

# Opportunistic Sensing: Security Challenges for the New Paradigm\*

(Invited Paper)

Apu Kapadia<sup>†</sup>  
MIT Lincoln Laboratory  
Lexington, MA USA

David Kotz  
Institute for Security, Technology, and Society  
Dartmouth College  
Hanover, NH USA

Nikos Triandopoulos<sup>‡</sup>  
Boston University  
Boston, MA USA

**Abstract**—We study the security challenges that arise in *opportunistic people-centric sensing*, a new sensing paradigm leveraging humans as part of the sensing infrastructure. Most prior sensor-network research has focused on collecting and processing environmental data using a static topology and an application-aware infrastructure, whereas opportunistic sensing involves collecting, storing, processing and fusing large volumes of data related to everyday human activities. This highly dynamic and mobile setting, where humans are the central focus, presents new challenges for information security, because data originates from sensors carried by people— not tiny sensors thrown in the forest or attached to animals. In this paper we aim to instigate discussion of this critical issue, because opportunistic people-centric sensing will never succeed without adequate provisions for security and privacy. To that end, we outline several important challenges and suggest general solutions that hold promise in this new sensing paradigm.

## I. INTRODUCTION

Sensor networks provide tremendous potential for information collection and processing in a variety of application domains. The first generation of sensor-network scenarios included stationary devices sensing ephemeral features of the environment around them. In this paper we focus on a new kind of sensing, aimed at applications in daily life, employing the mobile devices people already carry. *Opportunistic people-centric sensing* has been introduced as a term to describe this new paradigm: small computational devices, carried by individuals in their daily activities, sensing information directly or indirectly related to human activity, as well as aspects of the environment around them [1].

This new brand of sensing induces a different set of assumptions and trade-offs than in much of the prior work on sensor networks, requiring new thought about the communications infrastructure. Likewise, these new capabilities and architectures pose different challenges and therefore require new solutions for information security. First, applications will probably deal with highly personal information, requiring a deeper attention to privacy and anonymity than in most prior work. Second,

modified assumptions about device and network capabilities—including high mobility, opportunistic networking, strong but not continuous connectivity, and relatively plentiful power—lead to new opportunities for sensor-network architecture and push different security solutions. For example, previous work has focused on security solutions for resource-constrained devices [2], [3], secure routing techniques for static sensor networks [4]–[6], and secure data collection and aggregation in static and fixed tree topologies [7], [8]. Other work has focused on providing anonymity in location-based applications [9], but we feel the trade-offs made do not allow for the rich applications we envision; for example, there is no notion of dynamic, anonymous tasking, thus limiting the breadth of applications supported by a large deployment of sensors.

Given the increasing interest in people-centric sensing applications, the time is ripe to explore new security and privacy challenges. In this setting, sensor devices will (i) be highly mobile, so existing security solutions assuming fixed topologies cannot be employed; (ii) have limited resources, but not severely so: nodes may be in frequent contact with the Internet via Wi-Fi access points or the cellular infrastructure, and strong cryptographic techniques can be employed (e.g., we may assume that public-key cryptography will not significantly degrade a sensor node’s lifetime [10]); and, critically, (iii) most sensors will be carried by people, thus introducing privacy concerns as well as new adversarial threat models.

This paper contributes to the discussion by surveying these security and privacy challenges, and identifying some key solution concepts. We hope this work will instigate further research on this critical topic, and encourage application and system designers to embed security and privacy features prior to deploying systems. Security and privacy has been identified as an important new design challenge in some urban-sensing projects (examples: [11]–[13]) but only in passing.

In Section II, we overview opportunistic people-centric sensing and identify the key features and assumptions driving a need for new solutions. In Section IV, we discuss the applications enabled by this form of sensor network. In Section V, we describe what we believe to be the new challenges and discuss some conceptual solutions. We summarize the paper in Section VI.

\*Authors are listed alphabetically. Many thanks to Peter C. Johnson (Dartmouth College) for his contributions to earlier versions of this paper.

<sup>†</sup>This research was performed while the author was at Dartmouth College.

<sup>‡</sup>Center for Reliable Information Systems & Cyber Security and Department of Computer Science, Boston University. This research was primarily performed while the author was at Dartmouth College.

## II. OPPORTUNISTIC PEOPLE-CENTRIC SENSING

Wireless sensor networks provide a structure for gathering data on scales and in environments previously unattainable. Forests and battlegrounds, oft-cited settings for sensor networks, are not the only interestingly “sensible” locales: new research is focusing on urban areas, sensing aspects of the environment, infrastructure, and society. In an urban setting, one could leverage millions of personal mobile phones, and a near-pervasive wireless-network infrastructure, to collect sensor data on a grand scale without the need to deploy thousands of static sensors. Thus, many researchers propose the *opportunistic-sensing* model, in which people volunteer their mobile devices to transparently collect sensor data as they go about their daily life. Nodes adopt opportunistic practices for sensing and networking, allowing their sensors to be remotely tasked on someone else’s behalf, collecting and reporting sensor data on a best-effort basis when the conditions permit. (In a variant, *participatory sensing*, the participating user is directly involved in the sensing action, e.g., to photograph certain locations or events [14]; although we focus on opportunistic sensing in this paper, we briefly summarize some challenges for this variant in Section V-D.) For example, an application measuring average temperature in a specific room could emit a task requesting temperature readings every five minutes from sensor nodes whose location coordinates indicate they are within the room.

In the opportunistic-sensing model, sensor nodes are carried by people, either directly or in their vehicles, and therefore are conceptually tied to specific individuals. We envision sensors built into devices such as cell phones, which are more powerful than traditional “mote” platforms and can be conveniently recharged on a daily basis, thus mitigating the extreme compute and energy constraints. The sensors are inherently mobile and the sensor data is necessarily “people-centric”; that is, sensing not only the surrounding environment, but also aspects of the person and his or her social setting: where people are and where they are going, what they are doing and what they are seeing. While these new aspects conjure up images of an amazing array of applications, they also present significant security challenges.

Perhaps the most obvious concern is the security of the sensed data itself, especially in correlation to the owner of the device. For example, from CarTel data it is trivial to deduce the location of a participant’s home within a small radius. How can this information be protected? Additionally, is it possible to allow selected people to see personal data at high fidelity, but to present a less accurate view to less trusted viewers?

A second concern is the integrity of the sensed data. An application based on sensor data is useless if data consumers cannot trust the accuracy or timeliness of the data. Is it possible to operate an open, cooperative network of human-carried sensor devices when some of the people, or some elements of the infrastructure, cannot be trusted to communicate sensor data accurately and quickly?

A third concern is the privacy of the user with respect to his

or her participation in the system. Can a participant trust the systems not to track their location, which sensing tasks they execute, or which reports they submit?

On the other hand, many of the security solutions already addressed in the sensor-network literature do not apply to opportunistic, people-centric, urban-sensing scenarios. In “traditional” sensor networks, such as a stationary collection of mote-class devices connected by an ad hoc wireless network, the nodes are configured, deployed, and operated by the same party that desires the sensor data, the nodes are not sensing human behavior, and the main concerns are about an adversary that wishes to intercept or tamper with the data in transit, disrupt the routing of packets in the network, re-task the nodes inappropriately, or prevent the legitimate retasking of nodes. Substantial research attention has been paid, for example, to secure key management and distribution schemes [15]–[17], or to secure retasking solutions such as Deluge [18] and Sluice [19]. The network infrastructure used by most opportunistic-sensing systems, however, is not owned or operated by any one party, and usually not by the same party that wishes to collect the sensor data nor by the individuals who own and operate the mobile devices used as sensor nodes. Neither party will necessarily trust the network infrastructure; the data consumers can not trust the sensor nodes or the people carrying them; conversely, those people will not necessarily trust the system that collects the data or the applications that use the data. The trust models are far more complex than those considered in most sensor-network literature.

Furthermore, an urban environment, the pervasive coverage of Wi-Fi and cellular networks means that multi-hop routing seems less necessary. [Through their experiments, the CarTel project [20] revealed the existence of more than enough open wireless access points to allow the uploading of data to an aggregation point on the Internet.] Although the multi-hop structure of early sensor networks inspired the development of secure routing protocols, in urban-sensing scenarios the focus should be on securing the data rather than securing the route.

Thus, the opportunistic-sensing model induces many challenges: inciting users to volunteer their device for data collection, leveraging the uncontrolled mobility of the sensors, managing a system that scales to millions of sensor nodes, protecting the privacy of the users who volunteer their devices, protecting the integrity of the data from human adversaries, and operating in a distributed-trust environment resulting from the involvement of multiple resource providers and administrating organizations. Other than our own AnonySense system [21], we are not aware of any proposed people-centric opportunistic sensing systems that incorporate security at the foundation of their design. System designs ignoring security and privacy, or attempting to achieve such through security by obscurity, are inadequate and are unlikely to succeed.

In the later sections we detail the challenges related to security and privacy. First, we describe example systems and example applications, and identify security and privacy issues therein.

### III. EXAMPLES: URBAN SENSING PROJECTS

Opportunistic, people-centric urban sensing has already been a topic of study in the community. Here are some current research projects working in this area; we often use the first two as points of reference in our exposition.

*CarTel:* The CarTel project [20] uses small Linux-based computers installed in vehicles to map traffic patterns in Boston and Seattle. The computer contains sensors of its own (e.g., GPS) and attaches to the vehicle's on-board diagnostic (OBD) interface to measure location, direction, speed, acceleration, and other parameters. CarTel's delay-tolerant networking stack, CafNet, takes advantage of randomly-encountered unrestricted wireless access points to upload sensed data to a relational database on the Internet. A web-based user interface provides access to the database, and allows for modification of data-gathering rules, such as time granularity. Applications enabled by the CarTel project include traffic monitoring, automotive diagnostics and notification, and road-surface diagnostics. Note that these applications are people-centric: the producers of the data are ordinary people (in their cars), and the consumers of this data are also ordinary users, not specialized scientists studying the environment.

*MetroSense:* Similarly, MetroSense [1] envisions a future in which the Internet is dominated by data from sensors carried by people during their daily life. The core ideas in MetroSense include opportunistic tasking of personal mobile devices, opportunistic sharing of sensors between neighboring nodes, and the use of smart phones as a large-scale platform for opportunistic people-centric sensing. Prototype applications include BikeNet [22], in which bicycles outfitted with multiple sensor nodes communicate with each other and with neighboring bikes as well as the network infrastructure, and CenceMe [23], in which mobile phones collect data about a user's activity and share it with buddies via a social network.

*Other urban-sensing projects:* Intel Research has collaborated with UC Berkeley to develop a stable of applications involving urban sensing, dubbed Urban Atmospheres [24]. These applications range from air quality testing to place-based ringtones. Additional projects related to urban sensing include Mobiscopes [11], Urbanet [12], CENS Urban Sensing [13], and SenseWeb [25].

### IV. EXAMPLES: APPLICATIONS

We give some examples of opportunistic-sensing applications to motivate that paradigm and to motivate the accompanying security challenges.

*Urban data collection and processing:* We anticipate that opportunistic sensing will encourage large-scale, on-line data collection and processing of context information related to aspects of everyday life, such as locating lost objects [26], or measuring the flow of bicycles in an urban center [27].

For instance, an Active Map [28, for example] represents the location and context of people on a geographical map: students at a university can view the location of their friends on campus, and possibly also information collected by personal and embedded sensors (indicating, for example, that they are in a

conversation or are asleep). Recent examples include commercial projects such as Dodgeball.com and BoostMobile.com, which provide location-based "friend finder" services where users are notified of friends in the vicinity, or can view the locations of their friends on a map. Other systems, such as CenceMe [23], allow sharing of activity information via social-network services. The Mobile Media Metadata project [29] attempts to leverage the location context of cameras to tag photographs with metadata describing their contents. Active Maps could also display historical information, such as the most frequently used running trails, indicating, e.g., which may be muddy or contain steep uphill sections. BikeNet [22] is one such project, aimed at cyclists who would like to share real-time sensor data. The map can also be an interface for registering context-sensitive triggers, such as "remind me to buy milk when I am next driving past the grocery store", "alert me when both Bob and Alice are back in the office," or "let me know when a tennis court is available at the park."

*Environmental monitoring at the human level:* With human-centric sensing, applications that monitor the human environment (such as building maintenance and HVAC) now may become more efficient and useful. It becomes possible to sense and collect information (such as indoor air temperature) at the human level. Thus, the environment is sensed exactly where its semantics are important: at the person, rather than on the wall. One could optimize energy usage, for example, by reducing effort to heat, cool, or ventilate classrooms when the classroom appears empty, or to adjust office temperature as needed for the person sitting at the desk in the corner.

Many have proposed efforts to collect data about traffic [30], pot holes [31], or noise pollution [32]. The Personal Environmental Impact Report provides a pollution-exposure report based on the a cell phone's location trace and public data from (stationary) air-quality sensors [33]. Eventually, if people carry devices that include environmental sensors, environmental monitoring can be achieved with far less static infrastructure. With the advent of mobile sensors, personal devices can help monitor levels of hazardous gases, generating alerts or collecting long-term trends. (The U.S. Department of Homeland Security is considering this very strategy in its Cell-All project [34].)

### V. SECURITY CHALLENGES

These people-centric sensing applications entail serious security and privacy risks. Unrestricted dissemination of users' sensor data results in breaches of *privacy*; users will want to control who may access information about themselves. Also, since data originates from sensors that are under the control of other people, the *integrity* of the data comes into question. For example, a user may tamper with a sensor device to cause it to report false data, or misrepresent the location or time the data was sensed. Furthermore, the *availability* of the infrastructure is critical for these applications to remain useful. In this section we outline specific challenges regarding privacy, integrity, and availability.

## A. Confidentiality and Privacy Issues

The confidentiality of sensed data goes far beyond the provision of a secure channel from the sensor node to some gateway node. Such encryption, and in particular key distribution, has already been well-discussed in sensor-network security literature. We focus on the privacy challenges associated with the collection and dissemination of sensor data.

*Challenge 1: Context privacy:* While several systems have been proposed to address specific types of sensor data (e.g., location privacy in pervasive environments [35], [36] and privacy of medical data from body sensors [37]), usable mechanisms to protect the privacy of more general types of context have been lacking. It is cumbersome for users to specify fine-grained policies, and users are not particularly good at it [38]. It will be harder to do when multiple types of context are involved—who should know whether they are awake, who should know whether they are in a conversation, who should get access to heart-rate information, and so on. In a people-centric sensing environment, context privacy is of high importance. To this end, we have proposed *virtual walls* [39] as a usable metaphor for controlling access to users’ context information.

Virtual walls intuitively map to real walls in the physical world. For example, *transparent* virtual walls allow the release of personal context information much as a person’s physical actions are fully visible through a real transparent wall. We also define intuitive semantics for *translucent* and *opaque* virtual walls. Still, for a real-world deployment, several challenges remain: users may want more precise access control for some types of context information, and hybrid mechanisms are needed that can balance ease of use and a larger degree of control. Another intriguing approach is based on a different metaphor: users pre-define a small set of disclosure policies, thinking of each one as a different public “face” they might wear [40]. In different situations, the user dons a different face.

There are several instances where groups of users should take ownership over data—for example, a team of athletes may want to control access to their sensor data as a group. Negotiation of group policies in a seamless and usable way will be a challenge. Furthermore, it may be desirable to have the negotiation preserve the privacy of users within the group, who may want to vote on policies without being implicated or singled out. We anticipate that schemes based on *secure multiparty computation (SMC)* may be helpful [41], [42]; in SMC, functions can be evaluated while maintaining the privacy of the individual inputs. We also recommend exploring the use of trusted hardware for performance reasons. For example, a trusted server can receive all the users inputs and compute the outcome within the aegis of the tamper-resistant trusted hardware [43], [44].

Consolvo et al. [45] found that when it comes to sharing location information, users were interested in knowing *why* the information was being demanded and in responding with the appropriate level of information to satisfy those requests. If these findings do indeed apply to other kinds of context, how

can solutions such as Virtual Walls incorporate this “why?” component without sacrificing usability? How these needs are expressed by the requester, and how the responder can trust that the data will not be used for some other reason, are interesting challenges. Khalil and Connelly [46] found that the type of context matters in making policy decisions. While models such as Virtual Walls can be extended to protect different types of context (beyond the “personal” and “general” categories), how to do so while maintaining the usability of such models remains a challenge. Lastly, an adversary may be able to infer restricted context information from other available context information. Care must be taken, therefore, to ensure that context is not inadvertently leaked. Preventing such inferences is indeed challenging, and is an important area for future research.

*Challenge 2: Anonymous tasking:* One of the exciting aspects of the new opportunistic-sensing paradigm is that user devices may be tasked to sense areas beyond the reach of a system’s static infrastructure. Tasking of users, however, can be a threat to their privacy. For example, assigning the task “measure temperature at location X” to Alice may reveal the time at which she visited location X. Furthermore, Alice may not want the system to know her location at the instant when she was tasked. It is important, therefore, to ensure that users can be tasked anonymously. In many cases, *what* is tasked is more important than *who* is tasked.

Delivering tasks to anonymous users may seem easy. For example, users can choose to not identify themselves with the system, and receive tasks through beacons. Users who wish to accept the tasks can notify the system (without revealing their identity) of their acceptance, and then finally report the data. In our *AnonySense* system [21], [47] users periodically download all available tasks from a *Tasking Service* when they are in public locations. The Tasking Service thus learns only that some user in some public location downloaded tasks, and thus the user’s privacy is maintained.

Such an approach, however, does not guarantee that the user is likely to deliver on that task. It would be useful if the system could predict which users are most likely to visit a particular location and opportunistically task those devices directly. Or, perhaps, the system may want to task certain “reliable” people, or classes of people, for a particular task. Reddy et al., for example, are developing a mechanism to quantify the reliability of participants in a participatory sensing application [48]. Such functionality is at odds with maintaining anonymity, and an approach that balances anonymity with location-prediction algorithms and reputation-based tasking is needed.

One solution that looks promising is attribute-based authentication, which ensures that users can authenticate themselves by revealing only their attributes, and not their identities. For example Alice might reveal that she is a “student at Dartmouth” without disclosing her identity. Several cryptographic primitives exist that allow users to prove they possess a certain set of credentials without revealing their identity. Following such authentication, users can be given specific tasks as long

as they possess the requisite attributes. Such techniques will interfere with any approach that relies on user identity for ensuring the integrity of sensed data. Solutions that address privacy, therefore, will need to be designed with data integrity in mind. We discuss challenges for integrity in Section V-B.

Another solution would be to allow the system to task specific users by identity, but without knowing their current location, providing them with location privacy [49], [50]. For example, the system may want to task Alice, specifically, to sense a given location X. A mechanism to deliver tasks to Alice, whose present location is hidden, could rely on an anonymizing network such as Tor [51].

The user being tasked is not the only one potentially at risk: the person issuing tasks needs to be considered as well. Perhaps users will be willing to accept tasks only from certain other users, which brings up the question of how much privacy should be supported for queriers. Perhaps techniques from trust negotiation could be used [52], in which both the querier and the user being tasked have their own requirements of each other. Some trust-negotiation techniques [53, for example] address the privacy of users in trust negotiation.

*Challenge 3: Anonymous data reporting:* Regardless of whether tasks were delivered anonymously, users may want to report sensor data without the system knowing their current location. For example, Alice may want to report her personal sensor data to a trusted database, without the system knowing where she is at the time she uploads the data. Protecting Alice’s location from the system infrastructure, however, can be challenging—simply identifying the wireless access point through which she makes a network connection can localize Alice.

As we do in AnonySense [21], one solution is to use an anonymizing network to hide Alice’s location while she is reporting data, e.g., by bouncing data between the anonymizing network’s nodes several times before the data goes to the database. If one organization manages all the nodes, however, a system administrator may be able to correlate routing information and infer Alice’s location. Trusted-computing hardware platforms (such as the Trusted Platform Module) may solve this problem by ensuring that even the human operators of access points are not privy to sensitive routing information.

Tang et al. [9] propose a technique called *hitchhiking*, where users report data anonymously, but only include characteristics about a location. We disagree with their claim that no personal information about the individual is leaked. For example, Alice may be known to be the only frequent user of a coffee shop at 2am on Sunday nights. A report from that location at around that time could imply that Alice is at the coffee shop. Additional techniques such as blurring [54] would be needed to reduce such inferences. If Alice is concerned about the privacy of her location at the time she *collected* the sensed data, she could blur the location and time of the sensed information—to add some random jitter, within acceptable limits. Such perturbations may make it difficult to correlate or aggregate sensor information, however, unless the amount of error can be constrained. It may be possible to use techniques developed

for privacy-sensitive data-mining and database queries [55].

Our AnonySense system [47] supports *automatic spatiotemporal blurring* of the time and location in a report by tessellating the geographical region into *tiles*. Users report sensor data at the granularity of tiles in discrete time intervals, which provide a statistical guarantee of  $k$ -anonymity [56] ( $k$  users are expected to visit that tile every 5 minutes, for example).

It is also important to anonymize the sensor data itself, of course. One approach, based on  $k$ -anonymity, would be to combine at least  $k$  reports before they are revealed. This would ensure that there are enough people present in a particular reporting area, adding enough “confusion” in the data to make it difficult to pinpoint exact times and locations (and sensor data) for the individuals reporting the data. Another approach is to develop protocols that allow people to contribute sensitive data to an untrustworthy data collector, but with mechanisms that allow them to prove negligence if a data leak occurs [57].

In some settings,  $k$ -anonymity may be threatened by *shadow attacks* [58], in which mobile users obtain service from an overly curious service provider. Even when requests are made anonymously, if the provider can observe future actions made by the user then the provider may be able to link requests to the user.

Another interesting problem related to anonymous reporting is that of status reporting. Users in the system may be interested in querying the *status* of various sensors in the network, and then making tasking decisions based on the characteristics of available sensors. Responses to status queries, however, may reveal enough information about the sensors’ carriers to threaten their privacy. Such functionality appears to be at odds with anonymous tasking, however; perhaps a standard subset of information can be reported to the system, e.g., location and availability of (the anonymous) nodes.

## B. Integrity Issues

If an opportunistic-sensing system provides anonymity to those nodes that are tasked, and to those nodes that submit reports, it is difficult to guarantee the integrity of information. If a user misbehaves by falsifying data, that user cannot be blocked from reporting more data if allowed full anonymity. Finding a solution that balances privacy with data integrity will be a major challenge in such systems. We discuss these integrity challenges next.

*Challenge 4: Reliable data readings:* Traditional sensor networks face the problem of data reliability; the literature already notes limitations on detecting or correcting data-pollution attacks [8] (for example). In people-centric environments, however, trustworthiness of data becomes more crucial: the adversary is no longer only a malicious outsider capturing a subset of sensor nodes; now any participant with an appropriately-configured device can report falsified data. Because users are in control of their own devices, they can more easily launch such an attack. Furthermore, given the more personal nature of applications, they have an increased interest in doing so. For instance, privacy issues or activities that expose anomalous behavior can now be considered strong

incentives for people to perform simple data-pollution attacks to protect themselves. How could a system prevent, detect or correct such behaviors?

Attacks against the sensed-data reliability can be launched using software or hardware. A misbehaving user may first attempt to tamper with the software running on the personal device hosting the sensor(s)—this is a relatively easy-to-perform attack requiring some advanced computer/hacker skills. In this case, the use of trusted hardware can provide a possible practical solution. This has been, for instance, the approach our AnonySense system [21] uses for preventing such software attacks: by leveraging a Trusted Platform Module (TPM) [43] within the sensor node, the proper configuration and sensing functionality can be verified by some designated registration authority and then ensured by the trustworthy hardware.

However, specialized hardware can be used also by misbehaving users to produce incorrect sensed-data values. Here, a promising solution area is *redundancy*. In a large-scale urban-sensing infrastructure, it may be possible to provide multiple sensor nodes with the same task. The system could then use appropriate statistical processing algorithms to estimate the validity of the data, or to remove polluted subsets. For instance, appropriate statistical mechanisms could detect bias or correct malicious responses in survey data. Moreover, since incorrect data can be introduced not only due to malicious user actions but also due to erroneous configurations of the sensor devices, additional correction mechanisms can be considered. Using these statistical mechanisms, the system could report confidence levels about the quality of reports arriving from sensors. The system could then provide anonymous feedback [59] to the sensor nodes that have returned low-confidence reports, following which honest nodes can recalibrate their sensors. The system could also temporarily “revoke” the node’s ability to report data using an anonymous blocking or revocation system [60]–[62]. Such mechanisms even support “three-strikes-out” policies, where an anonymous node, which has been given ample feedback, and yet has sent multiple bad reports, would be blocked. Finally, certain ground-truth enforcing techniques can be used to increase the confidence level of sensed data measurements. For instance, if the system infrastructure includes some physically secure access points with sensors of their own (as proposed in MetroSense [1]), one can use these measurements as representative high-weight or “ground-truth” data points that help in detecting and correcting falsification attacks. Alternately, support from the network infrastructure can attest to the time and location that sensor data was uploaded [63].

Another approach, based on game-theoretic principles, is to couple strong pollution-detection capabilities with punishment strategies (such as exclusion of misbehaving parties [64], [65]). If applied selectively, this approach could encourage fairness and provide a strong disincentive for misbehavior. Alternatively, a reputation-based algorithm could be used as incentive for cooperation by assigning more weight to data returned by trustworthy sensors, perhaps determined by consensus. A complicating constraint on any of these solutions

is to maintain the desired anonymity properties while also enforcing punishment or measuring reputation.

*Challenge 5: Data authenticity:* Another fundamental security problem is data authentication; that is, how can sensed data be delivered to a database or application with assurance that no intermediate devices have tampered with it? Furthermore, in applications that care about the identity of the user carrying the sensor (for example, to report Alice’s location on an active map), how can the system ensure the identity of the sensor node and its custodian?

A significant amount of research has been conducted on data authenticity in static sensor networks, especially in the context of secure in-network data aggregation. Several threat models have been studied, including fault-tolerant computation of aggregates [66, for example] and secure aggregation in the presence of single [67], [68] or multiple [7], [8], [69], [70, for example] adversarial sensor nodes. Also, various models of in-network data aggregation have been considered, including the star-based [8], [71], [72] and tree-based aggregation models [7], [73]. Aggregation of encrypted data has been also studied [74]–[76].

Unfortunately, these solutions are not feasible for mobile sensor networks. First, previous authentication techniques assume the availability of a stable topological tree that spans the network’s sensors. In a highly dynamic and mobile network no such tree structure exists or can be maintained. At the same time, opportunistic networking, tasking, data muling, and delay-tolerant data reporting create an inherently complex communication environment in which few assumptions can be made about node integrity. Moreover, the previous research using mobile sensor nodes has focused on sensors attached to animals and vehicles.

Human carriers introduce new challenges for maintaining node integrity. Although we still need low-level mechanisms to authenticate a sensor node, we need more sophisticated notions of user identity, group membership, and other attributes—for instance, *group signatures* can be employed as in the AnonySense system [21] to anonymously verify the validity of mobile sensing nodes. In many human-centric sensing applications these factors may evolve over time, perhaps quickly, due to privacy profile updates, role changes, revocation of reporting rights, or task-specific requirements. This complexity and dynamism in people-centric sensing introduces additional security concerns regarding system trustworthiness.

Thus we need new schemes to authenticate the user and her sensor device, and to ensure the integrity of the data, based on few assumptions about network topology. We have some advantages relative to existing sensor-network literature: we may assume that sensor nodes are fairly capable platforms, and that they frequently come into direct contact with network access points. Desired properties include computational- and bandwidth-efficient authentication, topology-independent techniques, resilience against malicious subsets of nodes, and resistance against denial-of-service attacks. Results on the authentication of streams in arbitrary adversarial networks using cryptographically-enhanced error-correcting techniques [77],

[78] may give some insight in this direction.

*Challenge 6: System integrity:* Another important aspect of integrity relates to the functionality of the underlying system. In principle, any operation (in addition to data collection) performed in a mobile, opportunistic, participatory<sup>1</sup>, and highly untrusted sensor network will need some form of correctness verification.

For instance, these people-centric sensor networks, likely based on pervasive devices such as cell phones, are an application platform as well as a sensing device. These applications need to be able to authenticate the information they receive from the system, and to determine when to trust the tasks they are asked to perform. One possible solution is to develop task-specification languages allowing a node to verify that execution of a particular task will not compromise the node itself. Such a procedure presents a trade-off between simplicity of verification and flexibility in describing complex tasks. A long history of research on the safe execution of mobile code [79] should be instructive.

Also, information centers will need to validate the integrity of actions (e.g., tasking, data muling, data collection) performed by mobile nodes at remote sites, often in a delay-tolerant fashion. How, for instance, can the system administrator of a sensing-based information system be sure that tasking is executed correctly or that reported data not only is authentic but also is temporally relevant? Suggested solutions on this front include developing efficient verification protocols that ensure secure data management and guarantee system integrity. Conceptually, the system administrator could maintain some secure cryptographic state of the entire system configuration, consisting of various topological, temporal, and user-related parameters. This cryptographic state could be updated and used for verification as network nodes execute queries, tasking commands, and other operations. This general approach recalls early results on memory consistency checking [80].

In general, any mechanism designed for system integrity depends on the specific threat or trust model. For instance, in an application where the focus lies on collection and off-line fusion of sensed data, certain components of infrastructure may be considered trusted by the system administrator; for example, “sensor access points” in MetroSense [1] may be considered to have higher levels of trust, whether implemented through physical protection or the use of secure hardware.

### C. Availability Issues

Researchers have addressed some of the denial of service (DoS) issues for static sensor networks—for example, adversaries may gain control of sensors and try to flood the network or subvert the routing protocols to affect the network’s reporting capabilities [81]. Opportunistic sensing, however, introduces different DoS issues because devices are already in the hands of potential adversaries. On the other hand,

<sup>1</sup>System integrity is a central concern in any participatory distributed system. For instance, in Seti@Home, the biggest distributed computing project on the planet, 50% of the project’s resources have been spent dealing with various security problems.

because we assume high connectivity (sensors are frequently in contact with a Wi-Fi or cellular access points), the sensing infrastructure itself is less susceptible—though not entirely immune—to DoS than in some sensor-network scenarios. We briefly mention three distinct aspects related to the availability of urban sensing systems.

*Challenge 7: Preventing data suppression:* Although most opportunistic sensor networks are designed as a best-effort service, in which sensor data are submitted by nodes volunteered by their human owners, DoS attacks limit the availability of data and are still a concern. If nodes are meant to be tasked opportunistically, their owners may configure their device to refuse to accept certain tasks, or to accept but then ignore the tasks. (Contrast this behavior with “traditional” sensor networks, where installed sensors faithfully report data unless they break or are compromised by an external adversary.) Data-consuming applications are left at the mercy and whims of the individual participants. This problem is exacerbated by privacy concerns—if users feel unable to control access to their data, they will be disinclined to carry the device or permit tasks that report such sensor data. This effect reinforces the importance of the privacy challenges above, and leads directly to the next challenge.

*Challenge 8: Participation:* Unlike static sensor networks, where sensors are embedded in the infrastructure, opportunistic people-centric sensing relies on the sensor devices carried by users. One of the major challenges today is to create incentives for users to carry such devices, especially if bulky or a threat to their privacy. The best approach is to incorporate the sensors into a device they want to carry (such as a cell phone or music player), and provide incentives that are compatible with users’ needs and interests. Depending on the application’s goals, the sensing coverage and density of the sensing system are also important issues, i.e., what percentage of a city is covered or what percentage of the population carries a certain sensor type. If the device has applications with clear direct and indirect benefits for the user, has strong, usable measures that protect privacy, and is easy to carry, people will use it. In general, the more data-secure and user-private a participatory human-centric sensing system is, the more individuals are likely to contribute to its services.

Another promising direction towards increasing people’s participation in these systems is the use of game-theoretic approaches. Users’ behavior in participatory systems can be studied using *privacy-aware hybrid payoff models*, where users care (primarily) about the services they benefit from, but also (secondarily) about the privacy loss they experience during their participation in the system. In particular, privacy-aware mechanism design using rational cryptography (cf. [82]) could lead to the design of concrete practical protocols for opportunistic context sensing, that would be followed by rational participants under reasonable assumptions about their system-usage preferences and their privacy concerns.

*Challenge 9: Fairness:* Many people-centric sensor applications provide direct benefits to the users who carry the sensor nodes. Thus, people have an incentive to cheat, to obtain

better service for themselves than they fairly deserve. For example, users may cause the applications to task many other sensors to collect information for their own needs, without being willing to take on tasks for other users. Research on incentive-compatible peer-to-peer systems [83] or, as above, privacy-aware game-theoretic mechanism design, may provide useful insight here.

#### D. Challenges in Participatory Sensing

In *participatory sensing*, the sensing operation requires some kind of human intervention [84]. For example, the task may ask the user to take a photograph of the menu when she visits a restaurant, or to comment on her opinion about the food at the restaurant, or to record the gas prices when she passes a filling station [85]. The human element adds additional security challenges. With respect to privacy, the user may leak more information about his or her identity by the nature of his or her response. For example, the (artistic) composition of the user's photographs may reveal information about the user, or the background may reveal additional context beyond the information being sensed. A user's textual response may contain grammatical errors, abbreviations, or other clues about the user's identity. With respect to integrity, the user now has an easy way to fabricate "sensor" data of his or her choosing, and the reliability of data now depends on the ability of the user to respond to tasks effectively. With respect to availability, users can easily suppress data by either not responding to tasks, or by responding to tasks with data that is not useful. Participation will be reduced because requiring human intervention is disruptive, and users would need incentives to participate in such sensing.

## VI. SUMMARY

In this paper we recognize the current interest in opportunistic people-centric sensing applications and infrastructure. Such applications have clear and substantial security and privacy challenges, which must be resolved if these systems have any hope of realizing their potential. We described 9 challenges, and offered conceptual solution approaches for each.

Challenge 1: Context privacy

Challenge 2: Anonymous tasking

Challenge 3: Anonymous data reporting

Challenge 4: Reliable data readings

Challenge 5: Data authenticity

Challenge 6: System integrity

Challenge 7: Preventing data suppression

Challenge 8: Participation

Challenge 9: Fairness

We hope that this paper will create a substantive discussion around these challenges and encourage all researchers involved with opportunistic sensing, participatory sensing, urban sensing, and people-centric sensor networks to address security and privacy at a fundamental level within their system and application design. The future of the paradigm depends on it.

## ACKNOWLEDGMENTS

Many thanks to Peter C. Johnson, of Dartmouth College, for his contributions to earlier versions of this paper. We also thank Vijay Bhuse, Andrew Campbell, George Cybenko, Shane Eisenman, Minh Shin and the rest of the MetroSense team at Dartmouth College for their helpful comments.

This research program is a part of the Institute for Security, Technology, and Society (ISTS), and was supported by the Bureau of Justice Assistance under grant number 2005-DD-BX-1091 from the U.S. Department of Justice, by NIST under award 60NANB6D6130 from the U.S. Department of Commerce, by NCSA under Grant Award Number 2006-CS-001-000001 from the U.S. Department of Homeland Security, and by the Institute for Information Infrastructure Protection (I3P) under an award from the U.S. Department of Homeland Security. Partial support was also provided by the Center for Algorithmic Game Theory at the University of Aarhus under an award from the Carlsberg Foundation. Points of view or opinions in this document are those of the authors and do not represent the sponsors' official position or policies.

## REFERENCES

- [1] A. T. Campbell, S. B. Eisenman, N. D. Lane, E. Miluzzo, and R. A. Peterson, "People-centric urban sensing," in *Proceedings of the Second Annual International Wireless Internet Conference (WICON)*. ACM Press, Aug. 2006, pp. 18–31. DOI: 10.1145/1234161.1234179
- [2] A. Perrig, R. Szewczyk, J. Tygar, V. Wen, and D. E. Culler, "SPINS: Security protocols for sensor networks," *Wireless Networks*, vol. 8, no. 5, pp. 521–534, Sep. 2002. DOI: 10.1023/A:1016598314198
- [3] S. Zhu, S. Setia, and S. Jajodia, "Leap: efficient security mechanisms for large-scale distributed sensor networks," in *Proceedings of the ACM Conference on Computer and Communications Security (CCS)*. ACM Press, 2003, pp. 62–72.
- [4] C. Karlof and D. Wagner, "Secure routing in wireless sensor networks: attacks and countermeasures," in *Proceedings of the IEEE International Workshop on Sensor Network Protocols and Applications (SNPA)*. IEEE Press, May 2003, pp. 113–127.
- [5] C. Yin, S. Huang, P. Su, and C. Gao, "Secure routing for large-scale wireless sensor networks," in *Proceedings of the International Conference on Communication Technology (ICCT)*, Apr. 2003.
- [6] J. Deng, R. Han, and S. Mishra, "A performance evaluation of intrusion-tolerant routing in wireless sensor networks," in *Proceedings of the International Workshop on Information Processing in Sensor Networks (IPSN)*, Apr. 2003.
- [7] H. Chan, A. Perrig, and D. Song, "Secure hierarchical in-network aggregation in sensor networks," in *Proceedings of the ACM Conference on Computer and Communications Security (CCS)*. ACM Press, Oct. 2006, pp. 278–287. DOI: 10.1145/1180405.1180440
- [8] B. Przydatek, D. Song, and A. Perrig, "SIA: Secure information aggregation in sensor networks," in *Proceedings of the International Conference on Embedded Networked Sensor Systems (SenSys)*. ACM Press, 2003.
- [9] K. P. Tang, P. Keyani, J. Fogarty, and J. I. Hong, "Putting people in their place: An anonymous and privacy-sensitive approach to collecting sensed data in location-based applications," in *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI)*. ACM Press, 2006, pp. 93–102. DOI: 10.1145/1124772.1124788
- [10] K. Piotrowski, P. Langendoerfer, and S. Peter, "How public key cryptography influences wireless sensor node lifetime," in *Proceedings of the ACM Workshop on Security of Ad Hoc and Sensor Networks (SASN)*. ACM Press, 2006, pp. 169–176.
- [11] T. Abdelzaher, Y. Anokwa, P. Boda, J. Burke, D. Estrin, L. Guibas, A. Kansal, S. Madden, and J. Reich, "Mobiscopes for human spaces," *IEEE Pervasive Computing*, vol. 6, no. 2, pp. 20–29, 2007. DOI: 10.1109/MPRV.2007.38
- [12] O. Riva and C. Borcea, "The Urbanet revolution: Sensor power to the people!" *IEEE Pervasive Computing*, vol. 6, no. 2, pp. 41–49, 2007. DOI: 10.1109/MPRV.2007.46

- [13] "CENS Urban Sensing project," [http://research.cens.ucla.edu/projects/2006/Systems/Urban\\_Sensing/](http://research.cens.ucla.edu/projects/2006/Systems/Urban_Sensing/), 2006, website visited December 2008.
- [14] N. D. Lane, S. B. Eisenman, M. Musolesi, E. Miluzzo, and A. T. Campbell, "Urban sensing systems: opportunistic or participatory?" in *Proceedings of the Workshop on Mobile Computing Systems and Applications (HotMobile)*. ACM Press, 2008, pp. 11–16. DOI: 10.1145/1411759.1411763
- [15] J. Hwang and Y. Kim, "Revisiting random key pre-distribution schemes for wireless sensor networks," in *Proceedings of the ACM Workshop on Security of Ad Hoc and Sensor Networks (SASN)*, Oct. 2004, pp. 43–52.
- [16] D. Huang, M. Mehta, D. Medhi, and L. Harn, "Location-aware key management scheme for wireless sensor networks," in *Proceedings of the ACM Workshop on Security of Ad Hoc and Sensor Networks (SASN)*, Oct. 2004, pp. 29–42.
- [17] S. Yi and R. Kravets, "Composite key management for ad hoc networks," in *Proceedings of the International Conference on Mobile and Ubiquitous Systems: Networking and Services (MobiQuitous)*, Aug. 2004, pp. 52–61.
- [18] P. K. Dutta, J. W. Hui, D. C. Chu, and D. E. Culler, "Securing the deluge network programming system," in *Proceedings of the International Workshop on Information Processing in Sensor Networks (IPSN)*. ACM Press, 2006, pp. 326–333. DOI: 10.1145/1127777.1127826
- [19] P. E. Lanigan, R. Gandhi, and P. Narasimhan, "Sluice: Secure dissemination of code updates in sensor networks," in *Proceedings of the IEEE International Conference on Distributed Computing Systems (ICDCS)*. IEEE Press, 2006, p. 53. DOI: 10.1109/ICDCS.2006.77
- [20] B. Hull, V. Bychkovsky, Y. Zhang, K. Chen, M. Goraczko, A. K. Miu, E. Shih, H. Balakrishnan, and S. Madden, "CarTel: A distributed mobile sensor computing system," in *Proceedings of the International Conference on Embedded Networked Sensor Systems (SenSys)*. ACM Press, Nov. 2006, pp. 125–138. DOI: 10.1145/1182807.1182821
- [21] C. Cornelius, A. Kapadia, D. Kotz, D. Peebles, M. Shin, and N. Triandopoulos, "AnonySense: Privacy-aware people-centric sensing," in *Proceedings of the International Conference on Mobile Systems, Applications and Services (MobiSys)*. ACM Press, Jun. 2008, pp. 211–224. DOI: 10.1145/1378600.1378624
- [22] S. B. Eisenman, E. Miluzzo, N. D. Lane, R. A. Peterson, G.-S. Ahn, and A. T. Campbell, "The BikeNet mobile sensing system for cyclist experience mapping," in *Proceedings of the International Conference on Embedded Networked Sensor Systems (SenSys)*. ACM Press, 2007, pp. 87–101. DOI: 10.1145/1322263.1322273
- [23] E. Miluzzo, N. Lane, K. Fodor, R. Peterson, S. Eisenman, H. Lu, M. Musolesi, X. Zheng, and A. Campbell, "Sensing meets mobile social networks: The design, implementation and evaluation of the CenceMe application," in *Proceedings of the International Conference on Embedded Networked Sensor Systems (SenSys)*. ACM Press, 2008, pp. 337–350. DOI: 10.1145/1460412.1460445
- [24] "Urban atmospheres," 2007, <http://www.urban-atmospheres.net>.
- [25] A. Kansal, S. Nath, J. Liu, and F. Zhao, "SenseWeb: An infrastructure for shared sensing," *IEEE MultiMedia*, vol. 14, no. 4, pp. 8–13, 2007. DOI: 10.1109/MMUL.2007.82
- [26] C. Frank, P. Bolliger, C. Roduner, and W. Kellerer, "Objects calling home: Locating objects using mobile phones," in *Proceedings of the International Conference on Pervasive Computing (Pervasive)*, ser. Lecture Notes in Computer Science, vol. 4480. Springer, May 2007, pp. 351–368. DOI: 10.1007/978-3-540-72037-9\_19
- [27] J. Froehlich, J. Neumann, and N. Oliver, "Measuring the pulse of the city through shared bicycle programs," in *Proceedings of the International Workshop on Urban, Community, and Social Applications of Networked Sensing Systems (UrbanSense08)*, Nov. 2008.
- [28] R. Want, A. Hopper, V. Falcão, and J. Gibbons, "The Active Badge location system," *ACM Transactions on Information Systems*, vol. 10, no. 1, pp. 91–102, Jan. 1992. DOI: 10.1145/128756.128759
- [29] M. Davis, S. King, N. Good, and R. Sarvas, "From context to content: leveraging context to infer media metadata," in *Proceedings of the ACM International Conference on Multimedia*. ACM Press, 2004, pp. 188–195. DOI: 10.1145/1027527.1027572
- [30] B. Hoh, M. Gruteser, R. Herring, J. Ban, D. Work, J.-C. Herrera, A. M. Bayen, M. Annavaram, and Q. Jacobson, "Virtual trip lines for distributed privacy-preserving traffic monitoring," in *Proceedings of the International Conference on Mobile Systems, Applications and Services (MobiSys)*. ACM Press, 2008, pp. 15–28. DOI: 10.1145/1378600.1378604
- [31] J. Eriksson, L. Girod, B. Hull, R. Newton, S. Madden, and H. Balakrishnan, "The Pothole Patrol: using a mobile sensor network for road surface monitoring," in *Proceedings of the International Conference on Mobile Systems, Applications and Services (MobiSys)*. ACM Press, 2008, pp. 29–39. DOI: 10.1145/1378600.1378605
- [32] P. Mohan, V. Padmanabhan, and R. Ramjee, "Nericell: Rich monitoring of road and traffic conditions using mobile smartphones," in *Proceedings of the International Conference on Embedded Networked Sensor Systems (SenSys)*. ACM Press, 2008, pp. 323–336. DOI: 10.1145/1460412.1460444
- [33] E. Agapie, G. Chen, D. Houston, E. Howard, J. Kim, M. Y. Mun, A. Mondschein, S. Reddy, R. Rosario, J. Ryder, A. Steiner, J. Burke, E. Estrin, M. Hansen, and M. Rahimi, "Seeing our signals: combining location traces and web-based models for personal discovery," in *Proceedings of the Workshop on Mobile Computing Systems and Applications (HotMobile)*. ACM Press, 2008, pp. 6–10. DOI: 10.1145/1411759.1411762
- [34] "Phones studied as attack detector," USA Today, [http://www.usatoday.com/tech/news/techpolicy/2007-05-03-cellphone-attack-detector\\_N.htm](http://www.usatoday.com/tech/news/techpolicy/2007-05-03-cellphone-attack-detector_N.htm), 3 May 2007.
- [35] B. Gedik and L. Liu, "Location privacy in mobile systems: A personalized anonymization model," in *Proceedings of the IEEE International Conference on Distributed Computing Systems (ICDCS)*, Jun. 2005, pp. 620–629. DOI: 10.1109/ICDCS.2005.48
- [36] B. Hoh and M. Gruteser, "Protecting location privacy through path confusion," in *Proceedings of the IEEE/CreateNet Intl. Conference on Security and Privacy for Emerging Areas in Communication Networks (SecureComm)*, Sep. 2005.
- [37] C. C. Tan, H. Wang, S. Zhong, and Q. Li, "Body sensor network security: an identity-based cryptography approach," in *Proceedings of the ACM Conference on Wireless Network Security (WiSec)*. ACM Press, 2008, pp. 148–153. DOI: 10.1145/1352533.1352557
- [38] J. Cornwell, I. Fette, G. Hsieh, M. Prabhakar, J. Rao, K. Tang, K. Vaniea, L. Bauer, L. Cranor, J. Hong, B. McLaren, M. Reiter, and N. Sadeh, "User-controllable security and privacy for pervasive computing," in *Proceedings of the Workshop on Mobile Computing Systems and Applications (HotMobile)*, Mar. 2007, pp. 14–19. DOI: 10.1109/HotMobile.2007.9
- [39] A. Kapadia, T. Henderson, J. Fielding, and D. Kotz, "Virtual walls: Protecting digital privacy in pervasive environments," in *Proceedings of the International Conference on Pervasive Computing (Pervasive)*, ser. Lecture Notes in Computer Science, vol. 4480. Springer-Verlag, May 2007, pp. 162–179. DOI: 10.1007/978-3-540-72037-9\_10
- [40] S. Lederer, A. K. Dey, and J. Mankoff, "A conceptual model and a metaphor of everyday privacy in ubiquitous computing," Intel Research Berkeley, Tech. Rep. IRB-TR-02-017, 2002.
- [41] O. Goldreich, S. Micali, and A. Wigderson, "How to play any mental game or a completeness theorem for protocols with honest majority," in *Proceedings of the ACM Symposium on Theory of Computing (STOC)*, 1987, pp. 218–229.
- [42] A. Yao, "Protocols for secure computations," in *Proceedings of the IEEE Symposium on Foundations of Computer Science (FOCS)*. IEEE Press, 1982, pp. 160–164.
- [43] "Trusted Computing Group (TCG)," <https://www.trustedcomputinggroup.org/home>, May 2005, website visited December 2008.
- [44] S. W. Smith and S. Weingart, "Building a high-performance, programmable secure coprocessor," *Computer Networks (Special Issue on Computer Network Security)*, vol. 31, pp. 831–860, Apr. 1999.
- [45] S. Consolvo, I. E. Smith, T. Matthews, A. LaMarca, J. Tabert, and P. Powledge, "Location disclosure to social relations: why, when, and what people want to share," in *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI)*, 2005, pp. 81–90.
- [46] A. Khalil and K. Connelly, "Context-aware telephony: privacy preferences and sharing patterns," in *Proceedings of the Conference on Computer Supported Cooperative Work (CSCW)*, 2006, pp. 469–478.
- [47] A. Kapadia, N. Triandopoulos, C. Cornelius, D. Peebles, and D. Kotz, "AnonySense: Opportunistic and privacy-preserving context collection," in *Proceedings of the International Conference on Pervasive Computing (Pervasive)*, ser. Lecture Notes in Computer Science, vol. 5013. Springer-Verlag, May 2008, pp. 280–297. DOI: 10.1007/978-3-540-79576-6\_17
- [48] S. Reddy, K. Shilton, J. Burke, D. Estrin, M. Hansen, and M. Srivastava, "Evaluating participation and performance in participatory sensing," in

- Proceedings of the International Workshop on Urban, Community, and Social Applications of Networked Sensing Systems (UrbanSense08)*, Nov. 2008.
- [49] J. Al-Muhtadi, R. H. Campbell, A. Kapadia, D. Mickunas, and S. Yi, "Routing through the Mist: Privacy preserving communication in ubiquitous computing environments," in *Proceedings of the IEEE International Conference on Distributed Computing Systems (ICDCS)*, 2002, pp. 74–83.
- [50] M. Gruteser and D. Grunwald, "Enhancing location privacy in wireless LAN through disposable interface identifiers: a quantitative analysis," *Mobile Networks and Applications*, vol. 10, no. 3, pp. 315–325, 2005. DOI: 10.1145/1145911.1145917
- [51] R. Dingleline, N. Mathewson, and P. Syverson, "Tor: The second-generation onion router," in *Proceedings of the USENIX Security Symposium*, Aug. 2004, pp. 303–320.
- [52] W. H. Winsborough, K. E. Seamons, and V. E. Jones, "Automated trust negotiation," in *Proceedings of the DARPA Information Survivability Conference and Exposition (DISCEX)*, Jan. 2000, pp. 88–102. DOI: 10.1109/DISCEX.2000.824965
- [53] K. Frikken, J. Li, and M. Atallah, "Trust negotiation with hidden credentials, hidden policies, and policy cycles," in *Proceedings of the Annual Symposium on Network and Distributed System Security (NDSS)*, Feb. 2006, pp. 157–172.
- [54] G. Iachello, I. Smith, S. Consolvo, M. Chen, and G. D. Abowd, "Developing privacy guidelines for social location disclosure applications and services," in *Proceedings of the Symposium on Usable Privacy and Security (SOUPS)*, Jul. 2005, pp. 65–76. DOI: 10.1145/1073001.1073008
- [55] V. S. Verykios, E. Bertino, I. N. Fovino, L. P. Provenza, Y. Saygin, and Y. Theodoridis, "State-of-the-art in privacy preserving data mining," *ACM SIGMOD Record*, vol. 3, no. 1, pp. 50–57, Mar. 2004. DOI: 10.1145/974121.974131
- [56] L. Sweeney, "k-anonymity: A model for protecting privacy," *International Journal of Uncertainty, Fuzziness, and Knowledge-Based Systems*, vol. 10, no. 5, pp. 557–570, October 2002.
- [57] P. Golle, F. McSherry, and I. Mironov, "Data collection with self-enforcing privacy," in *Proceedings of the ACM Conference on Computer and Communications Security (CCS)*. ACM Press, Oct. 2006, pp. 69–78. DOI: 10.1145/1180405.1180416
- [58] D. Riboni, L. Pareschi, and C. Bettini, "Shadow attacks on users' anonymity in pervasive computing environments," *Pervasive and Mobile Computing*, vol. 4, no. 6, pp. 819–835, Dec. 2008. DOI: 10.1016/j.pmcj.2008.04.008
- [59] L. Sassaman, B. Cohen, and N. Mathewson, "The Pynchon Gate: A secure method of pseudonymous mail retrieval," in *Proceedings of the Workshop on Privacy in the Electronic Society (WPES)*, Nov. 2005, pp. 1–9. DOI: 10.1145/1102199.1102201
- [60] P. C. Johnson, A. Kapadia, P. P. Tsang, and S. W. Smith, "Nymble: Anonymous IP-address blocking," in *Proceedings of the International Symposium on Privacy Enhancing Technologies (PET)*, ser. Lecture Notes in Computer Science, vol. 4776. Springer-Verlag, Jun. 2007, pp. 113–133. DOI: 10.1007/978-3-540-75551-7\_8
- [61] P. P. Tsang, M. H. Au, A. Kapadia, and S. W. Smith, "Blacklistable anonymous credentials: blocking misbehaving users without TTPs," in *Proceedings of the ACM Conference on Computer and Communications Security (CCS)*. ACM Press, Oct. 2007, pp. 72–81. DOI: 10.1145/1315245.1315256
- [62] —, "BLAC: Revoking repeatedly misbehaving anonymous users without relying on TTPs," Dartmouth College, Tech. Rep. TR2008-635, Oct. 2008.
- [63] A. Parker, S. Reddy, T. Schmid, K. Chang, G. Saurabh, M. Srivastava, M. Hansen, J. Burke, D. Estrin, M. Allman, and V. Paxson, "Network system challenges in selective sharing and verification for personal, social, and urban-scale sensing applications," in *Proceedings of the Workshop on Hot Topics in Networks (HotNets)*, Nov. 2006, pp. 37–42.
- [64] I. Abraham, D. Dolev, R. Gonen, and J. Halpern, "Distributed computing meets game theory: Robust mechanisms for rational secret sharing and multiparty computation," in *Proceedings of the ACM Symposium on Principles of Distributed Computing (PODC)*, 2006.
- [65] A. Lysyanskaya and N. Triandopoulos, "Rationality and adversarial behavior in multi-party computation," in *Proceedings of Advances in Cryptology (CRYPTO)*, 2006, pp. 180–197.
- [66] J.-Y. Chen, G. Pandurangan, and D. Xu, "Robust computation of aggregates in wireless sensor networks: distributed randomized algorithms and analysis," in *Proceedings of the International Workshop on Information Processing in Sensor Networks (IPSN)*. IEEE Press, 2005, p. 46.
- [67] L. Hu and D. Evans, "Secure aggregation for wireless networks," in *Proceedings of Workshop on Security and Assurance in Ad hoc Networks*, Jan. 2003.
- [68] P. Jadia and A. Mathuria, "Efficient secure aggregation in sensor networks," in *Proceedings of the International Conference on High Performance Computing (HiPC)*, 2005, pp. 40–49.
- [69] S. Roy, S. Setia, and S. Jajodia, "Attack-resilient hierarchical data aggregation in sensor networks," in *Proceedings of the ACM Workshop on Security of Ad Hoc and Sensor Networks (SASN)*. ACM Press, Oct. 2006, pp. 71–82. DOI: 10.1145/1180345.1180355
- [70] K. B. Frikken and J. A. Dougherty, IV, "An efficient integrity-preserving scheme for hierarchical sensor aggregation," in *Proceedings of the ACM Conference on Wireless Network Security (WiSec)*. ACM Press, 2008, pp. 68–76. DOI: 10.1145/1352533.1352546
- [71] W. Du, J. Deng, Y. S. Han, and P. Varshney, "A witness-based approach for data fusion assurance in wireless sensor networks," in *Proceedings of the IEEE Global Telecommunications Conference (GLOBECOM)*, 2003.
- [72] A. Mahimkar and T. S. Rappaport, "SecureDAV: a secure data aggregation and verification protocol for sensor networks," in *Proceedings of the IEEE Global Telecommunications Conference (GLOBECOM)*, 2004, pp. 2175–2179.
- [73] Y. Yang, X. Wang, S. Zhu, and G. Cao, "SDAP: a secure hop-by-hop data aggregation protocol for sensor networks," in *Proceedings of the ACM International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc)*. ACM Press, 2006, pp. 356–367.
- [74] H. Cam, S. Özdemir, P. Nair, D. Muthuavaniappan, and H. O. Sanli, "Energy-efficient secure pattern based data aggregation for wireless sensor networks," *Computer Communications*, vol. 29, no. 4, pp. 446–455, Feb. 2006. DOI: 10.1016/j.comcom.2004.12.029
- [75] C. Castelluccia, E. Mykletun, and G. Tsudik, "Efficient aggregation of encrypted data in wireless sensor networks," in *Proceedings of the International Conference on Mobile and Ubiquitous Systems: Networking and Services (MobiQuitous)*, 2005, pp. 109–117.
- [76] J. Girao, D. Westhoff, and M. Schneider, "CDA: Concealed data aggregation for reverse multicast traffic in wireless sensor networks," in *Proceedings of the IEEE International Conference on Communications*, 2005, pp. 3044–3049.
- [77] S. Micali, C. Peikert, M. Sudan, and D. A. Wilson, "Optimal error correction against computationally bounded noise," in *Proceedings of the Theory of Cryptography Conference (TCC)*, 2005, pp. 1–16.
- [78] A. Lysyanskaya, R. Tamassia, and N. Triandopoulos, "Multicast authentication in fully adversarial networks," in *Proceedings of the IEEE Symposium on Security and Privacy (S&P)*, May 2004, pp. 241–255.
- [79] C. F. Tschudin, "Mobile agent security," in *Intelligent Information Agents*, M. Klusch, Ed. Springer-Verlag, 1999, ch. 18, pp. 431–445.
- [80] M. Blum, W. Evans, P. Gemmel, S. Kannan, and M. Naor, "Checking the correctness of memories," in *Proceedings of the IEEE Symposium on Foundations of Computer Science (FOCS)*. IEEE Press, Oct. 1991, pp. 90–99. DOI: 10.1109/SFCS.1991.18535
- [81] A. D. Wood and J. A. Stankovic, "Denial of service in sensor networks," *IEEE Computer*, vol. 35, no. 10, pp. 54–62, Oct. 2002.
- [82] P. B. Miltersen, J. B. Nielsen, and N. Triandopoulos, "Privacy-enhancing first-price auctions using rational cryptography," Manuscript. Available at <http://eprint.iacr.org/2008/418>, 2008.
- [83] P. Golle, K. Leyton-Brown, and I. Mironov, "Incentives for sharing in peer-to-peer networks," in *Proceedings of the ACM Conference on Electronic Commerce*, 2001. DOI: 10.1145/501158.501193
- [84] J. Burke, D. Estrin, M. Hansen, A. Parker, N. Ramanathan, S. Reddy, and M. B. Srivastava, "Participatory sensing," in *World Sensor Web Workshop (at Sensys)*, 2006.
- [85] N. Bulusu, C. T. Chou, and S. Kanhere, "Participatory sensing in commerce: Using mobile camera phones to track market price dispersion," in *Proceedings of the International Workshop on Urban, Community, and Social Applications of Networked Sensing Systems (UrbanSense08)*, Nov. 2008.