drive-by sensing of road-side parking statistics. In *MobiSys*, 2010.
- [20] P. Mohan, V. Padmanabhan, and R. Ramjee. Rich monitoring of road and traffic conditions using mobile smartphones. In *Sensys*, 2008.
- [21] E. Paulos, R. Honicky, and E. Goodman. Sensing atmosphere. In *Sensys Workshops*, 2007.
- [22] S. Reddy, K. Shilton, G. Denisov, C. Cenizal, D. Estrin, and M. Srivastava. Bikitastic: sensing and mapping for better biking. In *CHI*, 2010.
- [23] J. Shi, R. Zhang, Y. Liu, and Y. Zhang. PriSense: Privacy-Preserving Data Aggregation in People-Centric Urban Sensing Systems. In *Infocom*, 2010.
- [24] K. Shilton. Four billion little brothers?: Privacy, mobile phones, and ubiquitous data collection. *Communications of the ACM*, 52(11), 2009.
- [25] L. Sweeney. k-Anonymity: A model for Protecting Privacy. *Int. J. Uncertain. Fuzziness Knowl.-Based Syst.*, 2002.