Verifiable Privacy-Preserving Aggregation in People-Centric Urban Sensing Systems

Rui Zhang, *Member, IEEE,* Jing Shi, Yanchao Zhang, *Senior Member, IEEE,* and Chi Zhang, *Member, IEEE*

***Abstract*—People-centric urban sensing systems (PC-USSs) refer to using human-carried mobile devices such as smartphones and tablets for urban-scale distributed data collection, analysis, and sharing to facilitate interaction between humans and their surrounding environments. A main obstacle to the widespread deployment and adoption of PC-USSs are the privacy concerns of participating individuals as well as the concerns about data integrity. To tackle this open challenge, this paper presents the design and evaluation of VPA, a novel peer-to-peer based solution to verifiable privacy-preserving data aggregation in PC-USSs. VPA achieves strong user privacy by letting each user exchange random shares of its datum with other peers, while at the same time ensures data integrity through a combination of Trusted Platform Module and homomorphic message authentication code. VPA can support a wide range of statistical additive and non-additive aggregation functions such as Sum, Average, Variance, Count, Max/Min, Median, Histogram, and Percentile with accurate aggregation results. The efficacy and efficiency of VPA are confirmed by thorough analytical and simulation results.**

***Index Terms*—People-centric urban sensing system (PC-USS), peer-to-peer, aggregation, security, privacy.**

I. Introduction

PEOPLE-centric urban sensing systems (PC-USSs) refer to using human-carried mobile devices such as smartphones and tablets with ever-growing capabilities in sensing, computation, storage, and communications for urban-scale distributed data collection, analysis, and sharing to facilitate the interaction between humans and their surrounding environments. PC-USSs are expected to open a new era of exciting scientific, social, and commercial applications [2]–[8]. PC-USSs differ significantly from traditional wireless sensor networks that focus on environment sensing and data collection. First, system devices are no longer owned and managed by a single authority but belong to individuals with diverse interests. Second, system devices have much more powerful resources than sensor nodes and can be charged regularly. Third, the system features dynamic node mobility. Fourth, sensing data are more related to the interactions among humans and between humans and their surroundings instead of only about some physical phenomena of interest. Fifth, but not the last, humans are no longer just passive data users but also active data contributors.

The widespread deployment and adoption of PC-USSs face many obstacles, of which user privacy and data integrity are among the most critical [5]–[7]. For instance, in a study of the relationship between air quality and public health, researchers desire some aggregate statistics of personal health data such as heart rates, blood pressure levels, and weights at different sections of an urban area. Individuals may be unwilling to disclose their personal data if there were no guarantee that their data would not be used to invade their privacy. As an example for data-integrity breach, consider applications like CarTel [3] and VTrack [8] that use traffic statistics such as average speed as an indicator of congestion to help system users do route planning. A selfish and malicious driver may prevent other users from choosing his current road by manipulating the aggregation result, i.e., cheating the server into accepting a lower-than-actual average speed that indicates road congestion. These two examples highlight the necessity for verifiable privacy-preserving data aggregation techniques that can ensure strong user privacy and also aggregation integrity.[1]

Designing a verifiable privacy-preserving aggregation scheme for PC-USSs is particulary challenging. On the one hand, ensuring user privacy means that a user's original data cannot be disclosed. This requirement makes it hard to detect if a user has faithfully participated in data aggregation. On the other hand, ensuring aggregation integrity requires any misbehavior during data aggregation to be detected with overwhelming probability. This requirement is extremely difficult to satisfy without knowing users' original data.

The contribution of this paper is the design and evaluation of VPA, a novel peer-to-peer based solution to verifiable privacy-preserving data aggregation in PC-USSs. VPA consists of the following two components.

- The first component VPA$^+$ targets additive aggregation functions such as Sum, Average, and Variance. Its basic idea is to divide the aggregation process into two phases. In the first phase, each node submits a commitment to the aggregation server, which is a homomorphic message authentication code of its original datum. The homomorphic property of commitments enables the aggregation server to compute the aggregate commitment corresponding to the final aggregate without the ability to recover any

Manuscript received February 11, 2012; revised July 14, 2012.

R. Zhang is with the Department of Electrical Engineering, University of Hawaii, Honolulu, HI, 96822 (e-mail: ruizhang@hawaii.edu).

J. Shi is with the School of Public Administration, Huazhong University of Science and Technology, China (e-mail: shi.jing@mail.hust.edu.cn).

Y. Zhang is with the School of Electrical, Computer, and Energy Engineering, Arizona State University, Tempe, AZ, 85287 (e-mail: yczhang@asu.edu).

C. Zhang is with the School of Information Science and Technology, University of Science and Technology of China, China (e-mail: chizhang@ustc.edu.cn).

This work was supported in part by the US National Science Foundation under grants CNS-1117462 and CNS-0844972 (CAREER) and by the National Natural Science Foundation of China under grant 61202140.

The preliminary version of this work was published in INFOCOM'10 [1].

Digital Object Identifier 10.1109/JSAC.2013.SUP.0513024

[1]We use "user privacy" and "data privacy" as well as "aggregation integrity" and "data integrity" interchangeably in this paper.

node's original datum. In the second phase, the original datum of each node is aggregated in a privacy-preserving manner, in which users first exchange random shares of their data with selected peers and then submit mixed data to the aggregation server. The aggregation server can then verify the aggregation-result integrity using the aggregate commitment derived in the first phase.

- The second component VPA$^{\oplus}$ is designed for various non-additive aggregation functions like Max/Min, Median, Histogram, and Percentile through a unique combination of the binary search and verifiable privacy-preserving Count queries.

VPA is the first work of its kind as far as we know. The performance of VPA is thoroughly analyzed and evaluated with detailed simulations.

The rest of this paper is organized as follows. Section II reviews the related work. Section III gives the system and adversary models and the design objectives. Section IV and Section V present the solutions to additive and non-additive aggregation, respectively. Section VI evaluates the performance of VPA using extensive simulation results. This paper is finally concluded in Section VII.

II. Related Work

Although PC-USSs, also known as participatory or opportunistic sensing systems, have received extensive attention (e.g., [2]–[8]), there is relatively little work focusing on their security and privacy aspects. Kapadia *et al.* [7] surveyed the security and privacy challenges in opportunistic sensing systems. Cornelius *et al.* [5] presented the AnonySense architecture for anonymous tasking and reporting in people-centric sensing systems. AnonySense relies on a Mix network like Minimaster [9] to ensure user privacy, which we will not assume in our scheme. Ganti *et al.* [6] proposed PoolView for computing community statistics of time-series data in a privacy-preserving manner without considering aggregation-result integrity. More recently, Cristofaro and Soriente [10] proposed PEPSI to protect data and query privacy from unauthorized subscribers. None of these schemes could achieve the same objectives as our VPA.

Privacy-preserving aggregation in sensor networks has been extensively studied. The work [11]–[14] can support additive aggregation functions such as Sum and Average. GP^{2}S [15] can support both additive aggregation functions and non-additive ones such as Max/Min, Median, and Histogram at the sacrifice in data accuracy. The work [16] applies a particular class of encryption transformations to compute two aggregation functions, Average and "movement detection" specific to sensor networks. These schemes [11]–[16] do not address aggregation-result integrity and cannot be directly applied to PC-USSs due to different application scenarios.

There is also a big chunk of work on secure aggregation in sensor networks, see [17]–[23] for example. Such work ensures that aggregation results are not so different from the true values despite malicious intermediate aggregation nodes and does not address individual nodes' data privacy.

To the best of our knowledge, the work in [24] is the only one that simultaneously addresses data confidentiality and aggregation-result integrity. VPA differs from [24] significantly in following aspects. First, the scheme proposed in [24] targets histogram aggregates in traditional sensor networks with static topology, while VPA can support a large family of aggregates, including Sum, Average, Max/Min, Median, Histogram, and Percentile. Second, the scheme proposed in [24] can only detect ill-performed aggregation with some probability and protect users' data privacy against other users. In contrast, VPA can detect any false aggregation result with certainty and ensure user data confidentiality against both curious users and aggregation servers.

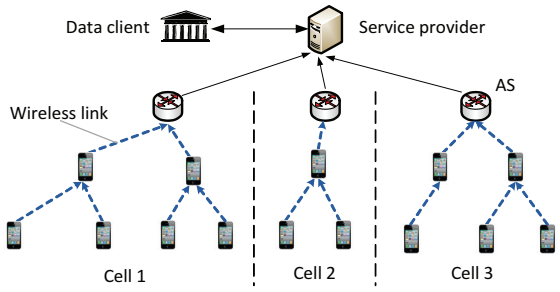

Fig. 1. The abstract architecture of a people-centric urban sensing system (PC-USS).

III. Models and Design Goals

In this section, we present the system and adversary models as well as our design goals.

A. System Model

There is no universally accepted model for a PC-USS. For ease of illustration, we assume an urban-sensing service provider which deploys a large-scale system similar to a metro-scale wireless mesh network [25], as shown in Fig. 1. Our solution can be easily extended to work with other system models such as cellular networks. The PC-USS features a high-speed wireless backbone consisting of M powerful aggregation servers (ASs for short) which also provide network access services for system nodes. Each AS is in charge of a certain region referred to as a *cell* and interacts with nodes therein. Here we use the term "node" to indicate a human who carries a portable device such as a smartphone and tablet. The devices have different communication and computation capabilities as well as various embedded sensors such as accelerator, digital compass, proximity sensors, and humidity sensors.

A node may participate in data sensing/sharing and also enjoy network access at will. To prevent fraudulent use of system resources and also provide basic privacy assurance to nodes, the system and nodes need mutually authenticate each other each time a node moves into a new cell. We assume a similar mutual authentication protocol as in [25]. Assume that an AS, denoted by $\mathcal{A}$, can simultaneously accommodate up to 2^λ users. After achieving mutual authentication with a node, say i, $\mathcal{A}$ assigns node i a secret key k_i, a temporal integer-valued ID ID_i (which is an unused one between $[0, 2^\lambda - 1]$), and also an ID-based private key K_i^{-1}. The pair ID_i/K_i^{-1}

will serve as the temporal public/private keys of node i which are valid only in $\mathcal{A}$'s cell.

In addition, we assume an efficient method for $\mathcal{A}$ to track node mobility in its cell. For example, node i need periodically notify $\mathcal{A}$ about its existence; otherwise, $\mathcal{A}$ would assume that i has left its cell and then reclaim ID_i for allocation to new nodes. In the latter case, $\mathcal{A}$ updates all the private keys of the remaining nodes in its cell using a single broadcast message with the approach in [26]. The use of such ID-based public/private keys will be illustrated soon. Note that the mutual authentication process is performed whenever a node enters a new cell.

About the communication capabilities, we assume that each node and the AS can directly communicate. In addition, each node can communicate with neighboring nodes through WiFi or Bluetooth interfaces, for which very efficient protocols are available such as in [27]. Moreover, each node can transmit to the AS in a multi-hop fashion through other nodes if necessary.

We consider the following scenario throughout. Assume that the service provider, on behalf of a data client, wants to get statistical aggregates of some personal data such as heart rates, blood pressure levels, glucose levels, weights, and moving speeds. A query will be sent to selected ASs which in turn broadcast the query to the nodes inside their respective cells. If some nodes have data satisfying the query, they will participate in aggregation if provided with privacy guarantees. The ASs can then aggregate the returned data and forward the aggregation result to the service provider. The service provider often need provide extra incentives such as credits to motivate participation in data aggregation, but the design of sound incentives is beyond our scope.

B. Adversary Model

This paper focuses on thwarting attacks on breaching nodes' data privacy as well as aggregation-result integrity. Other important issues such as DoS defenses [7] are beyond the scope of this paper.

We assume that ASs are trusted to follow aggregation operations for generating correct aggregation results, but they may be curious about individual user data. A curious AS may collude with other curious nodes to attempt deducing the sensing data of target nodes.

In contrast, a node could be curious, malicious, or both. Like a curious AS, a curious node is interested in discovering other nodes' data but faithfully follow aggregation operations, while a malicious node intends to make the AS derive false aggregation result. More specifically, a malicious node may launch two types of *false-data injection* attacks [20]. First, a malicious node may forge its own datum. Second, a malicious node may forge a false intermediate aggregation result that could significantly affect the final aggregation result. Most recent research [28]–[31] has shown that perhaps the only feasible defense against the former (i.e., ensuring the integrity of sensor readings from human-carried mobile devices) is to use some trusted hardware such as a Trusted Platform Module (TPM). We thus follow this line of research and assume that every participating mobile device has an embedded TPM. To keep the TPM cost as low as possible, we only require the TPM to have a minimal set of functionalities during aggregation, which include collecting sensor readings and generating a message authentication code (MAC). For this purpose, we assume that every TPM has a unique public/private key pair bound to the affiliated mobile device. After achieving mutual authentication with node i, the AS sends another secret key κ_i encrypted with the public key of node i's TPM. The TPM can then decrypt the ciphertext using its private key and store κ_i for later use. Our focus is thus on mitigating the forgery of intermediate aggregation results.

There might also be external eavesdroppers not involved in data aggregation. We shall use end-to-end encryption to withstand eavesdropping and focus on counteracting internal attackers hereafter, which includes both curious ASs and curious/malicious participating nodes.

C. Design Goals

Given the aforementioned adversary model, VPA is designed with the following objectives.

- *Aggregation accuracy*: VPA should output accurate aggregation results in the absence of malicious attacks.
- *Aggregation/data integrity*: any attempt of injecting false data should be detected with certainty.
- *Data/user privacy*: Each user's datum should be hidden from all the other parties with high probability.
- *Efficiency*: VPA should incur low communication and computation overhead.

IV. VPA$^+$: Verifiable Privacy-Preserving Additive Aggregation

This section presents VPA$^+$, a novel scheme to enable verifiable privacy-preserving additive aggregation. Our discussion focuses on a cell with AS $\mathcal{A}$ and a set of n nodes, denoted by $\mathcal{U}$. We will also use Sum aggregation as an example, based on which other additive aggregation functions such as Average and Variance [12] can be easily realized.

A. Overview

We observe that either of user privacy and aggregation integrity alone can be easily achieved if we ignore the other. On the one hand, if aggregation integrity is the only concern, a straightforward solution is to let each node submit its datum directly to $\mathcal{A}$ along with a message authentication code (MAC). The AS can then verify the authenticity of each datum and compute the correct sum. This naive approach, however, offers no data privacy to users. On the other hand, many existing techniques such as [13], [14] can realize privacy-preserving data aggregation, but a malicious node can launch the false-data injection attack without being detected.

Inspired by the above observation, we divide the whole aggregation process into two phases. In the first phase, each node submits to $\mathcal{A}$ a commitment, which is a homomorphic MAC of its datum and has a nice *one-way* property that $\mathcal{A}$ cannot deduce the corresponding datum. The homomorphic property of individual commitments enables $\mathcal{A}$ to compute an aggregate commitment corresponding to the Sum aggregate of all nodes' data. In the second phase, nodes perform privacy-preserving in-network data aggregation for $\mathcal{A}$ to derive the

Sum aggregate without disclosing any individual datum with overwhelming probability. Finally, $\mathcal{A}$ can verify the integrity of the Sum aggregate by using the aggregate commitment derived in the first phase. In what follows, we detail the design of VPA$^+$, which includes *aggregation initialization*, *commitment submission*, *privacy-preserving in-network aggregation*, and *aggregation verification*.

B. Aggregation Initialization

The AS $\mathcal{A}$ initializes the aggregation process by selecting a large prime p and a generator g of the group $\mathbb{Z}_p^* = \{1, \ldots, p-1\}$. The parameters p and g should ensure the computational hardness of the *discrete logarithm problem*, that is, given a random $y \in \mathbb{Z}_p^*$, it is computationally infeasible to find the unique integer $x \in [0, p-2]$ such that $g^x = y \mod p$. Assume that each node has reported to $\mathcal{A}$ what kinds of data it could generate when moving into $\mathcal{A}$'s cell. Let $\mathcal{U}$ denote the set of $n = |\mathcal{U}|$ users that $\mathcal{A}$ has selected and motivated to participate in data aggregation.[2] Finally, $\mathcal{A}$ broadcasts an aggregation request $\langle p, g, \mathcal{U}, r \rangle$, where r is a random nonce for message freshness. It is worth noting that the aggregation request can be sent as part of $\mathcal{A}$'s periodic service beacons and need be authenticated properly as in [25] to prevent attackers from sending fake aggregation requests, which we have ignored here for the focus of this paper. In addition, there can be various methods to transmit a condensed version of $\mathcal{U}$, which is also not discussed here for simplicity.

C. Commitment Submission

In this phase, each node $i \in \mathcal{U}$ submits to $\mathcal{A}$ a commitment, which is a homomorphic MAC of its datum d_i after appropriate expansion. A homomorphic MAC function $H(\cdot)$ has a desired property that the given the homomorphic MACs of two messages, say $H(m_1)$ and $H(m_2)$, anyone can derive $H(m_1+m_2)$ without knowing m_1 or m_2. VPA$^+$ uses a simple homomorphic MAC construction as follows,

$$H(m) = g^m \mod p,$$

where $m \in [0, p-2]$. It is easy to see that $H(\cdot)$ is homomorphic because $\forall m_1, m_2 \in [0, p-2]$,

$$H(m_1 + m_2) = g^{m_1+m_2} = H(m_1)H(m_2) \mod p.$$

Before generating the commitment, each node i first need expand its datum d_i to introduce sufficient randomness. Note that the data range in many PC-USS applications is usually limited. For instance, in a traffic monitoring application, the driving speed is between 0 and 100 miles. If node i directly submits $H(d_i)$ to $\mathcal{A}$, then $\mathcal{A}$ can deduce d_i by exhaustive search. To avoid this situation, each node i expands d_i by adding a random number. In particular, assume that each datum d_i is of l bits. Node i generates a random number r_i of ϕ bits known only to itself and computes

$$e_i = 2^{l+\lfloor \log_2 n \rfloor} \cdot r_i + d_i \ . \tag{1}$$

[2]How $\mathcal{A}$ selects $\mathcal{U}$ from candidate users and appropriately stimulate their participation is an orthogonal topic deserving independent investigation.

where ϕ is a system parameter determining the difficulty of exhaustive search. Alternatively, we can view e_i as the concatenation of r_i, $\lfloor \log_2 n \rfloor$ zeros, and d_i as follows

$$e_i = r_i, \overbrace{0, \ldots, 0}^{\lfloor \log_2 n \rfloor}, d_i \ .$$

The reason to separate r_i and d_i by $\lfloor \log_2 n \rfloor$ zeros can be explained as follows. If we perform Sum aggregation over all e_i, then we have

$$\begin{aligned} \sum_{i \in \mathcal{U}} e_i &= 2^{l+\lfloor \log_2 n \rfloor} \cdot \sum_{i \in \mathcal{U}} r_i + \sum_{i \in \mathcal{U}} d_i \\ &\leq 2^{l+\lfloor \log_2 n \rfloor} \cdot \sum_{i \in \mathcal{U}} r_i + n(2^l - 1) \\ &< 2^{l+\lfloor \log_2 n \rfloor} \cdot \sum_{i \in \mathcal{U}} r_i + 2^{l+\lfloor \log_2 n \rfloor} \ . \end{aligned}$$

It follows that

$$\sum_{i \in \mathcal{U}} d_i = \sum_{i \in \mathcal{U}} e_i \mod 2^{l+\lfloor \log_2 n \rfloor} \ . \tag{2}$$

This property will be used later by $\mathcal{A}$ to derive the correct aggregation result without knowing $\{r_i\}_{i \in \mathcal{U}}$.

To prevent malicious nodes from submitting arbitrary data, we require that d_i and e_i be generated and authenticated by node i's TPM. Recall that $\mathcal{A}$ has assigned a secret key k_i to node i and another secret key κ_i to node i's TPM after mutual authentication (see Section III-A). Node i submits to $\mathcal{A}$ the following message.

$$i \to \mathcal{A} : i, \langle H(e_i), h(\kappa_i || H(e_i)) \rangle_{k_i} \ ,$$

where $h(\cdot)$ denotes a hash function, and $\langle \cdot \rangle_*$ denotes a symmetric-key encryption operation using the subscript key.

On receiving the message, the AS locates k_i and κ_i using node ID i, uses k_i to decrypt the message, and then verifies $h(\kappa_i || H(e_i))$ using κ_i. If the verification succeeds, $\mathcal{A}$ considers $H(e_i)$ authentic and drops it otherwise. If $\{H(e_i)\}_{i=1}^n$ are all authentic, the AS proceeds to derive the aggregate commitment corresponding to $\sum_{i=1}^n e_i$ by computing

$$\begin{aligned} H(\sum_{i \in \mathcal{U}} e_i) &= \prod_{i \in \mathcal{U}} H(e_i) \mod p \\ &= \prod_{i \in \mathcal{U}} g^{e_i} \mod p \\ &= g^{\sum_{i \in \mathcal{U}} e_i} \mod p \ . \end{aligned}$$

D. Privacy-Preserving In-Network Data Aggregation

In this phase, nodes jointly perform in-network aggregation over their expanded data without disclosing them. This phase requires the establishment of an on-demand temporary aggregation tree. In particular, the AS $\mathcal{A}$ broadcasts an *aggregation-tree formation* request, which specifies any node, say $v \in \mathcal{U}$, as the root of the aggregation tree. On receiving the request, node v rebroadcasts it via its Bluetooth or WiFi interface, depending on the particular method (e.g., [27]) it uses to communicate with neighboring nodes. Upon receiving the request for the first time, each node further rebroadcasts it and records the parent node from which this request came from. In this way,

an aggregation tree is formed and rooted at node v which can directly communicate with $\mathcal{A}$.

In what follows, we present two techniques for privacy-preserving in-network Sum aggregation over all expanded data with different user-privacy guarantees and communication overhead. To facilitate presentation, we define node i's *aggregation neighbors* as i's neighboring nodes on the aggregation tree, denoted by $\mathcal{T}_i$.

*1) Method 1: Data Perturbation (*DP*):* In this method, each node i perturbs its expanded datum e_i before actual aggregation. Since e_i is of $l + \lfloor \log_2 n \rfloor + \phi$ bits, we have

$$\sum_{i \in \mathcal{U}} e_i \leq n \cdot (2^{l+\lfloor \log_2 n \rfloor + \phi} - 1) < 2^{l+2\lfloor \log_2 n \rfloor + \phi},$$

i.e., that $\sum_{i \in \mathcal{U}} e_i$ is at most of $l + 2\lfloor \log_2 n \rfloor + \phi$ bits.

Denote by $h_1(\cdot)$ a good hash function of $l + 2\lfloor \log_2 n \rfloor + \phi$ bits. Each node i generates a perturbed datum α_i by computing

$$\alpha_i = h_1(k_i || r) + e_i \mod 2^{l+2\lfloor \log_2 n \rfloor + \phi}, \quad (3)$$

where k_i is the secret key shared between node i and the AS and r is the nonce broadcasted by $\mathcal{A}$.

Each node then performs in-network aggregation over its perturbed datum by adding it to the values received from its children on the aggregation tree and then transmitting the result to its parent. Finally, the AS $\mathcal{A}$ can obtain $\sum_{i \in \mathcal{U}} \alpha_i$ by summing the values received from the root of the aggregation tree, i.e., node v. Since $\mathcal{A}$ knows k_i for each $i \in \mathcal{U}$, it can compute all $h_1(k_i||r)$ and derive $\sum_{i \in \mathcal{U}} e_i$ by computing

$$\sum_{i \in \mathcal{U}} e_i = \sum_{i \in \mathcal{U}} \alpha_i - \sum_{i \in \mathcal{U}} h_1(k_i||r) \mod 2^{l+2\lfloor \log_2 n \rfloor + \phi}. \quad (4)$$

Since e_i is completely concealed by $h_1(k_i||r)$, which is only known by $\mathcal{A}$, other curious nodes, e.g., node i's neighbors on the aggregation tree, cannot derive e_i by monitoring i's incoming and outgoing transmission. Unfortunately, node i's data privacy can still be breached if $\mathcal{A}$ colludes with node i's aggregation neighbors.

2) Method 2: Peer-to-Peer Slicing and Mixing: To defend against $\mathcal{A}$ colluding with other curious nodes, we further propose another approach based on peer-to-peer data slicing and mixing. In this approach, before participating in in-network aggregation, each node i randomly divides its expanded datum e_i into multiple slices and mixes them with those from selected peers, such that data privacy can be preserved without affecting the correctness of the final aggregation result.

Specifically, before answering the query, each node i slices e_i into $t+1$ random slices $\{s_{i,j}\}_{j=1}^{t+1}$ with $t \leq n-1$, such that

$$e_i = \sum_{j=1}^{t+1} s_{i,j} \mod 2^{l+2\lfloor \log_2 n \rfloor + \phi}.$$

Then node i keeps $s_{i,t+1}$ to itself while sending each other slice to a unique peer called a *cover node*. Next, each node i adds the slices received from other nodes to its remained slice $s_{i,t+1}$ and conducts in-network aggregation as in Method 1. Finally, $\mathcal{A}$ adds up all the received values. It is easy to see that the result is exactly the Sum aggregate of interest.

This slicing technique shares the similar idea as PDA [13], while our application scenario is totally different. In particular, PDA is designed for sensor networks with relatively static network topology, where all the nodes know each other and have pairwise shared keys whereby to encrypt/decrypt data slices transmitted from any node to its chosen cover nodes. Such assumptions no longer holds in our target scenarios, where nodes in a cell are dynamically changing. Since the nodes do not know each other beforehand, they have no pre-shared keys for end-to-end encryption. This significant difference necessitates novel cover-selection strategies. In what follows, we detail two cover-selection approaches that specially tailored for our target scenario.

*a) **Random cover selection (**RCS**) :*** In this approach, each node randomly chooses t cover nodes from $\mathcal{U}$ and sends a data slice to each of them. The challenge is how a node can establish a shared key with each of its cover nodes for end-to-end encryption of its shares. VPA$^+$ uses the following method. Consider node i as an example with data e_i to share. It first slices e_i into $\{s_{i,j}\}_{j=1}^{t+1}$ and then randomly chooses a set of t nodes from $\mathcal{U}$ as its cover nodes, denoted by $\mathcal{C}_i \subseteq \mathcal{U}$. For any cover node $j \in \mathcal{C}_i$, node i computes a shared key $k_{i,j}$ based on its temporal public/private keys ID_i/K_i^{-1} and ID_j by using the method in our previous work [25] and then sends an encrypted unique slice $s_{i,\tau_{i,j}}$ to node j as follows.

$$i \rightarrow j : ID_i, \langle s_{i,\tau_{i,j}}, h(s_{i,\tau_{i,j}}) \rangle_{k_{i,j}}$$

Since the route to j may be unknown, the packet transmission is normally preceded by a route discovery process using protocols like AODV [32]. On receiving the message, node j can derive the same key $k_{i,j}$ using its temporal public/private keys ID_j/K_j^{-1} and ID_i according to [25] and then decrypts the packet to get $s_{i,\tau_{i,j}}$. Node i does this for all its cover nodes, and so does every other node in $\mathcal{U}$.

Each node waits sufficiently long for receiving all the slices from other nodes choosing it as cover. Let $\mathcal{S}_i \subset \mathcal{U}$ denote those selecting i as a cover node. Each node i computes its share as

$$\beta_i = s_{i,t+1} + \sum_{j \in \mathcal{S}_i} s_{j,\tau_{j,i}} \mod 2^{l+2\lfloor \log_2 n \rfloor + \phi}. \quad (5)$$

Finally, all the nodes perform in-network aggregation over there shares as in Method 1 so that the AS $\mathcal{A}$ finally receives $\sum_{i \in \mathcal{U}} \beta_i$ which equals $\sum_{i \in \mathcal{U}} e_i$.

*b) **μ-hop cover selection (**μCS**):*** Random cover selection may not be efficient because cover nodes are randomly chosen regardless of their locations. An on-demand energy-consuming route discovery process is thus often incurred to find a route to a chosen cover node multi-hop away. We note that it is unnecessary for each node i to predetermine the slices $\{s_{i,j}\}_{j=1}^{t+1}$ and send each of them to a cover node. Instead, node i can broadcast a random seed within its μ-hop neighborhood, in which every node is chosen as a cover node and can compute a slice using their shared key.

Specifically, in μ-hop cover selection, each node i initiates the slicing process by broadcasting a slicing request with a random seed r_i and a TTL value set to μ. Upon receiving a request with a TTL larger than one, each node further broadcasts it after decreasing the TTL by one. A node should

only process the first copy of the same request which may be heard multiple times. In addition, each node memorizes the parent node from which this request came from. In this way, a routing tree of depth μ is formed and rooted at node i. When a node receives a request with the TTL value equal to one, the node should send a slicing response to its parent node which in turn forwards the response via the routing tree back to node i after appending its ID.

Each node waits sufficiently long and then updates its share as follows. Consider node i as an example. Suppose that node i has received slicing responses from nodes $\mathcal{C}_i$ and slicing requests from nodes $\mathcal{S}_i$ (i.e., the set of nodes choosing i as covers). Node i derives a shared key $k_{i,j}$ for each $j \in \mathcal{C}_i \bigcup \mathcal{S}_i$ according to [25] and updates its share by computing

$$\begin{aligned}\beta_i = e_i &- \sum_{j\in\mathcal{C}_i} h_1(r_i||k_{i,j}) \\ &+ \sum_{j\in\mathcal{S}_i} h_1(r_j||k_{i,j}) \mod 2^{l+2\lfloor\log_2 n\rfloor+\phi} .\end{aligned} \tag{6}$$

Finally, all the nodes perform in-network aggregation over their shares so that the AS $\mathcal{A}$ finally obtains $\sum_{i\in\mathcal{U}} \beta_i$ which equals $\sum_{i\in\mathcal{U}} e_i$.

Unlike in random cover selection, the number of cover nodes in μ-hop cover selection is a random variable which cannot be determined before the process is completed. Intuitively, the larger μ, the more cover nodes discovered, the higher privacy and the communication cost, and vice versa.

E. Aggregation-Result Verification

After in-network aggregation via Method 1 or 2, the AS obtain $\sum_{i\in\mathcal{U}} e_i$. It first verifies its integrity by checking if

$$g^{\sum_{i\in\mathcal{U}} e_i} = \prod_{i\in\mathcal{U}} H(e_i) \mod p.$$

If so, $\mathcal{A}$ considers $\sum_{i\in\mathcal{U}} e_i$ authentic and proceeds to derive $\sum_{i\in\mathcal{U}} d_i$ by computing

$$\sum_{i\in\mathcal{U}} d_i = \sum_{i\in\mathcal{U}} e_i \mod 2^l ,$$

which should hold according to Eq. (2).

F. Performance Analysis

Now we analyze the performance of VPA$^+$ with regard to its aggregation-integrity provision, data-privacy guarantee, and the associated overhead.

1) Aggregation Integrity: We first have the following theorem regarding the aggregation integrity of VPA$^+$.

Theorem 1: Assume that each node's datum is generated and authenticated by TPM and that $p > 2^{l+2\lfloor\log_2 n\rfloor+\phi}$. VPA$^+$ allows the AS to detect any false-data injection attack.

Proof: Assume that the AS receives $\{H(e_i) : i \in \mathcal{U}\}$ during the commitment submission phase, whereby it derives $H(\sum_{i\in\mathcal{U}} e_i) = g^{\sum_{i\in\mathcal{U}}} \mod p$. Suppose that some malicious nodes injected false data during the data-aggregation phase so that the AS receives $e' \neq \sum_{i\in\mathcal{U}} e_i$. The AS cannot detect the false-data injection attack if and only if

$$g^{e'} = g^{\sum_{i\in\mathcal{U}} e_i} \mod p .$$

However, for any $y \in \mathbb{Z}_p^*$, there is a unique $x \in [0, p-2]$ such that $g^x = y \mod p$. Since $p > 2^{l+2\lfloor\log_2 n\rfloor+\phi} > \sum_{i\in\mathcal{U}} e_i$, there is no $e' \neq \sum_{i\in\mathcal{U}} e_i$ can satisfy the above equation. It is thus impossible for the adversary to inject false data without being detected under VPA$^+$. ■

2) Data Privacy: To evaluate the data-privacy provision of VPA$^+$, we define *exposure probability*, denoted by P_{exp}, as the probability that a node i's data d_i is disclosed during aggregation. To enable quantitative analysis, we assume that each node has N_{tree} aggregation neighbors on average. We also assume that there are M_c out of M curious ASs and n_c out of n curious nodes.

We then have the following theorems regarding the exposure probability under VPA$^+$.

Theorem 2: The exposure probability under DP *is given by*

$$P_{exp} = \frac{M_c}{M} \cdot \frac{\binom{n-n_c}{n_c-N_{tree}}}{\binom{n}{n_c}} . \tag{7}$$

We given the proof in [33].

Theorem 3: The exposure probabilities under RCS *and* μCS *are bounded by*

$$P_{exp} \leq \frac{\binom{n-n_c}{n_c-w}}{\binom{n}{n_c}} , \tag{8}$$

where

$$w = \begin{cases} \max(N_{tree}, t) & \textit{for RCS,} \\ \max(N_{tree}, \sum_{x=1}^{\mu} N_x) & \textit{for } \mu\text{CS,} \end{cases} \tag{9}$$

is the minimum number of nodes colluding with $\mathcal{A}$.
We give the proof in [33].

3) Overhead Analysis: Now we analyze the computation and communication overhead incurred by VPA$^+$.

For computation overhead, each node need perform one exponentiation to generate one commitment of its data. In addition, each node i need compute the shared key $k_{i,j}$ for each node $j \in \mathcal{C}_i \bigcup \mathcal{S}_i$ under RCS and μCS.

We assume that the average distance two random chosen nodes is L hops. Also denote by l_{tree}, l_{seed}, l_{hmac}, l_{h} the length of a aggregation tree formation request, a slicing request in μCS, a homomorphic MAC, and $h(\cdot)$, respectively. We then have the following theorem regarding the communication overhead incurred by VPA$^+$.

Theorem 4: The communication overhead incurred by VPA$^+$ in bits is given by

$$\mathsf{T}_{VPA^+} = nl_{tree} + \mathsf{T}_{commit} + \mathsf{T}_{agg} , \tag{10}$$

where

$$\mathsf{T}_{commit} = n(\lambda + l_{hmac} + l_h) \tag{11}$$

is the overhead incurred by transmitting commitments to $\mathcal{A}$, and

$$\mathsf{T}_{agg} = \begin{cases} nl_{data} & \textit{for DP,} \\ nt(nl_{req} + L(l_{rsp} + l_{data} + \lambda)) & \\ \quad + nl_{data} & \textit{for RCS,} \\ n((1 + \sum_{x=1}^{\mu-1} N_x(\lambda + l_{seed}) & \\ \quad + \sum_{x=1}^{\mu} N_x\lambda) + nl_{data} & \textit{for } \mu\text{CS,} \end{cases} \tag{12}$$

is the overhead incurred by in-network aggregation, and $l_{data} = l + 2\lfloor\log_2 n\rfloor + \phi$.
We give the proof in [33].

V. VPA$^{\oplus}$: Verifiable Privacy-Preserving Non-additive Aggregation

VPA^{+} cannot be directly applied to non-additive aggregation functions such as Max/Min, Median, Percentile, and Histogram, which have wide applications in practice. In this section, we propose VPA$^{\oplus}$ as an extension of VPA^{+} to support non-additive aggregation.

A. Basic Idea

Our key observation is that all the above non-additive aggregation functions are closely related to Count aggregation that ask for the number of nodes whose values are above, below, or equal to a certain value. In particular, let $\mathsf{Count}[Q]$ be the number of nodes with data satisfying the condition Q. Also denote by $d_{\max}, d_{\min}, d_{\text{med}}, d_{\sigma-\text{per}}$, the Max, Min, Median, and σ-percentile of a data set, respectively. It is easy to see that the following conditions hold.

- *Max*:
$$\begin{cases} \mathsf{Count}[d > d_{\max}] = 0, \\ \mathsf{Count}[d = d_{\max}] > 0\ . \end{cases} \tag{13}$$
- *Min*:
$$\begin{cases} \mathsf{Count}[d < d_{\min}] = 0, \\ \mathsf{Count}[d = d_{\min}] > 0\ . \end{cases} \tag{14}$$
- *Median*:
  - If n is odd, then
$$\begin{cases} \mathsf{Count}[d \le d_{\text{med}}] \ge \lfloor n/2 \rfloor, \\ \mathsf{Count}[d \ge d_{\text{med}}] \ge \lfloor n/2 \rfloor\ . \end{cases} \tag{15}$$
  - If n is even, then there exists $i, j \in \mathcal{U}$, such that $d_i \le d_j$ and
$$\begin{cases} \mathsf{Count}[d \le d_i] \ge n/2, \\ \mathsf{Count}[d < d_i] < n/2, \\ \mathsf{Count}[d \ge d_j] \ge n/2, \\ \mathsf{Count}[d > d_j] < n/2\ , \end{cases} \tag{16}$$
    and $d_{\text{med}} = (d_i + d_j)/2$.
- *σ-percentile*: we only show the simplest case here
$$\begin{cases} \mathsf{Count}[d \le d_{\sigma-\text{per}}] \ge \lfloor \sigma n/100 \rfloor, \\ \mathsf{Count}[d \ge d_{\sigma-\text{per}}] \ge \lfloor (100-\sigma) n/100 \rfloor\ . \end{cases} \tag{17}$$

Conversely, if we can find d^* such that the conditions in Eq. (13) (respectively, (14), (15), (16) (17)) hold, then we have $d^* = d_{\max}$ (respectively, $d_{\min}$, d_{med}, $d_{\sigma-\text{per}}$). Since Count is an additive aggregation function, it can be realized by VPA^{+}. Built on the above observation, VPA$^{\oplus}$ combines VPA^{+} with binary search to realize non-additive aggregation functions through a series of verifiable privacy-preserving Count aggregations.

B. Scheme Description

Given a non-additive aggregation request, $\mathcal{A}$ transforms it into a series of Count queries with conditions $Q_1, Q_2, \ldots$, until the desired d^* is found, where Q_x is determined by the result of the previous Count query with condition Q_{x-1}. Each Count query Q_x asks how many nodes possess data above, below, or equal to a threshold, called a *count index*. Each node i with datum d_i satisfying condition Q_x gives an answer "yes", or "no" otherwise, by a single bit of value one or zero, respectively. The answers are then aggregated via VPA^{+} to let $\mathcal{A}$ get $\mathsf{Count}(Q_x)$ with both user-privacy and aggregation-integrity guarantees.

Below we brief how to realize privacy-preserving Max/Min, Median, Histogram, and Percentile aggregation queries under the assumption that each data value d_i is an integer between $[0, 2^l - 1]$. It is easy to extend our technique to other non-additive aggregation functions.

1) Max/Min: Since Min is opposite to Max, we just illustrate Max for brevity. Given a Max aggregation request, $\mathcal{A}$ first issues a Count query with $Q_1 = [d \ge 2^{l-1}]$ and then aggregates the received data via VPA^{+} to get the number of "yes" answers, denoted by θ_1. If $\theta_1 \ge 1$, the maximum value should be in $[2^{l-1}, 2^l - 1]$, so $\mathcal{A}$ will send a new Count query with $Q_2 = [d \ge 2^{l-1} + 2^{l-2}]$; otherwise, the maximum value should be in $[0, 2^{l-1} - 1]$, so $\mathcal{A}$ will send a new Count query with $Q_2 = [d > 2^{l-2}]$. The *suspicion range* in which the maximum value is located is reduced by half for each additional Count query. This process continues until the suspicion range is reduced to one, in which case the last count index is exactly the maximum value, and the last query result equals the number of nodes with the maximum value.

2) Median/Percentile: Since Median is a special case of Percentile, we illustrate the former for simplicity, which can be easily extended to the latter. A median value is described as the number separating the higher half of a sample, a population, or a probability distribution, from the lower half. Median aggregation can be realized in a similar fashion as Max. Here we present the case for n being odd for simplify, while the case of n being even can be realized accordingly.

Given a Median aggregation request, $\mathcal{A}$ first issues a Count query with $Q_1 = [d \ge 2^{l-1}]$ and obtains θ_1 via VPA^{+}. If $\theta_1 \ge (n+1)/2$, $\mathcal{A}$ sends the second Count query with $Q_2 = [d \ge 2^{l-1} + 2^{l-2}]$; otherwise, $\mathcal{A}$ sends the next query with $Q_2 = [d \ge 2^{l-2}]$. This process continues until the suspicion range of d_{med} is one, which takes total l queries. Suppose that the last two queries are $Q_{l-1} = [d \ge q_{l-1}]$ and $Q_l = [d \ge q_l]$ whereby $\mathcal{A}$ receives θ_{l-1} and θ_l, respectively. It follows that q_{l-1} and q_l differ by one. There are four cases.

- Case 1: if $q_{l-1} < q_l$ and $\theta_l \ge \lfloor n/2 \rfloor$, we have $d_{\text{med}} = q_l$ due to the following reasons. First, we must have $\theta_{l-1} < \lfloor n/2 \rfloor$, as otherwise $d_{\text{med}} \le q_l - 1$, and q_l should not be queried. Second, there must be a query $Q_x = [d \ge q_l + 1]$ with $x \in [1, l-2]$, due to the property of binary search. Third, it must hold that $\theta_x < \lfloor n/2 \rfloor$, as otherwise $d_{\text{med}} > q_l + 1$ and neither q_{l-1} nor q_l should be queried.
- Case 2: if $q_{l-1} > q_l$ and $\theta_l \ge \lfloor n/2 \rfloor$, we have $d_{\text{med}} = q_l$.
- Case 3: if $q_{l-1} < q_l$ and $\theta_l < \lfloor n/2 \rfloor$, $d_{\text{med}} = q_l + 1$.
- Case 4: if $q_{l-1} > q_l$ and $\theta_l < \lfloor n/2 \rfloor$, $d_{\text{med}} = q_l + 1$.

The reasoning for Cases 2∼4 are similar to that of Case 1.

3) Histogram: In statistics, a histogram is a graphical display of tabulated frequencies, shown as bars, and shows the proportion of cases falling into each of several categories. Using Count query to realize Histogram is straightforward. In particular, given a Histogram aggregation request, $\mathcal{A}$ partitions the data range $[0, 2^l - 1]$ into a certain number of consecutive, non-overlapping intervals according to the aggregation request. It then sends a Count query for each interval, and the corresponding query result will equal the number of nodes with data in that interval.

C. Performance Analysis

Since VPA$^{\oplus}$ is built upon VPA+, it can also ensure perfect aggregation integrity. We thus focus on analyzing the user-privacy provision and overhead of VPA$^{\oplus}$.

1) Data Privacy: The exposure probability P_{exp} used to analyze the performance of VPA^{+} can no longer precisely measure the privacy provision of non-additive aggregation. For example, even if the answer of node i to a Count query Q_x is disclosed, the adversary can only narrow down the search of d_i to a certain range instead of precisely determining d_i. Assume that the adversary knows that d_j is in a range of length ϵ after the whole query process. It is clear that the ratio $\rho = \epsilon/2^l$ can be used to analyze the privacy performance of the non-additive aggregation process: the larger ρ, the higher level of privacy provision, and vice versa.

In particular, when $\rho = 1$, the adversary has no clue about what d_i is; when $\rho = 2^{-l}$, i.e., $\epsilon = 1$, the adversary has precisely located d_i. We call ρ the *suspicion ratio* of d_i hereafter. Without loss of generality, we use Max as an example to evaluate the performance of the non-additive aggregation process. The studies about other non-additive aggregation functions can be conducted similarly. Before proceeding, we want to mention that the Max/Min aggregation functions naturally disclose some information: any user's data will be smaller or equal to $d_{\max}$ and larger or equal to $d_{\min}$. No scheme can prevent this kind of privacy breach which is due to the aggregate functions themselves. In the following, we will ignore such natural privacy breach and focus on the loss of privacy occurring in the query process.

We make the following assumptions for analytical tractability. We assume that the aggregation tree is static for the entire sequence of l Count queries. For clarity, we consider a special case where the maximum value $d_{max} = 2^l - 1$. The similar process can be used to analyze the more general case that d_{max} may be any value in $[0, 2^l - 1]$, which is not reported here due to space limitations. We then have the following theorems regarding the expected suspicion ratio of VPA$^{\oplus}$.

Theorem 5: Assuming that $d_{max} = 2^l - 1$, the expected suspicion ratio of VPA$^{\oplus}$ under DP *or* μCS *is given by*

$$\mathsf{E}[\rho] = 1 - P_{exp} + (2^{-2l} + \sum_{x=1}^{l} 2^{-2x}) P_{exp} \ , \tag{18}$$

where P_{exp} is given in Eq. (7) for DP *and Eq. (8) for* μCS. We give the proof in [33].

TABLE I
DEFAULT SIMULATION SETTINGS

Para.	Val.	Para.	Val.	Para.	Val.	Para.	Val.
M	10	M_c	5	n	200	n_c	50
λ	8	ϕ	160	μ	1	t	5
N_1	20.9	N_2	39.3	N_3	48.1	N_4	48.2
l	10	L	3.39	l_{tree}	160	l_{seed}	160
l_{req}	160	l_{rsp}	160	l_{hmac}	1024	N_{tree}	1.86

Theorem 6: Assuming that $d_{max} = 2^l - 1$, the expected suspicion ratio under of VPA$^{\oplus}$ under RCS *is given by*

$$\mathsf{E}[\rho] = \sum_{x=0}^{l-1} 2^{-x-1} \mathsf{E}[\rho_x] + 2^{-l} \mathsf{E}[\rho_l] \ , \tag{19}$$

where

$$\mathsf{E}[\rho_x] = \sum_{k_1=0}^{x} Pr(y_e = k_1) \sum_{k_2=x+1}^{l+1} Pr(n_e = k_2) \rho[y_e, n_e] \ , \tag{20}$$

$$Pr(y_e = k) = \begin{cases} (1 - P_{exp})^x & \text{if } k = 0, \\ P_{exp}(1 - P_{exp})^{k-1} & \text{if } 1 \le k \le x \ , \end{cases} \tag{21}$$

$$Pr(n_e = k) = \begin{cases} P_{exp}(1 - P_{exp})^{k-x-1} & \text{if } x + 1 \le k \le l, \\ (1 - P_{exp})^{l-x} & \text{if } k = l + 1 \ , \end{cases} \tag{22}$$

$$\rho[y_e, n_e] = \begin{cases} 2^{-y_e} - 2^{-n_e} & \text{if } y_e + 1 \le n_e \le l, \\ 2^{-y_e} & \text{if } n_e = l + 1 \ , \end{cases} \tag{23}$$

P_{exp} is given in Eq. (8).
We give the proof in [33].

2) Overhead Analysis: VPA$^{\oplus}$ differs from VPA+ mainly in the communication overhead. Since it takes l queries to complete the aggregation process, each of which incurs communication overhead of $\mathsf{T}_{\text{VPA}^+}$, we thus have

$$\mathsf{T}_{\text{VPA}^{\oplus}} = l \cdot \mathsf{T}_{\text{VPA}^+} \ ,$$

where $\mathsf{T}_{\text{VPA}^+}$ is given in Eq. (10).

VI. Performance Evaluation

In this section, we evaluate VPA+ and VPA$^{\oplus}$ using extensive simulations.

A. Simulation Setting

We simulate 10 cells of 1 km^2, each with an AS located at the center and 200 nodes randomly distributed within the cell. The transmission range of each node is 200m. This gives the average hop distance between two random nodes $L = 3.39$.

For our purpose, the simulation code is written in C++ and each data point represents an average of 50 simulation runs with different random seeds. Table I summarizes the default setting used in our simulation if not mentioned otherwise.

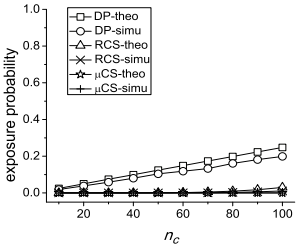

(a) exposure probability

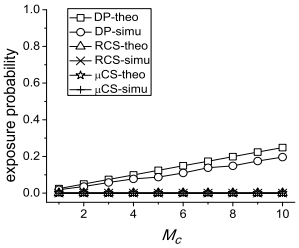

(b) exposure probability

Fig. 2. Impact of n_c and M_c.

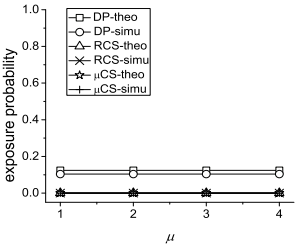

(a) exposure probability

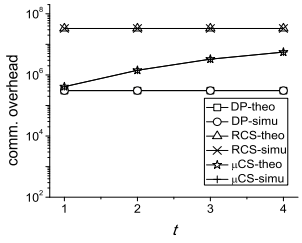

(b) exposure probability

Fig. 4. Impact of μ.

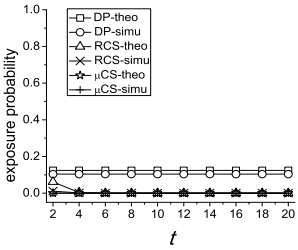

(a) exposure probability

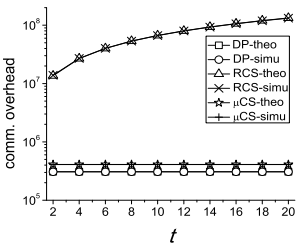

(b) communication overhead

Fig. 3. Impact of t, the number of cover nodes on RCS.

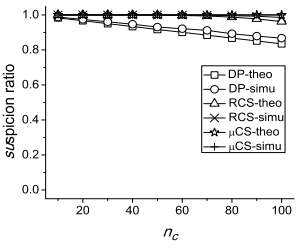

(a) suspicion ratio

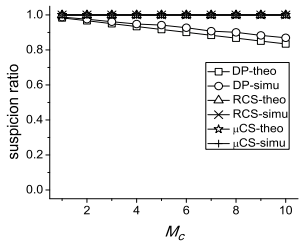

(b) suspicion ratio

Fig. 5. Impact of n_c and M_c on suspicion ratio.

B. Evaluation of VPA$^+$

Fig. 2(a) shows both the theoretical and simulation results of the exposure probabilities of DP, RCS and μCS varying with n_c, the number of curious nodes. We can see that the exposure probabilities of all three schemes decrease as n_c increases. Among three schemes, DP has the highest exposure probability, followed by RCS and μCS. The reason is that on average, each node has only less than two neighbors on the aggregation tree (i.e., a spanning tree), making it easier for the adversary to compromise (or collude with) all the neighbors of a target node under DP. In contrast, it is much more difficult to compromise all the cover nodes under RCS and μCS. In addition, we can see that the P_{exp} of DP obtained via theoretical analysis is slightly higher than that obtained by simulations. The reason is that we round N_{tree} to $\lfloor N_{\text{tree}} \rfloor$ when computing $\binom{n-n_c}{n_c-N_{\text{tree}}}$ in Eq. (7), leading to higher P_{exp}.

Fig. 2(b) shows the impact of M_c, the number of curious ASs on the exposure probability of DP. Since curious ASs has no impact on RCS and μCS, there exposure probabilities are shown only for references. We can see that the exposure probability of DP increases linearly with the number of curious AS increases, which is expected.

Fig. 3(a) shows the impact of t, the number of cover nodes, on the exposure probability of RCS, where the P_{exp}s of DP and μCS are shown only for reference. We can see that the P_{exp} of RCS decreases as t increases, and quickly drops to zero when $t > 4$. The reason is that the probability of all the t cover nodes being compromised decreases exponentially as t increases.

Fig. 3(b) shows the communication overhead of RCS varying with t. We can see that under the default settings, RCS incurs significantly higher communication overhead than that of DP and RCS. This is anticipated since finding a cover node under RCS requires an AODV-like route discovery that involves a network-wide flooding.

Fig. 4 shows the impact of μ on the exposure probability and communication overhead of μCS, where the results of DP and RCS are only shown for reference. We can see from Fig. 4(a) that the exposure probability of μCS is not much affected by μ because P_{exp} is already close to zero when $\mu = 1$. In addition, We can see that the communication overhead of μCS increases moderately as μ increases, which is of no surprise.

C. Evaluation of VPA$^\oplus$

Fig. 5(a) shows the suspicion ratios of DP, RCS and μCS, varying with n_c. We can see that the suspicion ratios of all three schemes decrease as n_c increases. The reason is that the higher n_c, the lower P_{exp}, and the lower suspicion ratio, and vice versa. In addition, under the default setting, μCS has the highest suspicion ratio, followed by that of RCS and DP.

Fig. 5(b) shows the impact of M_c on the suspicion ratio of DP, where the performance of RCS and μCS are only shown for reference. We cans see that the larger M_c, the lower suspicion ratio, and vice versa, which is easy to understand.

Fig. 6 shows the impact of l on the suspicion ratio and communication overhead of VPA$^\oplus$. We can see from Fig. 6(a) that the change in data range has negligible impact on the suspicion ratio of VPA$^\oplus$. The reason is that the suspicion ratio is determined by the last disclosed yes answer and the first disclosed no answer. Under the default setting, DP has the lowest suspicion ratio due to its highest P_{exp} among the three schemes (cf. Fig. 2(a)), while the suspicion ratios of both RCS and μCS are close to one. In addition, we can see from Fig. 6(b) that the communication overhead of VPA$^\oplus$ increases linearly as l increases, as it takes l Count queries to locate the desired aggregate.

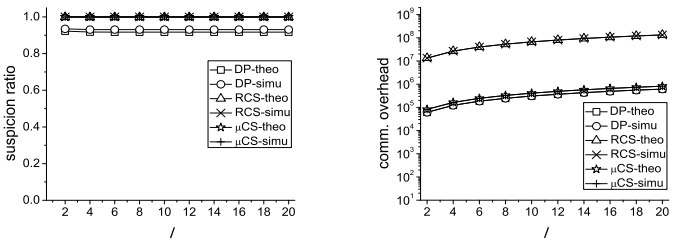

(a) suspicion ratio (b) communication overhead

Fig. 6. Impact of l.

D. Discussion

We summarize the evaluation results as follows.

- All three variants of VPA$^+$ (i.e., DP, RCS, and μCS) can ensure aggregation integrity by detecting any false-data injection attempt.
- DP can provide user/data privacy with high probability while incurring the minimum communication overhead.
- RCS can provide user/data privacy against curious ASs with overwhelming probability while incurring the highest communication overhead.
- μCS can provide user/data privacy against curious ASs with overwhelming probability while incurring relatively low communication overhead.
- Built upon VPA$^+$ and binary search, VPA$^\oplus$ can ensure both aggregation integrity and user/data privacy with communication overhead linear to the bit length of data.

In practice, μCS and the resulting VPA$^\oplus$ may be the best choices whose performance can be adjusted as needed.

VII. Conclusion

In this paper, we have presented the design and evaluation of VPA, a novel peer-to-peer approach to verifiable privacy-preserving aggregation for people-centric urban sensing systems. VPA can support a wide range of additive and non-additive aggregation functions with strong user-privacy and aggregation-integrity guarantees. The high efficacy and efficiency of VPA are confirmed by thorough theoretical analysis and simulation results.

References

[1] J. Shi, R. Zhang, Y. Liu, and Y. Zhang, "PriSense: Privacy-preserving data aggregation in people-centric urban sensing systems," in *Proc. IEEE INFOCOM*, Mar. 2010.

[2] A. T. Campbell, S. B. Eisenman, N. D. Lane, E. Miluzzo, and R. A. Peterson, "People-centric urban sensing," in *Proc. ICST WICON*, Aug. 2006.

[3] B. Hull, V. Bychkovsky, Y. Zhang, K. Chen, M. Goraczko, A. Miu, E. Shih, H. Balakrishnan, and S. Madden, "CarTel: A distributed mobile sensor computing system," in *Proc. ACM SENSYS*, Oct. 2006, pp. 125–138.

[4] 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.

[5] C. Cornelius, A. Kapadia, D. Kotz, D. Peebles, M. Shin, and N. Triandopoulos, "Anonysense: Privacy-aware people-centric sensing," in *Proc. ACM MobiSys*, June 2008, pp. 211–224.

[6] R. K. Ganti, N. Pham, Y.-E. Tsai, and T. F. Abdelzaher, "Poolview: Stream privacy for grassroots participatory sensing," in *Proc. ACM SenSys*, Nov. 2008, pp. 281–294.

[7] A. Kapadia, D. Kotz, and N. Triandopoulos, "Opportunistic sensing: Security challenges for the new paradigm," in *Proc. COMSNETS*, Jan. 2009.

[8] A. Thiagarajan, L. Ravindranath, K. LaCurts, S. Madden, H. Balakrishnan, S. Toledo, and J. Eriksson, "VTrack: Accurate, energy-aware road traffic delay estimation using mobile phones," in *Proc. ACM SenSys*, Nov. 2009, pp. 85–98.

[9] U. Moller, L. Cottrell, P. Palfrader, and L. Sassaman, "Mixmaster protocol version 2," IETF Internet Draft, July 2003.

[10] E. Cristofaro and C. Soriente, "PEPSI: Privacy enhancing participatory sensing infrastructure," in *Proc. ACM WiSec*, June 2011.

[11] J. Girao, M. Schneider, and D. Westhoff, "CDA: Concealed data aggregation in wireless sensor networks," in *Proc. ACM WiSec*, Oct. 2004.

[12] C. Castelluccia, E. Mykletun, and G. Tsudik, "Efficient aggregation of encrypted data in wireless sensor networks," in *Proc. MobiQuitous*, July 2005, pp. 109–117.

[13] W. He, X. Liu, H. Nguyen, K. Nahrstedt, and T. F. Abdelzaher, "PDA: Privacy-preserving data aggregation in wireless sensor networks," in *Proc. IEEE INFOCOM*, May 2007, pp. 2045–2053.

[14] T. Feng, C. Wang, W. Zhang, and L. Ruan, "Confidentiality protection for distributed sensor data aggregation," in *Proc. IEEE INFOCOM*, Apr. 2008, pp. 475–483.

[15] W. Zhang, C. Wang, and T. Feng, "Gp^{2}s: Generic privacy-preservation solutions for approximate aggregation of sensor data (concise contribution)," in *Proc. IEEE PerCom*, Mar. 2008, pp. 179–184.

[16] D. Westhoff, J. Girao, and M. Acharya, "Concealed data aggregation for reverse multicast traffic in sensor networks: Encryption, key distribution, and routing adaptation," *IEEE Trans. Mobile Comput.*, vol. 5, no. 10, pp. 1417–1431, Oct. 2006.

[17] B. Przydatek, D. Song, and A. Perrig, "SIA: Secure information aggregation in sensor networks," in *Proc. ACM SenSys*, Nov. 2003, pp. 255–265.

[18] H. Chan, A. Perrig, and D. Song, "Secure hierarchical in-network aggregation in sensor networks," in *Proc. ACM CCS*, Oct. 2006, pp. 278–287.

[19] S. Roy, S. Setia, and S. Jajodia, "Attack-resilient hierarchical data aggregation in sensor networks," in *Proc. SASN*, Oct. 2006, pp. 71–82.

[20] Y. Yang, X. Wang, S. Zhu, and G. Cao, "SDAP: A secure hop-by-hop data aggregation protocol for sensor networks," in *Proc. ACM MobiHoc*, May 2006, pp. 356–367.

[21] S. Roy, M. Conti, S. Setia, and S. Jajodia, "Securely computing an approximate median in wireless sensor networks," in *Proc. ICST SecureComm*, 2008, pp. 6:1–6:10.

[22] H. Yu, "Secure and highly-available aggregation queries in large-scale sensor networks via set sampling," in *Proc. ACM/IEEE IPSN Conf.*, Apr. 2009, pp. 1–12.

[23] B. Chen and H. Yu, "Secure aggregation with malicious node revocation in sensor networks," in *Proc. ICDCS*, June 2011, pp. 581–592.

[24] C. Wang, G. Wang, W. Zhang, and T. Feng, "Reconciling privacy preservation and intrusion detection in sensory data aggregation," in *Proc. IEEE INFOCOM*, Apr. 2011, pp. 336–340.

[25] Y. Zhang and Y. Fang, "ARSA: An attack-resilient security architecture for multi-hop wireless mesh networks," *IEEE J. Sel. Areas Commun., Special Issue on High-Speed Network Security - Architecture, Algorithms, and Implementation*, vol. 24, no. 10, pp. 1916–1928, Oct. 2006.

[26] Y. Zhang, W. Liu, W. Lou, and Y. Fang, "Securing mobile ad hoc networks with certificateless public keys," *IEEE Trans. Dependable Secure Comput.*, vol. 3, no. 4, pp. 386–399, Oct.–Dec. 2006.

[27] Z. Yang, B. Zhang, J. Dai, A. Champion, D. Xuan, and D. Li, "E-SmallTalker: A distributed mobile system for social networking in physical proximity," in *Proc. IEEE ICDCS*, June 2010, pp. 468–477.

[28] A. Dua, N. Bulusu, W.-C. Feng, and W. Hu, "Towards trustworthy parcitipatory sensing," in *Proc. USENIX HotSec*, Aug. 2009.

[29] S. Saroiu and A. Wolman, "I am a sensor, and I approve this message," in *Proc. ACM HotMobile*, Mar. 2010, pp. 37–42.

[30] P. Gilbert, L. Cox, J. Jung, and D. Wetherall, "Toward trustworthy mobile sensing," in *Proc. ACM HotMobile*, Feb. 2010, pp. 31–36.

[31] P. Gilbert, J. Jung, K. Lee, H. Qin, D. Sharkey, A. Sheth, and L. Cox, "YouProve: Authenticity and fidelity in mobile sensing," in *Proc. SenSys*, Nov. 2011.

[32] C. Perkins, E. Belding-Royer, and S. Das, "Ad hoc on-demand distance vector (AODV) routing," RFC 3561, July 2003.

[33] R. Zhang, J. Shi, Y. Zhang, and C. Zhang, "A peer-to-peer approach to verifiable privacy-preserving aggregation in people-centric urban sensing systems," Arizona State University, Tempe, AZ, Technical Report, 2012 [Online]. Available: http://wins.lab.asu.edu/files/rui-JSACPeer12Full.pdf

Rui Zhang received the B.E. in communication engineering and the M.E. in communication and information systems from Huazhong University of Science and Technology, China, in 2001 and 2005, respectively, and the Ph.D. degree in electrical engineering from Arizona State University in 2013. He was a software engineer at the UT Starcom Shenzhen R&D Center from 2005 to 2007. He has been an assistant professor in the Department of Electrical Engineering at the University of Hawaii since July 2013. His primary research interests are network and distributed system security, wireless networking, and mobile computing.

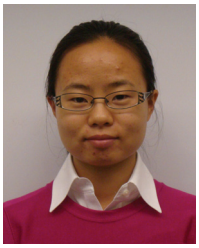

Jing Shi received the B.E. in communication engineering and the M.E. in communication and information systems from Huazhong University of Science and Technology, China, in 2003 and 2006, respectively, and the Ph.D. in electrical and computer engineering from the New Jersey Institute of Technology in 2010. She is currently a lecturer in the School of Public Administration at Huazhong University of Science and Technology, China. Her research interests are network and distributed system security, wireless networking, and mobile computing.

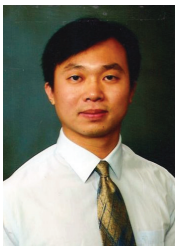

Yanchao Zhang received the B.E. in computer science and technology from Nanjing University of Posts and Telecommunications, China, in 1999; the M.E. in computer science and technology from Beijing University of Posts and Telecommunications, China, in 2002; and the Ph.D. in electrical and computer engineering from the University of Florida in 2006. He is an Associate Professor in the School of Electrical, Computer, and Energy Engineering at Arizona State University and was an Assistant Professor of Electrical and Computer Engineering at the New Jersey Institute of Technology from 2006 to 2010. His primary research interests are network and distributed system security, wireless networking, and mobile computing. He is an Associate Editor of the IEEE TRANSACTIONS ON VEHICULAR TECHNOLOGY and a Technical Editor of *IEEE Wireless Communications*. He received the NSF CAREER Award in 2009.

Chi Zhang received the B.E. and M.E. degrees in electrical and information engineering from Huazhong University of Science and Technology, China, in 1999 and 2002, respectively, and the Ph.D. degree in electrical and computer engineering from the University of Florida in 2011. He joined the University of Science and Technology of China in September 2011 as an Associate Professor of the School of Information Science and Technology. His research interests are in the areas of network protocol design, network performance analysis, and network security guarantee, particularly for wireless networks and social networks. He has served on the Technical Program Committee (TPC) for several international conferences including IEEE INFOCOM, ICC, GLOBECOM, WCNC, and PIMRC. He received the 7th IEEE ComSoc Asia-Pacific Outstanding Young Researcher Award in 2012.