Privacy-Preserving Set-Based Estimation Using Differential Privacy and Zonotopes

Mohammed M. Dawoud, Changxin Liu, Karl H. Johansson, and Amr Alanwar

***Abstract*—For large-scale cyber-physical systems, the collaboration of spatially distributed sensors is often needed to perform the state estimation process. Privacy concerns arise from disclosing sensitive measurements to a cloud estimator. To solve this issue, we propose a differentially private set-based estimation protocol that guarantees true state containment in the estimated set and differential privacy (DP) for the sensitive measurements throughout the set-based state estimation process within the central DP (CDP) and local DP (LDP) models. Zonotopes are employed in the proposed differentially private set-based estimator, offering computational advantages in set operations. We consider a plant of a nonlinear discrete-time dynamical system with bounded modeling uncertainties, sensors that provide sensitive measurements with bounded measurement uncertainties, and a cloud estimator that predicts the system's state. The privacy-preserving noise perturbs the centers of measurement zonotopes, thereby concealing the precise position of these zonotopes, i.e., ensuring privacy preservation for the sets containing sensitive measurements. Compared to existing research, our approach achieves less privacy loss and utility loss through the CDP and LDP models by leveraging a numerically optimized truncated noise distribution. The proposed estimator is perturbed by weaker noise than the analytical approaches in the literature to guarantee the same level of privacy, therefore improving the estimation utility. Numerical and comparison experiments with truncated Laplace noise are presented to support our approach.**

***Index Terms*—Differential privacy (DP), set-based estimation, truncated noise distribution, zonotopes.**

I. Introduction

PRESERVING privacy in cyber-physical systems that rely on sensor measurements, such as autonomous vehicles, drones, and medical monitors, remains a critical challenge, especially when such measurements contain sensitive information. Existing differential privacy (DP) methods often focus on pointwise estimation or assume centralized processing with strong trust assumptions. However, in many safety-critical systems, the true state must lie within a guaranteed set despite measurement noise and modeling uncertainty, requiring a privacy-preserving approach tailored for set-based estimation.

This challenge is increasingly critical in intelligent systems driven by the Internet of Things (IoT), where sharing sensor data is essential for real-time control and coordination. However, such data often include sensitive information such as user locations, medical signals, or financial patterns that must be protected from inference attacks during processing [1]. In intelligent transportation systems (ITSs), for instance, location sharing services or vehicular communication may inadvertently leak drivers' identities or trajectories, raising significant privacy concerns [2], [3], [4].

Motivated by these challenges, this work introduces a differentially private set-based estimator tailored for sets containing sensitive measurements. This differentially private set-based estimator integrates DP with zonotope-based set representations to safeguard the privacy of these sensitive measurements throughout the set-based state estimation process while minimizing utility loss and is evaluated in localization scenarios using both root-mean-square error (RMSE) and state zonotope deviation as performance metrics.

A. Set-Based Estimation

State estimation is essential in modern control and monitoring systems, supporting decision-making in domains such as robotics and autonomous vehicles [5]. It can be broadly classified into point-wise estimation, which computes a single-valued state estimate, and set-based estimation, which computes a set guaranteed to contain the true system state based on measurements and system models, if available.

Set-based estimators rely on different set representations, including intervals, ellipsoids, polytopes, zonotopes, and constrained zonotopes [6], [7], [8]. Among these, zonotopes offer computational efficiency and are closed under common set operations, making them suitable for real-time applications [9].

Estimation techniques can further be divided into model-based, which utilize known system models, and data-driven approaches, which bypass explicit modeling in favor of direct data analysis. As systems grow more complex and sensor-rich, constructing accurate system models becomes increasingly difficult and costly [10], making data-driven set-based estimation an attractive alternative [11].

Received 7 April 2025; revised 22 May 2025 and 7 August 2025; accepted 7 September 2025. Date of publication 16 September 2025; date of current version 8 December 2025. This work was supported by European Union's Horizon 2020 Research and Innovation Program under Agreement 830927. *(Corresponding author: Mohammed M. Dawoud.)*

Mohammed M. Dawoud is with the School of Computer Science and Engineering, Constructor University, 28759 Bremen, Germany (e-mail: mdawoud@constructor.university).

Changxin Liu is with the Key Laboratory of Smart Manufacturing in Energy Chemical Process, Ministry of Education, East China University of Science and Technology, Shanghai 200231, China (e-mail: changxinl@ecust.edu.cn).

Karl H. Johansson is with the School of Electrical Engineering and Computer Science, KTH Royal Institute of Technology, 114 28 Stockholm, Sweden (e-mail: kallej@kth.se).

Amr Alanwar is with the School of Computer Science and Engineering, Constructor University, 28759 Bremen, Germany, and also with the School of Computation, Information and Technology, Technical University of Munich, 80333 Heilbronn, Germany (e-mail: aalanwar@constructor.university).

Digital Object Identifier 10.1109/JIOT.2025.3610366

B. State Estimation Under Privacy Constraints

Private state estimation strategies are broadly categorized into encryption- and nonencryption-based approaches. Encryption-based schemes, such as those employing homomorphic encryption [4], [12], [13], [14], secure data by computation on encrypted measurements, albeit often at significant computational cost.

In contrast, most recent works in the latter are based on DP, a powerful noncryptographic technique for quantifying and preserving individuals' privacy. In essence, DP characterizes a property of a randomized algorithm, ensuring that its output remains stable regardless of any changes made to an individual's information in the database. This stability protects individuals' privacy by mitigating privacy-related attacks. DP operates under two models: central DP (CDP), where a trusted data curator applies DP mechanisms to a centralized dataset, and local DP (LDP), where each participant randomizes their data locally before sharing [15]. The tradeoff between privacy and accuracy in these models is governed by the magnitude and distribution of the injected noise. Prior works have extended DP to filtering and observer designs for pointwise estimation, such as differentially private Kalman filters for linear and nonlinear systems [16], [17], [18]. However, these approaches do not directly address systems with bounded disturbances and uncertainties, where set-based estimators are more appropriate.

Recently, a differentially private interval observer was proposed in [19] for such cases, employing bounded input perturbations. Building upon this foundation, our previous work [20] introduced a differentially private set-based estimator using a truncated additive mechanism with a numerically optimized noise distribution under the CDP model, leveraging zonotopes for efficient set representation. This article extends and refines that framework to enhance privacy protection while maintaining estimation utility.

C. Contributions

This article presents a differentially private set-based estimator that ensures true state containment within the estimated set and provides DP for sensitive measurements throughout the set-based state estimation process. It extends the work in [20] to nonlinear discrete-time systems within the context of the CDP and LDP models and evaluates this extension on real-world data. Specifically, we continue to utilize zonotopes for set representation and employ a truncated additive mechanism with a numerically optimized noise distribution in both models to present a differentially private set-based estimator for nonlinear discrete-time systems with bounded modeling and measurement uncertainties. The proposed estimator ensures the privacy of sets containing sensitive sensors' measurements throughout the estimation process with minimal utility loss. Comparative results provided in this article demonstrate the enhancement in estimation utility.

The main contributions of this article can be summarized as follows.

1) We introduce a differentially private set-based estimator leveraging a truncated additive mechanism with a numerically optimized noise distribution [21] and employing zonotopes for set representation. The proposed estimator ensures the privacy of sets containing sensors' measurements throughout the estimation process with minimal utility loss.
2) We apply the privacy-preserving noise to the zonotopes containing the sensitive measurements within the context of the CDP and LDP models, concealing the precise positions of these zonotopes.
3) Through our comprehensive evaluation of the proposed differentially private set-based estimator, we illustrate that the truncated Laplace distribution [19] necessitates a more significant amount of noise compared to the truncated optimal noise distribution to achieve an equivalent level of privacy under both the CDP and LDP models.

All utilized data and code are publicly available.[1]

The rest of this article is organized as follows. The preliminaries and problem statement are presented in Section II. The algorithms are designed and evaluated in Sections III and IV, respectively. Finally, we conclude the work in Section V.

II. Preliminaries and Problem Setup

In this section, we introduce some notation, present the needed preliminaries, and then formulate the problem.

A. Notation

Throughout this article, we denote vectors and scalars by lower case letters, matrices by upper case letters, the set of real numbers by $\mathbb{R}$, the set of integers by $\mathbb{Z}$, and the set of positive integers by $\mathbb{Z}^+$. For a given vector v, we denote the matrix with v on the diagonal as $\text{diag}(v)$. For a given matrix $M \in \mathbb{R}^{m\times n}$, its transpose is given by M^{T} and its Frobenius norm is given by $\|M\|_F = (\text{Tr}(M^{\mathrm{T}}M))^{1/2}$. We denote the ith element of a vector or list a by $a^{(i)}$. For a vector $a \in \mathbb{R}^n$, we denote its L_2 norm by $|a|_2 = \left(\sum_{i=1}^{n}(a^{(i)})^2\right)^{1/2}$. For a scalar or a vector a_k, we use $\{a_k\}$ to denote a list of multiple a_k such that $k \in \{1,\ldots,k_1\}$ and $k_1 \in \mathbb{Z}^+$. We denote the L_2 norm of a vector-valued signal y composed of a list of $y_k \in \mathbb{R}^n$, i.e., $y = \{y_k\}$, by $\|y\|_2 = \left(\sum_{k=1}^{\infty}(|y_k|_2)^2\right)^{1/2}$.

B. Set Representation and Set-Based Estimation

Next, we review set representation, set operations, and set-based estimation based on zonotopes.

1) Set Representation and Operations: The zonotope is defined as follows.

Definition 1 (Zonotope [22]): Given a center $c_{\mathcal{Z}} \in \mathbb{R}^n$ and $\gamma_{\mathcal{Z}} \in \mathbb{N}$ generator vectors in a generator matrix $G_{\mathcal{Z}} = \left[g_{\mathcal{Z}}^{(1)} \cdots g_{\mathcal{Z}}^{(\gamma_{\mathcal{Z}})}\right] \in \mathbb{R}^{n\times\gamma_{\mathcal{Z}}}$, a zonotope is defined as

$$\mathcal{Z} = \left\{ x \in \mathbb{R}^n \;\middle|\; x = c_{\mathcal{Z}} + \sum_{i=1}^{\gamma_{\mathcal{Z}}} \beta^{(i)} g_{\mathcal{Z}}^{(i)}, -1 \le \beta^{(i)} \le 1 \right\}. \quad (1)$$

We use the shorthand notation $\mathcal{Z} = \langle c_{\mathcal{Z}}, G_{\mathcal{Z}} \rangle$ for such a zonotope.

[1] https://github.com/mohammed-dawoud/Differentially-Private-Set-Based-Estimation-Using-Zonotopes

Zonotopes are closed under linear maps and Minkowski sum [23]. The linear map $L \in \mathbb{R}^{m\times n}$ for zonotope $\mathcal{Z}$ is defined and computed as follows:

$$L\mathcal{Z} = \{Lz | z \in \mathcal{Z}\} = \langle Lc_{\mathcal{Z}}, L\, G_{\mathcal{Z}}\rangle. \tag{2}$$

Given two zonotopes $\mathcal{Z}_1 = \langle c_{\mathcal{Z}_1}, G_{\mathcal{Z}_1}\rangle$ and $\mathcal{Z}_2 = \langle c_{\mathcal{Z}_2}, G_{\mathcal{Z}_2}\rangle$, the Minkowski sum is defined and computed as

$$\begin{aligned}\mathcal{Z}_1 \oplus \mathcal{Z}_2 &= \{z_1 + z_2 | z_1 \in \mathcal{Z}_1, z_2 \in \mathcal{Z}_2\} \\ &= \left\langle c_{\mathcal{Z}_1} + c_{\mathcal{Z}_2}, \left[G_{\mathcal{Z}_1}, G_{\mathcal{Z}_2}\right]\right\rangle.\end{aligned} \tag{3}$$

The Cartesian product is

$$\begin{aligned}\mathcal{Z}_1 \times \mathcal{Z}_2 &= \left\{\begin{bmatrix} z_1 \\ z_2\end{bmatrix} \middle| z_1 \in \mathcal{Z}_1, z_2 \in \mathcal{Z}_2\right\} \\ &= \left\langle \begin{bmatrix} c_{\mathcal{Z}_1} \\ c_{\mathcal{Z}_2}\end{bmatrix}, \begin{bmatrix} G_{\mathcal{Z}_1} & 0 \\ 0 & G_{\mathcal{Z}_2}\end{bmatrix}\right\rangle.\end{aligned} \tag{4}$$

Zonotopes are used in the proposed scheme because they provide an efficient and compact representation of uncertainty sets. Specifically, they are [23], [24] given as follows.

1) Closed under linear transformations and Minkowski sums, enabling efficient state prediction and correction.
2) Computationally more efficient than other set representations (e.g., polytopes or ellipsoids) because their complexity scales linearly with the number of generators.
3) The linear complexity in terms of the number of generators offers a significant computational advantage to zonotopes, especially for high-dimensional systems.
4) Zonotopes have compact memory usage as they only require storing their generators and centers. In contrast, polytopes require storing their vertices, whose number grows exponentially with the dimension.
5) Compared to polytopes or ellipsoids, zonotopes achieve tighter bounding of the state space, improving the accuracy of the estimator.

2) Set-Based Estimation: Consider the following nonlinear discrete-time dynamical system with bounded modeling and measurement uncertainties:

$$\begin{aligned}x_{k+1} &= f(x_k) + w_k \\ y_k^{(i)} &= h^{(i)}(x_k) + v_k^{(i)}\end{aligned} \tag{5}$$

where $x_k \in \mathbb{R}^n$ is the system state at time step $k \in \mathbb{Z}^+$ and $y_k^{(i)} \in \mathbb{R}$ is the measurement of sensor $i \in \{1, \ldots, m\}$ with m equals the number of available sensors. The functions f and $h^{(i)}$ are assumed to be differentiable. The vector w_k and the scalar $v_k^{(i)}$ are process and measurement noise, respectively. They are assumed to be unknown but bounded by the zonotopes $\mathcal{Z}_w = \langle 0, G_w\rangle$ and $\mathcal{Z}_v^{(i)} = \langle 0, G_v^{(i)}\rangle$ (if the noise zonotopes are not centered around zero, the resulting estimates will be shifted). The system has a bounded initial state $x_0 \in \bar{\mathcal{Z}}_0 = \langle \bar{c}_0, \bar{G}_0\rangle$. At each time step k, the set-based state estimator aims to find the corrected state set $\bar{\mathcal{Z}}_k$ by finding the intersection between the predicted state set $\hat{\mathcal{Z}}_k$ and the measurement sets $\hat{\mathcal{Z}}_{y_k}^{(i)} = \langle y_k^{(i)}, G_v^{(i)}\rangle$ corresponding to the sensor measurements $y_k^{(i)}$ and measurement uncertainties $\mathcal{Z}_v^{(i)}$, $i = 1, \ldots, m$ [11], [25], [26]. The set-based state estimator is described by Algorithm 1.

Algorithm 1 Set Estimation Using Zonotopes

Input: Process noise zonotope $\mathcal{Z}_w = \langle 0, G_w\rangle$, and measurement noise zonotopes $\mathcal{Z}_v^{(i)} = \langle 0, G_v^{(i)}\rangle$, where $i \in \{1, \ldots, \}$.
Output: $\bar{\mathcal{Z}}_k = \langle \bar{c}_k, \bar{G}_k\rangle$.
Initialization: Set $k = 1$ and $\bar{\mathcal{Z}}_0 = \langle \bar{c}_0, \bar{G}_0\rangle$.
while *True* **do**
The cloud estimator uses the m measurement sets $\mathcal{Z}_{y_k}^{(i)}$ containing the sensitive measurements to perform the estimation process according to the following steps:
- Prediction step: The set-based state estimator determines the predicted state set $\hat{\mathcal{Z}}_k$ using a Taylor series expansion as described in [23]. The predicted state set is determined based on the past corrected state set $\bar{\mathcal{Z}}_{k-1}$ and a process noise zonotope $\mathcal{Z}_w$. This is given by
$$\hat{\mathcal{Z}}_k = f\left(\bar{\mathcal{Z}}_{k-1}\right) \oplus \mathcal{Z}_w. \tag{6}$$
- Correction step: The corrected state set $\bar{\mathcal{Z}}_k$ is determined by the reduction of $\acute{\mathcal{Z}}_k = \langle \acute{c}_k, \acute{G}_k\rangle$, which overapproximates the intersection between the predicted state set $\hat{\mathcal{Z}}_k = \langle \hat{c}_k, \hat{G}_k\rangle$ and the m measurement sets $\mathcal{Z}_{y_k}^{(i)}$. This is given by
$$\acute{\mathcal{Z}}_k \supseteq \hat{\mathcal{Z}}_k \cap_{i=1}^{m} \mathcal{Z}_{y_k}^{(i)}. \tag{7}$$
- Update the time: $k = k + 1$.

end while

To obtain the corrected state set zonotope $\bar{\mathcal{Z}}_k = \langle \bar{c}_k, \bar{G}_k\rangle$ with $\bar{c}_k = \acute{c}_k$, the order of the generator matrix $\acute{G}_k$ of $\acute{\mathcal{Z}}_k$ is reduced as follows:

$$\bar{G}_k = \downarrow_q \acute{G}_k \tag{8}$$

where $\downarrow_q$ denotes the reduction in the order of the generator matrix according to [9].

C. Differential Privacy

We will present a few notions on DP, which will be used later to develop a differentially private set-based estimator. Let $\mathcal{H}$ denote the space of datasets of interest (e.g., sensor measurements) [17]. We define a symmetric binary relation on $\mathcal{H}$, called adjacency and denoted by Adj, in which two datasets $h, \acute{h} \in \mathcal{H}$ are called adjacent, denoted by $\text{Adj}(h, \acute{h})$, if they differ by the value of exactly one individual's data [27]. Given a pair of adjacent datasets $h, \acute{h} \in \mathcal{H}$, a differentially private randomized mechanism M with an anonymized measurable output space $\mathcal{O}$, aims to prevent an adversary from inferring knowledge about an individual's data by generating randomized outputs $M(h), M(\acute{h}) \in \mathcal{O}$ with close distributions for adjacent inputs. We use a pair of nonnegative constants (ϵ, δ) to quantify the privacy loss [17], [28], [29]. As the size of a dataset increases, the relative contribution of an individual's data diminishes, requiring less stringent privacy-preserving measures to maintain a given level of an individual's privacy.

Definition 2 (Approximate DP [17], [28]): A randomized mechanism M, which maps $\mathcal{H}$ equipped with the adjacency relation $\text{Adj}(h, \acute{h})$ to a measurable space

$\mathcal{O}$, is (ϵ, δ)-approximate differential privacy (ADP), $\epsilon, \delta \geq 0$, if $\forall S \in \mathcal{O}$ and $\forall\, h, \acute{h} \in \mathcal{H}$ such that $\text{Adj}(h, \acute{h})$

$$\Pr\left[M\left(h\right) \in S\right] \leq e^{\epsilon} \Pr\left[M\left(\acute{h}\right) \in S\right] + \delta. \tag{9}$$

If $\delta = 0$, M is said to be ϵ-differentially private.

For the set-based estimation problem, the adjacent datasets $h, \acute{h} \in \mathcal{H}$ can be the centers of measurement sets. In the CDP, as the dataset size increases, the relative impact of an individual sensor's measurement on the overall data decreases. While the sensitivity of individual measurements may remain constant, the overall effect of any single measurement on aggregated results diminishes. Consequently, less noise needs to be added to each measurement to achieve the same level of DP, improving the accuracy of the protected data while still ensuring privacy. In contrast, in the LDP, each sensor perturbs its own data before sending it to an aggregator, meaning that noise is added at the individual level rather than at the aggregate level. Since data of each sensor are already noisy before aggregation, increasing the dataset size does not reduce the impact of individual noise in the same way it does in CDP. While a larger dataset may help stabilize aggregated statistics, each sensor's contribution remains independently perturbed, preventing the noise from averaging out as efficiently as in CDP.

To ensure privacy in the set-based estimation process, we consider two primary models of DP: CDP and LDP. Each model offers different privacy guarantees and has distinct practical implications.

1) *Central DP:*
   a) In CDP, a trusted entity (e.g., a sensor manager) aggregates raw measurements from multiple sensors and then applies a differentially private noise mechanism before sharing the sanitized data with the cloud estimator.
   b) This model assumes that the data curator is trusted and follows the privacy-preserving protocol correctly.
   c) Since noise is added only once at the aggregate level, less noise is required, resulting in higher estimation accuracy compared to LDP.
2) *Local DP:*
   a) In LDP, each sensor individually perturbs its measurement before transmission, ensuring that even an untrusted cloud estimator never sees raw sensor readings.
   b) Since noise is added at the individual level, it needs to be stronger to mask each sensor's contribution independently, leading to higher utility loss compared to CDP.
   c) LDP is beneficial when the sensor manager is untrusted, but at the cost of higher utility loss.

Table I summarizes the key differences between CDP and LDP.

This distinction plays a crucial role in our proposed privacy-preserving set-based estimator. In our setup, the CDP model is used when a trusted sensor manager is available, allowing for better estimation accuracy with minimal noise. However, in scenarios where the sensor manager is untrusted, the LDP model ensures privacy at the cost of increased noise and reduced utility.

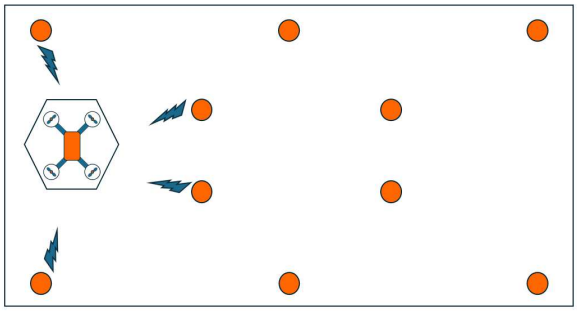

Fig. 1. Intrusion detection system installed over a region, : the quadcopter, : the estimated set, and : LIDAR sensors.

To motivate the work in this article, consider the following scenario.

Example 1: An intrusion detection system installed over a region consists of light imaging, detection, and ranging (LIDAR) sensors distributed throughout the area and an untrusted cloud estimator Fig. 1. This cloud estimator estimates the set of possible locations of an intruding quadcopter in the region using distance measurements with bounded noise provided by the LIDAR sensors. The cloud estimator utilizes these measurement sets to determine the set of possible locations of the intruding quadcopter. We aim to safeguard the measurements against untrusted parties throughout the estimation process.

D. Problem Setup

Consider the following entities for the two cloud-based state estimation setups visualized in Fig. 2(a) and (b) [25].

1) *Plant:* A system that we aim to estimate its set of possible states. We consider a system with a publicly known nonlinear discrete-time model described in (5).
2) *Sensors:* An array of m sensors, where each entity i produces a private measurement denoted as $y_k^{(i)}$, with $i \in 1, \ldots, m$.
3) *Sensor Manager:* An entity with computational capabilities that enable it to aggregate the measurements of all sensors and perturb them with DP noise.
4) *Cloud Estimator:* An untrusted entity that performs set-based estimation for the system state.

The setup in Fig. 2(a) uses the LDP model of DP, in which each sensor $i \in \{1, \ldots, m\}$ locally perturbs its measurement with the privacy-preserving noise before transmitting it to the cloud estimator. The datasets $h, \acute{h} \in \mathcal{H}$ in this setup are the centers of measurement sets of a single sensor i. The setup in Fig. 2(b) uses the CDP model of DP, in which the sensor manager acts as a trusted data curator, adding privacy-preserving noise to the measurements of m sensors. The datasets in this setup are the centers of measurement sets of m sensors. In both setups, "protected" refers to measurement sets that have been safeguarded with privacy-preserving noise, while

TABLE I
COMPARISON BETWEEN CDP AND LDP MODELS

Property	CDP	LDP
Privacy Guarantee	Assumes a trusted curator	No trusted entity required
Noise Addition	Applied at the aggregate level	Applied at each sensor locally
Estimation Accuracy	Higher (weaker noise)	Lower (stronger noise)
Utility Loss	Lower	Higher
Suitable Use Case	Trusted sensor manager	Untrusted cloud estimator

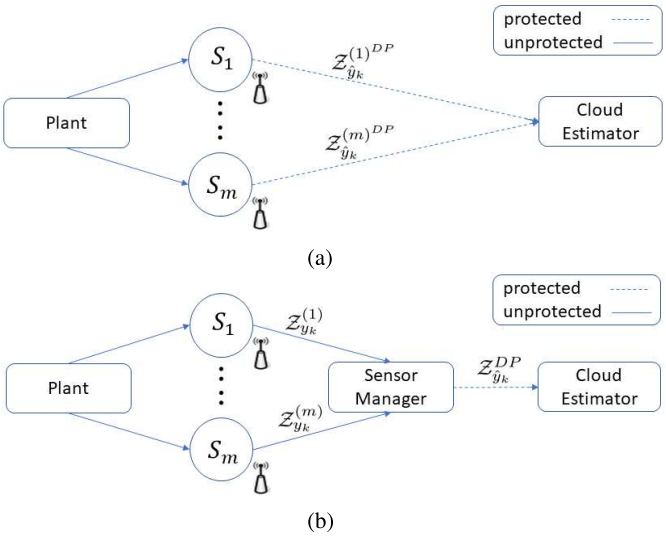

Fig. 2. Setups of the cloud estimator within the CDP and LDP models. (a) Cloud estimator is set up within the context of the LDP model. (b) Cloud estimator is set up within the context of the CDP model.

"unprotected" refers to those without any privacy-preserving noise.

Problem 1: We aim to design a differentially private set-based estimator for a plant with the model described in (5) such that the state estimates are obtained while preserving the privacy of the measurement sets. This estimator obtains state estimates while keeping each sensor's measurement private to the sensor and protecting it from any untrusted entity (e.g., the cloud estimator). This is accomplished within the context of the LDP model for the setup in Fig. 2(a) and within the context of the CDP model for the setup in Fig. 2(b).

III. DIFFERENTIALLY PRIVATE SET-BASED ESTIMATION

In this section, we develop the differentially private set-based estimator. Based on the characteristics of the zonotopic representation of sets, the actual positions of these sets in the measurement space can be concealed if we perturb their centers. Thus, we protect the measurement sets by adding privacy-preserving noise to their centers, thereby protecting their positions. In Section III-A, an additive noise mechanism employs the LDP model, while the other additive noise mechanism utilizes the CDP model to introduce perturbations to the sensitive measurements. In both cases, we incorporate optimal noise characterized by a numerically generated truncated distribution. These perturbations result in protected measurements. Then, in Section III-B, we deploy the measurement sets containing the protected measurements to the set-based estimator. This enables the derivation of state estimates while keeping the measurements protected against any untrusted party.

A. Design of Truncated Optimal Additive Noise

Additive noise mechanisms such as the Gaussian and Laplace mechanisms typically perturb the measurements with particular random noise to achieve DP. However, most of the results require the support of the noise distribution to be unbounded, except for a few cases such as [30] and [31]. This class of mechanisms is not applicable for some applications, such as safety-critical systems [21], which require state enclosure guarantees in a bounded set. One closely related study [19] developed an interval observer with DP based on a truncated Laplace mechanism. However, the design relies on an L_1 norm-based adjacency relation and an analytical bound for the noise variance, which may cause the design to be conservative, consuming more noise to achieve DP (a numerical comparison is provided in Section IV). To solve this issue, we follow the numerical approach, recently developed in [21], to optimize a noise distribution that is subjected to a bounded support constraint and the privacy constraint in Definition 2.

In our setups, all sets are defined over the continuous domain of all real numbers. However, the truncated optimal noise distribution we employ to achieve DP is generated numerically. This distribution consists of discrete noise occurrence probabilities, with noise values selected from the continuous domain. Consequently, even after adding this noise to the sets, the resultant domain remains continuous.

Next, we define a class of truncated noise distribution functions. Then, for a fixed DP parameter ϵ and a particular noise model, we present an optimization problem, where the objective function balances between the privacy loss parameter δ and the utility loss. Upon solving this optimization problem, an optimal noise distribution is generated.

Definition 3 (Truncated Noise Distribution [21]): Let $\Phi = [-d, d]$ define a bounded noise range such that $d \in \mathbb{R}$ and $\acute{\Phi} = \{\phi_l\}_{l\in\{1,\ldots,\,2\,N\}}$ be the discretization of Φ on $2\,N$ equidistant steps such that $\phi_l \in \Phi$, $N \in \mathbb{Z}^+$, and $l \in \{1, \ldots, \,2\,N\}$; then, a numerically generated, truncated, discrete, equidistant, symmetric, and monotonically decreasing from its zero center

noise distribution function, denoted by $P(\phi_l)$, has the following properties:

$$\sum_{\phi_l \in \acute{\Phi}} P(\phi_l) = 1 \quad \text{and} \quad P(\phi_l) \geq 0. \qquad \text{(Distribution)} \tag{10a}$$

$$P(\phi_l) \geq P(\phi_m) \quad \forall\, \phi_m > \phi_l > 0, \qquad \text{(Monotonicity)} \tag{10b}$$

where $l, m \in \{N+1, \ldots, 2N\}$

$$P(\phi_l) = P(-\phi_l), \qquad \text{(Symmetry)} \tag{10c}$$

where $l \in \{1, \ldots, N\}$.

For the LDP model, a dataset of interest is the local measurement signal $y^{(i)} \in \mathcal{H}$ released by sensor i as a list of measurements $\{y_k^{(i)}\}_{k\in\{1,\ldots,\mathcal{T}\}}$, where $\mathcal{T} = \infty$ is also of interest [19]. We aim to provide a certain level of privacy protection for a single sensor's measurement. Two local measurement signals $y^{(i)}$ and $\acute{y}^{(i)}$ are called adjacent and can be denoted by $\text{Adj}(y^{(i)}, \acute{y}^{(i)})$, if and only if they differ by the value of exactly one measurement $y_k^{(i)}$ [17]. In other words, $y^{(i)}$ and $\acute{y}^{(i)}$ are considered adjacent if there is any time step k at which $y_k^{(i)} \neq \acute{y}_k^{(i)}$. Since the privacy-preserving noise is added to the measurements themselves, the allowed variation within $\text{Adj}(y^{(i)}, \acute{y}^{(i)})$ is bounded by what is called sensitivity, which is defined formally as follows.

Definition 4 (Sensitivity—LDP [19], [32]): For a given sensor i, the allowed deviation for a single measurement $y_k^{(i)}$ between two adjacent local measurement signals $y^{(i)}$ and $\acute{y}^{(i)}$, i.e., $\text{Adj}(y^{(i)}, \acute{y}^{(i)})$, is bounded in the L_2 norm by s and given by

$$\left\| y^{(i)} - \acute{y}^{(i)} \right\|_2 \leq s \tag{11}$$

where $s \geq 0$.

Next, the following definition describes the additive noise mechanism for the LDP setup, shown in Fig. 2(a), in which each participating sensor locally perturbs its own measurement, and the sensor manager is not present in this setup since it is considered an untrusted entity.

Definition 5 (Additive Noise Mechanism—LDP [21]): Given a measurement of sensor $iy_k^{(i)}$ and a noise sample $\phi_k \in \acute{\Phi}$ such that the successive samples of ϕ_k are independent and identically distributed (IID) with the probability distribution in Definition 3 and satisfying Lemma 1, then the additive noise mechanism $M_y^{(i)}$ is defined as

$$M_y^{(i)} : \hat{y}_k^{(i)^{\text{DP}}} = y_k^{(i)} + \phi_k \tag{12}$$

where $\hat{y}_k^{(i)^{\text{DP}}}$ is a protected measurement.

Lemma 1 ([21, Theorem 15]): Let $M_y^{(i)}$ be an additive noise mechanism with a sensitivity s (Definition 4) and $\acute{\Phi} = \{\phi_l\}_{l\in\{1,\ldots,2N\}}$ be the discretization of Φ with the truncated optimal noise distribution $P(\phi_l)$ (Definition 3). If $\forall\, \hat{\Phi} \subseteq \acute{\Phi}$

$$\sum_{\phi_l \in \hat{\Phi}} P(\phi_l) \leq e^{\epsilon} \sum_{\phi_l \in \hat{\Phi}} P(\phi_l + s) + \delta \tag{13}$$

then the additive noise mechanism $M_y^{(i)}$ is (ϵ, δ)-ADP for any $y_k^{(i)} \in \mathcal{H}$.

It deserves noting that, given that the sensitivity s (Definition 4) is satisfied $\forall\, k \in \mathbb{Z}^+$, the mechanism $M_y^{(i)}$ is (ϵ, δ)-ADP at any time step k [16, Lemma 2].

Next, we present the loss function, denoted by $L_\gamma^{\Omega_t}$ [21]. This function balances between the privacy parameter δ and a utility loss U at a fixed ϵ and is given by

$$L_\gamma^{\Omega_t} = \delta + \Omega_t U \tag{14a}$$

where the utility loss U is given by

$$U = \left(\sum_{\phi_l \in \acute{\Phi}} |\phi_l|^\gamma P(\phi_l) \right)^{1/\gamma} \tag{14b}$$

with $\gamma \in \{1, 2\}$ such that γ selects between L_1 or L_2 norm-based utility loss, and

$$\Omega_t = \max\left(\frac{\Omega_{\text{start}}}{2^{t/\Gamma}}, \Omega_{\min} \right) \tag{14c}$$

is the utility weight at the training epoch t, where $t \in \mathbb{Z}^+$ with an exponentially decaying rate Γ from a starting value Ω_{start} and with a lower bound $\Omega_{\min}$. The optimal noise distribution $P(\phi_l)$ is then optimized by minimizing the weighted sum of the utility loss U and the privacy parameter δ. This noise distribution is generated through the following steps, and we are going to omit (ϕ_l) to ease the notation. We start by generating the first monotonically increasing half (i.e., $\{\phi_l\}_{l\in\{1,\ldots,N\}}$) of the noise distribution $P(\phi_l)$, which is given by

$$P_l = 1/2\ \text{SoftMax}(r_l);\ r_l \in \{r_0, \ldots, r_N\} \tag{15a}$$

where $\text{SoftMax}(r_l) = e^{r_l} / \sum_{i=0}^{N} e^{r_i}$ and N is the number of discretization steps in the half-width of the noise distribution $P(\phi_l)$. The SoftMax function normalizes the r_l values into a distribution, and the r_l values are generated using a model of v-stacked sigmoid functions (i.e., $\sigma(\phi_l) = (1 + e^{-\phi_l})^{-1}$), which is given by

$$r_l = \ln\left[A^2 + \sum_{j=0}^{v} B_j^2 \cdot \sigma\left(C \cdot \left(\phi_l - F_j \right) \right) \right] \tag{15b}$$

where $\phi_l = ld/N - d$. The parameters A, B_j, C, and F_j are randomly initialized and then learned to optimize the loss function in (14) using numerical optimization methods [e.g., stochastic gradient descent (SGD)].

Next, the first half (i.e., $\{\phi_l\}_{l\in\{1,\ldots,N\}}$) of the noise distribution generated by (15a) and denoted by P_l is mirrored according to the following:

$$P_j = P_{2N-j+1} \qquad \text{for } j \in \{N+1, \ldots, 2N\}. \tag{15c}$$

Then, P_j (i.e., $\{\phi_l\}_{l\in\{N+1,\ldots,2N\}}$) is concatenated to P_l (i.e., $\{\phi_l\}_{l\in\{1,\ldots,N\}}$) to obtain the symmetric noise distribution $P(\phi_l)$ (i.e., $\{\phi_l\}_{l\in\{1,\ldots,2N\}}$).

Note that the case with arbitrarily dimensional and spherically rotation-symmetric noise distributions and sensitivity conditions can be reduced to a 1-D privacy analysis [33]. Based on this claim, for the CDP setup, we can define the additive noise for the case in which the dataset of interest is the global measurement signal $y \in \mathcal{H}$ released by an array of m sensors as a list of measurements vectors $\{y_k\}_{k\in\{1,\ldots,\mathcal{T}\}}$, where $\mathcal{T} = \infty$ is also of interest and $y_k = [y_k^{(1)}, \ldots, y_k^{(m)}]^{\text{T}}$ [19]. Two global measurement signals y and $\acute{y}$ are called adjacent and can be denoted by $\text{Adj}(y, \acute{y})$, if and only if they differ by

the value of exactly one measurement vector y_k. Similarly, the allowed variation within $\text{Adj}(y, \acute{y})$ is bounded by what is called sensitivity, which is defined formally as follows.

Definition 6 (Sensitivity—CDP [19], [32]): The allowed deviation for a single measurement vector between two adjacent global measurement signals y and $\acute{y}$, i.e., $\text{Adj}(y, \acute{y})$, is bounded in the L_2 norm by s_g and given by

$$\|y - \acute{y}\|_2 \leq s_g \tag{16}$$

where $s_g \geq 0$.

Next, for the CDP setup, shown in Fig. 2(b), the sensor manager aggregates the measurements of all sensors into a measurement vector and perturbs them with the privacy-preserving noise according to the following additive noise mechanism.

Definition 7 (Additive Noise Mechanism—CDP [21]): Given a vector of measurements y_k and a noise vector $\Phi_k \in \acute{\Phi}$ of IID coordinates with the probability distribution in Definition 3 and satisfying Lemma 1 using the sensitivity s_g (Definition 6), and assuming that successive samples of Φ_k are also IID, then the additive noise mechanism M_y is defined as

$$M_y : \hat{y}_k^{\text{DP}} = y_k + \Phi_k \tag{17}$$

where $\hat{y}_k^{\text{DP}}$ is a vector of protected measurements.

Likewise, given that the sensitivity s_g (Definition 6) is satisfied $\forall\ k \in \mathbb{Z}^+$, the mechanism M_y is (ϵ, δ)-ADP at any time step k [16, Lemma 2].

1) Comparison With Other Noise Distributions [21], [30], [34]: To ensure DP while maximizing estimation utility, the proposed scheme employs a truncated optimal noise distribution. Unlike conventional noise models that rely on predefined probability distributions, the truncated optimal noise is numerically optimized using a gradient-descent-based tool. This approach allows the noise distribution to adapt to the specific sensitivity and privacy requirements of the application, ensuring a superior utility–privacy tradeoff.

Specific Features and Advantages.

1) *Numerically Optimized Noise Pattern:*
   a) The truncated optimal noise is learned by optimizing a discrete probability distribution rather than merely tuning hyperparameters of a predefined distribution.
   b) It minimizes utility loss by concentrating noise around the true value, thereby reducing the impact of noise on estimation accuracy.
   c) The noise distribution is designed to be symmetric and monotonically decreasing from the mean, which ensures DP while minimizing expected deviation.
2) *Comparison With Truncated Gaussian Noise:*
   a) Truncated Gaussian noise is widely used for (ϵ, δ)-DP due to its smoothness and well-studied privacy guarantees. However, it is not optimized for utility, leading to suboptimal tradeoffs.
   b) Truncated Gaussian noise exhibits a fixed tail behavior, leading to a higher probability of extreme values near the truncation boundary. The truncated optimal noise mitigates this issue by numerically optimizing the tail behavior based on the specific sensitivity of the data. Hence, the truncated optimal noise achieves better utility–privacy tradeoffs.
3) *Comparison With Truncated Laplace Noise:*
   a) Truncated Laplace noise is effective for achieving DP and is straightforward to implement. However, it typically introduces higher variance, which adversely affects estimation accuracy.
   b) The truncated optimal noise demonstrates a tighter concentration around the mean, reducing variance and leading to improved estimation utility.

B. Differentially Private Zonotope-Based Set-Membership Estimation

In this section, we introduce the differentially private set-based estimator. The corrupted measurements serve as the centers of the measurement sets. Hence, we have measurement sets with corrupted centers that guarantee DP for the contained sensitive measurements. The generators of all zonotopes are not corrupted with the privacy-preserving noise. We will deploy these measurement sets to the set-based estimator (Algorithm 1), i.e., the cloud estimator, to preserve the privacy of the sensitive measurements during the estimation process and get state estimates while keeping the measurements safeguarded against any untrusted entity. Hence, a differentially private version of that set-based estimator is summarized in Algorithm 2.

Algorithm 2 and Fig. 3 summarize our proposed differentially private set-based estimator. The inputs to the algorithm are the DP parameter ϵ, the noise range d, the sensitivity, the privacy-preserving noise zonotope $\mathcal{Z}_p = \langle 0, G_p \rangle$, where G_p is the generator matrix, created using the range of the optimal noise d, the process noise zonotope $\mathcal{Z}_w = \langle 0, G_w \rangle$, and the measurement noise zonotopes $\mathcal{Z}_v^{(i)} = \langle 0, G_v^{(i)} \rangle$, where $i \in \{1, \ldots, m\}$. At the initialization, the truncated optimal noise distribution $P(\phi_l)$ is generated using ϵ, d, and sensitivity according to (10) while optimizing the loss function in (14) by learning the parameters of the noise model in (15). Then, the privacy-preserving noise is added to the sensitive measurements using either of the two additive noise mechanisms described in Definitions 5 and 7 in lines 2–7. The cloud estimator uses the m measurement sets containing the protected measurements $\mathcal{Z}_{y_k}^{(i)^{\text{DP}}} = \langle y_k^{(i)^{\text{DP}}}, G_v^{(i)} \rangle$ to perform the estimation process. The cloud estimator computes the predicted state set $\hat{\mathcal{Z}}_k^{\text{DP}}$ in line 9, where the center of the state, x_{k-1}^*, is used as a linearization point for the state function f at time step $k - 1 \in \mathbb{Z}^+$. The infinite Taylor series is overapproximated by the first-order Taylor series and its Lagrange remainder $\mathcal{Z}_{L,k} = \langle c_{L,k}, G_{L,k} \rangle$. The corrected state set $\bar{\mathcal{Z}}_k^{\text{DP}}$ is computed in lines 11–16, where the center of the state, x_k^*, is used as a linearization point at time step $k \in \mathbb{Z}^+$. The infinite Taylor series is overapproximated by the first-order Taylor series and its Lagrange remainder $\mathcal{Z}_{L,k} = \langle c_{L,k}, G_{L,k} \rangle$ to guarantee true state inclusion. In line 13, as in [35], the weights $\bar{\Lambda}_k^*$ are optimized to reduce the Frobenius norm of the generator matrix $\acute{G}_k^{\text{DP}}$; therefore,

Algorithm 2 (ϵ, δ)-ADP Set Estimation Using Zonotopes

Input: ϵ, d, sensitivity, $\mathcal{Z}_p = \langle 0, G_p \rangle$, $\mathcal{Z}_w = \langle 0, G_w \rangle$, and $\mathcal{Z}_v^{(i)} = \langle 0, G_v^{(i)} \rangle$, $i \in \{1, \ldots, \}$.

Output: $\bar{\mathcal{Z}}_k^{DP} = \langle \bar{c}_k^{DP}, \bar{G}_k^{DP} \rangle$.

Initialization: The truncated optimal noise distribution $P(\phi_l)$ is generated. Set $k = 1$ and $\bar{\mathcal{Z}}_0 = \langle \bar{c}_0, \bar{G}_0 \rangle$.

1: **while** *True* **do**

2: Obtain $y_k^{(i)}$ from each sensor i, $i \in \{1, \ldots, \}$.

3: **if** the setup in Figure 2(b) is deployed **then**

4: The sensor manager aggregates the m measurements into a vector of measurements $y_k = [y_k^{(1)}, \ldots, y_k^{(m)}]^{\mathrm{T}}$, then adds the DP noise to y_k to obtain protected measurements $\hat{y}_k^{DP}$ following (17).

5: **else if** the setup in Figure 2(a) is deployed **then**

6: Each sensor locally perturbs its measurement to obtain a protected measurement $\hat{y}_k^{(i)^{DP}}$ according to (12).

7: **end if**

8: The cloud estimator uses $\mathcal{Z}_{y_k}^{(i)DP} = \left\langle y_k^{(i)^{DP}}, G_v^{(i)} \right\rangle, i \in \{1, \ldots, \}$ to perform the estimation process as follows:

9: Compute the predicted state set $\hat{\mathcal{Z}}_k^{DP} = \langle \hat{c}_k^{DP}, \hat{G}_k^{DP} \rangle$ as follows:

$$\hat{c}_k^{DP} = f\left(x_{k-1}^*\right) + \frac{\partial f_{k-1}}{\partial x}|_{x_{k-1}^*} \left(\hat{c}_{k-1}^{DP} - x_{k-1}^*\right) + c_{L,k}, \tag{18}$$

$$\hat{G}_k^{DP} = \left[\frac{\partial f_{k-1}}{\partial x}|_{x_{k-1}^*} \hat{G}_{k-1}^{DP}, G_{L,k}, G_w\right]. \tag{19}$$

10: Compute the corrected state set $\bar{\mathcal{Z}}_k^{DP} = \langle \bar{c}_k^{DP}, \bar{G}_k^{DP} \rangle$ in lines 11 to 16.

11:

$$\acute{c}_k^{DP} = \hat{c}_k^{DP} + \sum_{i=1}^{m} \lambda_k^{(i)} \left(y_k^{(i)^{DP}} - h^{(i)}\left(x_k^*\right) - \frac{\partial h_k^{(i)}}{\partial x}|_{x_k^*} \left(\hat{c}_k^{DP} - x_k^*\right) - c_{L,k} \right), \tag{20}$$

12:

$$\acute{G}_k^{DP} = \left[\left(I - \sum_{i=1}^{m} \lambda_k^{(i)} \frac{\partial h_k^{(i)}}{\partial x}|_{x_k^*} \right) \hat{G}_k^{DP}, \; -\lambda_k^{(1)} G_{L,k}, \ldots, -\lambda_k^{(m)} G_{L,k}, -\lambda_k^{(1)} G_p, \ldots, -\lambda_k^{(m)} G_p, -\lambda_k^{(1)} G_v^{(1)}, \ldots, -\lambda_k^{(m)} G_v^{(m)} \right]. \tag{21}$$

13: Compute the weights $\bar{\Lambda}_k^*$ as follows:

$$\bar{\Lambda}_k^* = \operatorname{argmin} B_{\lambda_k} \left\| \acute{G}_k^{DP} \right\|_F^2, \tag{22}$$

where $\bar{\Lambda}_k^* = [\lambda_k^{(1)}, \ldots, \lambda_k^{(m)}]$.

14: Reduce the order of $\hat{G}_k^{DP}$ as in [9]: $\bar{G}_k^{DP} = \downarrow_q \acute{G}_k^{DP}$.

15: $\bar{c}_k^{DP} = \acute{c}_k^{DP}$.

16: $\bar{\mathcal{Z}}_k^{DP} = \langle \bar{c}_k^{DP}, \bar{G}_k^{DP} \rangle$.

17: Update the time: $k = k + 1$.

18: **end while**

they reduce uncertainty around estimated values. The time is updated for the next time step in line 17.

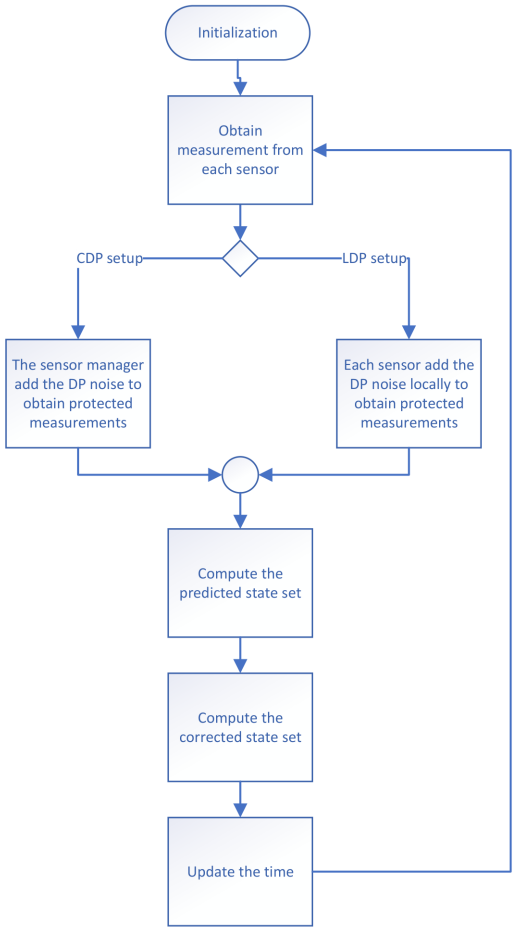

Fig. 3. Flowchart of the proposed differentially private set-based estimator.

Theorem 1: Given that the two additive noise mechanisms described in Definitions 5 and 7 are (ϵ, δ)-ADP, Algorithm 2 guarantees the true state containment in the estimated set and DP for the sensitive measurements throughout the set-based state estimation process.

Proof: The state estimates are obtained, while the sensitive measurements are kept DP-protected since the DP guarantees are preserved by postprocessing [27]. The true state containment in the estimated set follows from the boundness of the process noise, the measurement noise, and the privacy-preserving noise, which are all added as a Minkowski sum to the estimated set.■

In our work, estimation utility refers to the accuracy of the estimated state while ensuring privacy preservation. It quantifies how well the estimated set approximates the true system state after applying the privacy-preserving noise [36].

1) Estimation utility is evaluated by measuring the error between the estimated state and the true state. Specifically, we consider the following key metrics.

a) *Center Deviation:* The Euclidean distance between the center of the estimated zonotope and the true state.

b) *RMSE:*

$$\text{RMSE} = \sqrt{\frac{1}{K}\sum_{k=1}^{K} \|x_k - \hat{x}_k\|^2} \tag{23}$$

where x_k is the true state, $\hat{x}_k$ is the estimated state, and K is the number of time steps.

2) *Measurement of the Estimation Utility:*
   a) *Baseline (Without DP Noise):* We first evaluate the estimator's accuracy in the absence of privacy-preserving noise to establish an upper bound on utility.
   b) *Impact of DP Noise:* We then introduce the truncated optimal noise and truncated Laplace noise separately, comparing the resulting estimation error.

To ensure guaranteed state inclusion regardless of the actual sampled noise values, the privacy-preserving noise zonotope $\mathcal{Z}_p$, the process noise zonotope $\mathcal{Z}_w$, and the measurement noise zonotopes $\mathcal{Z}_v^{(i)}$, $i \in \{1, \ldots, m\}$, are added as a Minkowski sum to the estimated zonotope at each time step k. We guarantee the inclusion of the true state in the estimated zonotope $\bar{\mathcal{Z}}_k^{\text{DP}}$. Thus, the upper bound of the estimation error is double the radius of the estimated zonotope $\bar{\mathcal{Z}}_k^{\text{DP}}$ presented in [37]. This provides a worst case bound on the deviation from the estimated state, accounting for all sources of bounded uncertainty, including process, measurement, and privacy-preserving noise. Analytically, the estimation error at each time step k is conservatively bounded by

$$\|x_k - \hat{x}_k\|_2 \leq 2 \text{ rad}\left(\bar{\mathcal{Z}}_k^{\text{DP}}\right) \tag{24}$$

where $\text{rad}(\bar{\mathcal{Z}}_k^{\text{DP}}) = \sup_{p \in \bar{\mathcal{Z}}_k^{\text{DP}}} \|\bar{c}_k^{\text{DP}} - p\|_\infty$ and p is any point such that $p \in \bar{\mathcal{Z}}_k^{\text{DP}}$. This bound incorporates all sources of uncertainty, including the injected privacy-preserving noise. This approach guarantees verifiable overapproximation of the true state even in worst case realizations, aligning with safety-critical estimation requirements. Additionally, the proposed differentially private set-based estimator is applicable to a wide range of IoT scenarios. In smart healthcare, for instance, patient data from wearable devices can be processed in a privacy-preserving manner to ensure both the accuracy of health monitoring and the protection of sensitive physiological signals [38]. In industrial IoT systems, where sensor networks monitor processes under bounded disturbances, the proposed mechanism can secure operational data without degrading control performance. Similarly, in collaborative robotics or swarm UAVs, state sharing among agents can be protected using our approach to guarantee privacy while maintaining safe coordination. These cases illustrate the broader potential of the proposed estimator in privacy-sensitive IoT applications [39], [40].

C. Computational Complexity

In this section, we analyze the computational complexity of the proposed differentially private set-based estimator. As summarized in Table II, the proposed differentially private set-based estimator consists of two main computational components.

1) *Noise Optimization (One-Time Computation):* The optimal noise distribution is generated once using numerical optimization (e.g., SGD). This step involves solving an optimization problem to balance privacy loss and utility loss, which depends on privacy parameters (ϵ, δ), sensitivity, and noise range. The computational complexity of the optimization problem depends on the optimization algorithm used. If SGD is used, the complexity is typically $\mathcal{O}(T.2\,N.\text{Par})$, where T is the number of iterations performed, $2\,N$ is the number of discretization steps, and Par is the number of parameters being updated iteratively during the gradient descent [41].
2) *Real-Time Estimation (per Time Step Computation):*
   a) *Noise Addition:* The precomputed noise values are sampled and added to the sensor measurements. This is $\mathcal{O}(1)$ since it involves a simple lookup and addition.
   b) *Set-Based Estimation Using Zonotopes:* The estimation process consists of prediction and correction steps.
      i) *Prediction Step:* Use a first-order Taylor expansion, linear maps, and Minkowski sums for zonotopes. Zonotopic operations typically scale as $\mathcal{O}(n^2)$ for n-dimensional systems.
      ii) *Correction Step:* Find the intersection of measurement sets with the predicted state set. Zonotope intersection is approximated using overbounding techniques, which have complexity $\mathcal{O}(n^2\gamma_{\mathcal{Z}})$, where $\gamma_{\mathcal{Z}}$ is the number of generators.
      iii) *Total Complexity per Time Step:* Approximately $\mathcal{O}(n^2\gamma_{\mathcal{Z}})$, which is efficient for real-time set-based estimation.

IV. Example

This section evaluates the proposed differentially private set-based estimator in both the LDP and CDP setups visualized in Fig. 2(a) and (b), respectively.

A. Experimental Settings

The experimental results are presented based on the implementation in MATLAB R2023b, with zonotope operations performed using the CORA toolbox [42]. We evaluate the differentially private set-based estimator through the localization of a quadcopter navigating at a speed of 0.35 m/s through arbitrary nonlinear motion within 3-D space measuring $10 \times 10 \times 10$ m^3. According to the following model, the anchor nodes provide measurements as relative distances to the intruding quadcopter. The proposed differentially private set-based estimator can handle the bounded measurement uncertainties and anonymize these measurements with privacy-preserving noise, thus concealing the exact locations of the anchor nodes from the cloud estimator or any untrusted party. This localization is achieved using a high rate of real-world data measurements provided by a set of 8 anchor nodes distributed across the motion area [43]. The simulation was conducted using these measurements provided over a total duration of 1443 time steps, where each time step was set

TABLE II
COMPUTATIONAL COMPLEXITY ANALYSIS OF THE PROPOSED DIFFERENTIALLY PRIVATE SET-BASED ESTIMATOR

Component	Complexity	Impact
Noise Optimization (One-time)	$\mathcal{O}(T.2N.Par)$	Happens once before deployment, does not affect real-time performance.
Noise Addition (Per Step)	$\mathcal{O}(1)$	Negligible computational cost.
Set-Based Estimation (Per Step)	$\mathcal{O}(n^2\gamma_{\mathcal{Z}})$	Efficient for real-time state estimation.

to 0.00693 s. The model of this system incorporates a linear state function represented by the matrix f, which is given by

$$f = \text{diag}\left([1\ 1\ 1]^{\mathrm{T}}\right).$$

Additionally, the measurement function $h^{(i)}(x_k)$ is defined as

$$h^{(i)}(x_k) = \left\| x_A^{(i)} - x_k \right\|_2$$

where $x_A^{(i)}$ represents the location of anchor node i and x_k is an estimate for the location of the quadcopter at time step k. The measurement and process noise zonotopes $\mathcal{Z}_v^{(i)}$ and $\mathcal{Z}_w$, respectively, are set to

$$z_v^{(i)} = \left\langle 0, \left[0.01 \quad 0.02 \quad 0.01\right]\right\rangle$$
$$z_w = \left\langle [0\ 0\ 0]^{\mathrm{T}}, \text{diag}\left([0.50\ 0.50\ 0.50]^{\mathrm{T}}\right)\right\rangle.$$

Additionally, Fig. 4 shows the true values, upper bound, and lower bound for each dimension of the estimated state using the proposed differentially private set-based estimator within the context of the CDP setup. Also, Fig. 5 shows the true values, upper bound, and lower bound for each dimension of the estimated state using the proposed differentially private set-based estimator within the context of the LDP setup.

B. Evaluation Scope

In this experimental study, our primary objective is to evaluate the impact of two differentially private noise mechanisms, i.e., the truncated Laplace noise and the numerically optimized (ϵ, δ)-ADP truncated optimal noise, on the estimation utility within the proposed differentially private set-based estimator. The estimation utility is assessed through localization accuracy metrics, including the RMSE and center deviation of the estimated state zonotopes. Additionally, we include a DP version of the extended Kalman filter (DP-EKF) in [44] in the comparison to substantiate the minimized utility loss. Since the proposed estimator employs zonotopes for set representation, we also demonstrate their computational advantages in terms of efficient Minkowski sum and linear mapping operations during set prediction and correction steps. Compared to traditional set representations such as polytopes or ellipsoids, zonotopes enable scalable, real-time estimation performance in high-frequency multisensor settings. This comparative analysis and computational efficiency validation form the basis of the experimental results presented in Sections IV-D and IV-E.

C. Adversary Model

In the experimental evaluation, we assume an adversary model consistent with the problem formulation in Section II-D. Specifically, the untrusted cloud estimator is considered as a passive adversary that receives the differentially private state estimates and attempts to infer sensitive sensor measurements or their locations. In the CDP model, the trusted sensor manager perturbs the measurements before forwarding them to the cloud, while in the LDP model, each sensor independently applies the privacy-preserving mechanism before transmission. This adversary model represents typical inference threats in privacy-preserving state estimation scenarios and is used as the basis for evaluating the privacy–utility tradeoff in our simulations.

D. CDP Model

In this section, we evaluate the differentially private set-based estimator within the context of the CDP setup. The (ϵ, δ)-ADP truncated optimal noise distribution (Definition 3) is generated according to the noise model in (15) with $C = 500$. We set $d \in [-1,\ 1]$, $\epsilon = 0.3$, and $s = 1$. The training process was configured to run for 15 000 epochs with a learning rate of 0.01, after which the algorithm terminates. Additionally, we set $v = 10\,000$, $\gamma = 1$, $\Omega_{\text{start}} = 0.5$, $\Omega_{\min} = 0.0000001$, $\Gamma = 1.0$, and $N = 5000$. The parameters A, B_j, F_j, and δ are learned to optimize the loss function in (14) using the SGD tool [21]. We conducted a sensitivity analysis across a set of ϵ values shown in Table III. We found that $\epsilon = 0.3$ provides a suitable tradeoff, ensuring a high degree of privacy while maintaining acceptable estimation utility (i.e., low RMSE).

Fig. 6 represents a random snapshot for the localization of the quadcopter. We notice that the quadcopter is enclosed by the estimated zonotope, which indicates that the state containment is still guaranteed. Also, the center of the estimated zonotope is very close to the quadcopter, which is a good utility indicator. In industrial applications, for a certain privacy budget ϵ, the selection of the optimal noise range d, illustrated in Table III, should be guided by the acceptable error ranges. Also, Fig. 4 shows the true values, upper bound, and lower bound for each dimension of the estimated state using the proposed differentially private set-based estimator within the context of the CDP setup. For comparison, we consider the work in [19], where a differentially private interval estimator deploying truncated Laplace noise is presented. The truncated Laplace noise range, for a given ϵ, δ, and sensitivity s, is determined by

$$a = \frac{s}{\epsilon} \ln\left(1 + e^{\epsilon} \frac{1 - e^{-\epsilon}}{2\delta}\right). \tag{25}$$

Fig. 7 indicates that the truncated Laplace noise needed to achieve a certain δ is wider than the (ϵ, δ)-ADP truncated optimal noise needed to achieve the same δ. Indeed, for a

TABLE III

OPTIMAL δ VALUES CORRESPONDING TO DIFFERENT RANGES $d(m)$ OF (ϵ, δ)-ADP OPTIMAL NOISE AT $\epsilon = 0.1, 0.3, 0.5, 0.7$. THE NOISE DISTRIBUTIONS WERE GENERATED USING THE PUBLIC IMPLEMENTATION FROM [21], WITH TRAINING PARAMETERS SET AS $s = 1$, $\nu = 10\,000$, $\gamma = 1$, $\Omega_{\text{START}} = 0.5$, $\Omega_{\min} = 1 \times 10^{-7}$, $\Gamma = 1.0$, $N = 5000$, LEARNING RATE = 0.01, AND 15 000 TRAINING EPOCHS. THESE VALUES CAN BE INDEPENDENTLY REPRODUCED USING THE CITED IMPLEMENTATION AND PARAMETERS

ϵ / **Noise Range** $d(m)$	3	5	7	9	11	13	15
0.1	0.1502	0.0811	0.0518	0.0360	0.0262	0.0197	0.0151
0.3	0.1198	0.0503	0.0244	0.0126	0.0067	0.0036	0.0020
0.5	0.0931	0.0290	0.0101	0.0036	0.0013	0.0005	0.0002
0.7	0.0707	0.0158	0.0038	0.0009	0.0002	$5.64e^{-5}$	$1.39e^{-5}$

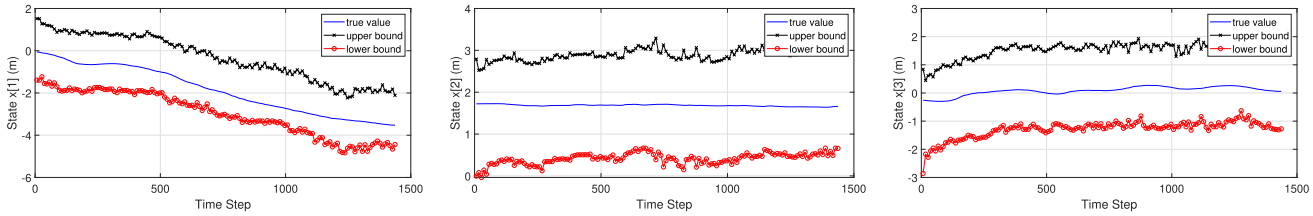

Fig. 4. True values, upper bounds, and lower bounds of the 3-D estimated states, i.e., (a) $x[1]$, (b) $x[2]$, and (c) $x[3]$, using the differentially private set-based estimator within the context of the CDP setup.

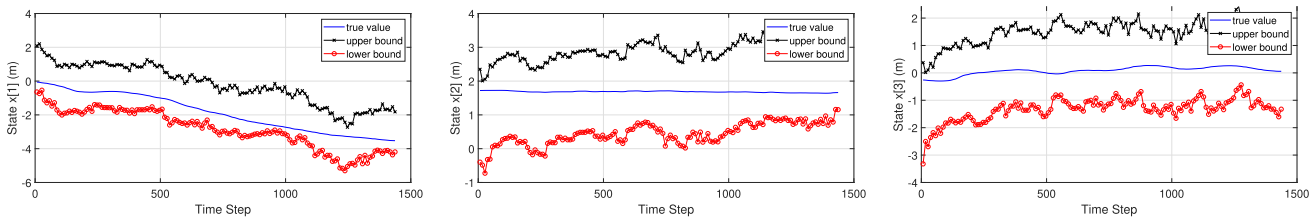

Fig. 5. True values, upper bounds, and lower bounds of the 3-D estimated states, i.e., (a) $x[1]$, (b) $x[2]$, and (c) $x[3]$, using the differentially private set-based estimator within the context of the LDP setup.

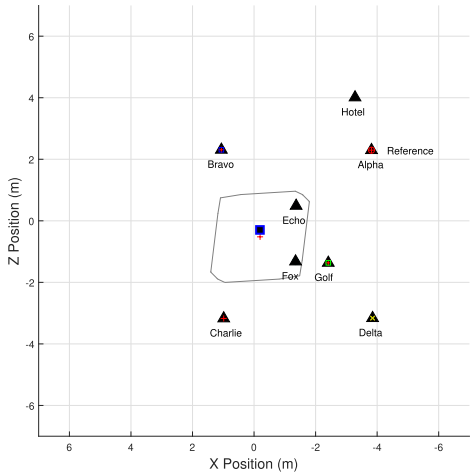

Fig. 6. Localization of a quadcopter navigating through arbitrary nonlinear motion, ■: the quadcopter, ▲: anchor nodes, and +: center of the estimated zonotope.

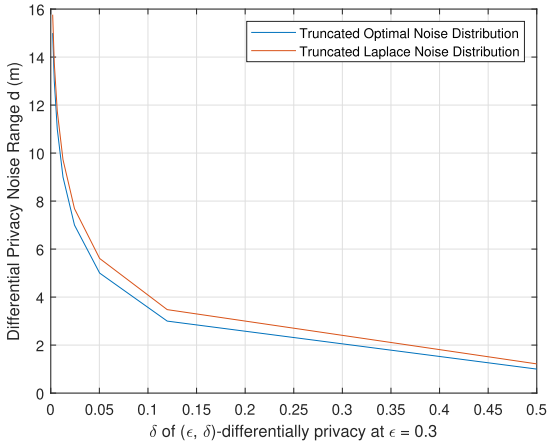

Fig. 7. Comparison of δ values with needed range of both (ϵ, δ)-ADP truncated optimal noise and truncated Laplace noise.

certain privacy budget ϵ, learning the truncated optimal noise distribution $P(\phi_l)$ using the SGD tool in [21] allows our proposed differentially private set-based estimator (Algorithm 2) to minimize loss of privacy and utility simultaneously. Hence, at a certain δ value, we find that the truncated Laplace noise causes a higher utility loss than the (ϵ, δ)-ADP truncated optimal noise.

In particular, we calculate the error in the estimated location as the distance between the estimated zonotope's center and the quadcopter's true location. Then, we compare the utility

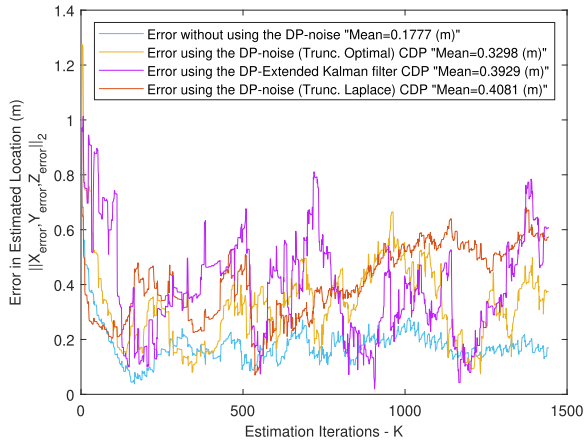

Fig. 8. Comparison of the impact of (ϵ, δ)-ADP truncated optimal noise, DP-EKF, and truncated Laplace noise on the error in the estimated location of the quadcopter at $\epsilon = 0.3$, within the CDP setup.

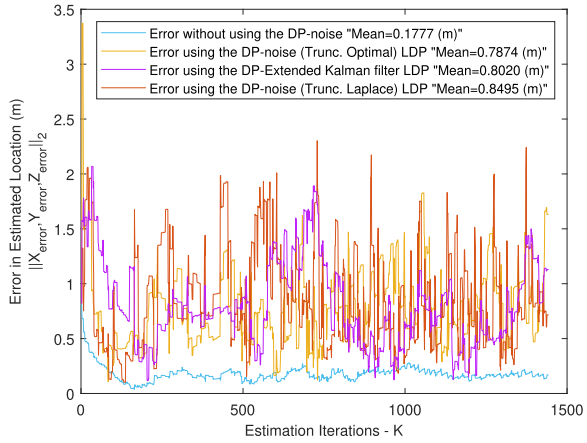

Fig. 9. Comparison of the impact of (ϵ, δ)-ADP truncated optimal noise, DP-EKF, and truncated Laplace noise on the error in the estimated location of the quadcopter at $\epsilon = 0.3$, within the LDP setup.

loss, represented by the second norm of the error in the estimated location of the quadcopter, associated with the (ϵ, δ)-ADP truncated optimal noise and the truncated Laplace noise [19] at $\epsilon = 0.3$. The results presented in Fig. 8 demonstrate that the utility loss incurred when employing the (ϵ, δ)-ADP truncated optimal noise is lower than that observed with the truncated Laplace noise and with a DP-EKF at a certain privacy level. Based on the comparative results, the proposed differentially private set-based estimator employs (ϵ, δ)-ADP truncated optimal noise, demonstrates reduced utility loss compared to using truncated Laplace noise [19], and outperforms the DP-EKF in the context of the CDP setup.

E. LDP Model

In this section, we evaluate the proposed differentially private set-based estimator within the context of the LDP setup. It deserves mentioning that the sensor manager is not present to aggregate the measurements of the sensors since it is treated as an untrusted entity in the LDP setup. Instead, each sensor locally perturbs its own measurement with the privacy-preserving noise. As in the CDP setup of this example, the (ϵ, δ)-ADP truncated optimal noise distribution is regenerated for the LDP setup. Fig. 5 shows the true values, upper bound, and lower bound for each dimension of the estimated state using the proposed differentially private set-based estimator within the context of the LDP setup. In the context of the LDP setup, we again compare the utility loss of (ϵ, δ)-ADP truncated optimal noise and truncated Laplace noise [19] at $\epsilon = 0.3$. The comparative results in Fig. 9 demonstrate that the proposed differentially private set-based estimator, which utilizes (ϵ, δ)-ADP truncated optimal noise, incurs less utility loss compared to using truncated Laplace noise [19], and outperforms the DP-EKF in the context of the LDP setup.

Furthermore, our proposed differentially private set-based estimator (Algorithm 2) utilizes zonotopes for set representation. Zonotopes offer a less conservative set representation that enables efficient computation of linear maps and Minkowski sums, both essential operations in the set-based estimation process. Consequently, this confers a computational advantage to the proposed differentially private set-based estimator.

V. Conclusion

We have proposed a differentially private set-based estimator that performs set-based estimation in a privacy-preserving manner, ensuring that an adversary cannot infer the actual values of sensitive measurements from the state estimates. The proposed estimator safeguards the privacy of measurement sets throughout the estimation process while maintaining minimal utility loss.

Comprehensive simulation results demonstrated that, for a given value of δ, a significantly wider range of truncated Laplace noise is required compared to the (ϵ, δ)-ADP truncated optimal noise distribution. This highlights the improved efficiency and reduced utility degradation of the proposed noise mechanism. Furthermore, under both CDP and LDP models, the proposed estimator consistently achieved lower RMSEs in state estimation compared to its truncated Laplace counterpart at equivalent privacy budgets.

Additionally, employing zonotopes for set representation in the proposed estimator provided a computational advantage over other set representations, owing to their favorable closure properties under linear transformations and Minkowski sum operations. These advantages collectively position the proposed mechanism as a practical and effective solution for privacy-preserving set-based estimation in IoT and cyber-physical system applications. As future work, we plan to extend our analysis to more complex threat models, including scenarios involving active adversaries such as sensor-cloud collusion. Additionally, we aim to investigate the impact of sequential (ϵ, δ)-composition across multiple inference rounds, which is particularly relevant to distributed or multiround estimation scenarios. Exploring these directions will further enhance the robustness and applicability of the proposed privacy-preserving set-based estimation framework.

References

[1] D. Hahn, A. Munir, and V. Behzadan, "Security and privacy issues in intelligent transportation systems: Classification and challenges," *IEEE Intell. Transp. Syst. Mag.*, vol. 13, no. 1, pp. 181–196, Spring 2021.

[2] M. E. Andrés, N. E. Bordenabe, K. Chatzikokolakis, and C. Palamidessi, "Geo-indistinguishability: Differential privacy for location-based systems," in *Proc. ACM SIGSAC Conf. Comput. Commun. Security*, 2013, pp. 901–914.

[3] C. Y. T. Ma, D. K. Y. Yau, N. K. Yip, and N. S. V. Rao, "Privacy vulnerability of published anonymous mobility traces," *IEEE/ACM Trans. Netw.*, vol. 21, no. 3, pp. 720–733, Jun. 2013.

[4] A. Alanwar, Y. Shoukry, S. Chakraborty, P. Martin, P. Tabuada, and M. Srivastava, "PrOLoc: Resilient localization with private observers using partial homomorphic encryption," in *Proc. 16th ACM/IEEE Int. Conf. Inf. Process. Sensor Netw.*, Apr. 2017, pp. 41–52.

[5] P. Bouron, D. Meizel, and P. Bonnifait, "Set-membership non-linear observers with application to vehicle localisation," in *Proc. Eur. Control Conf. (ECC)*, Sep. 2001, pp. 1255–1260.

[6] M. Althoff, O. Stursberg, and M. Buss, "Computing reachable sets of hybrid systems using a combination of zonotopes and polytopes," *Nonlinear Anal., Hybrid Syst.*, vol. 4, no. 2, pp. 233–249, May 2010.

[7] A. A. Kurzhanskiy and P. Varaiya, "Ellipsoidal techniques for reachability analysis of discrete-time linear systems," *IEEE Trans. Autom. Control*, vol. 52, no. 1, pp. 26–38, Jan. 2007.

[8] J. K. Scott, D. M. Raimondo, G. R. Marseglia, and R. D. Braatz, "Constrained zonotopes: A new tool for set-based estimation and fault detection," *Automatica*, vol. 69, pp. 126–136, Jul. 2016.

[9] A. Girard, "Reachability of uncertain linear systems using zonotopes," in *Hybrid Systems: Computation and Control*. Berlin, Germany: Springer, 2005, pp. 291–305.

[10] B. Zhang, X. Sun, S. Liu, and X. Deng, "Tracking control of multiple unmanned aerial vehicles incorporating disturbance observer and model predictive approach," *Trans. Inst. Meas. Control*, vol. 42, no. 5, pp. 951–964, Mar. 2020, doi: 10.1177/0142331219879858.

[11] A. Alanwar, A. Berndt, K. H. Johansson, and H. Sandberg, "Data-driven set-based estimation using matrix zonotopes with set containment guarantees," in *Proc. Eur. Control Conf. (ECC)*, Jul. 2022, pp. 875–881.

[12] J. Kim and H. Shim, "Encrypted state estimation in networked control systems," in *Proc. IEEE 58th Conf. Decis. Control (CDC)*, Dec. 2019, pp. 7190–7195.

[13] Z. Zhang, P. Cheng, J. Wu, and J. Chen, "Secure state estimation using hybrid homomorphic encryption scheme," *IEEE Trans. Control Syst. Technol.*, vol. 29, no. 4, pp. 1704–1720, Jul. 2021.

[14] S. Emad, A. Alanwar, Y. Alkabani, M. W. El-Kharashi, H. Sandberg, and K. H. Johansson, "Privacy guarantees for cloud-based state estimation using partially homomorphic encryption," in *Proc. Eur. Control Conf. (ECC)*, Jul. 2022, pp. 98–105.

[15] Y. Zhao et al., "Local differential privacy-based federated learning for Internet of Things," *IEEE Internet Things J.*, vol. 8, no. 11, pp. 8836–8853, Jun. 2021.

[16] J. Le Ny and G. J. Pappas, "Differentially private filtering," *IEEE Trans. Autom. Control*, vol. 59, no. 2, pp. 341–354, Feb. 2014.

[17] K. H. Degue and J. Le Ny, "On differentially private Kalman filtering," in *Proc. IEEE Global Conf. Signal Inf. Process. (GlobalSIP)*, Nov. 2017, pp. 487–491.

[18] J. Le Ny, "Differentially private nonlinear observer design using contraction analysis," *Int. J. Robust Nonlinear Control*, vol. 30, no. 11, pp. 4225–4243, 2020.

[19] K. H. Degue and J. Le Ny, "Differentially private interval observer design with bounded input perturbation," in *Proc. Amer. Control Conf. (ACC)*, Jul. 2020, pp. 1465–1470.

[20] M. M. Dawoud, C. Liu, A. Alanwar, and K. H. Johansson, "Differentially private set-based estimation using zonotopes," in *Proc. Eur. Control Conf. (ECC)*, Jun. 2023, pp. 1–8.

[21] D. M. Sommer, L. Abfalterer, S. Zingg, and E. Mohammadi, "Learning numeric optimal differentially private truncated additive mechanisms," 2021, *arXiv:2107.12957*.

[22] G. M. Ziegler, *Lectures on Polytopes*. New York, NY, USA: Springer-Verlag, 1995.

[23] M. Althoff, "*Reachability analysis and its application to the safety assessment of autonomous cars*," Ph.D. thesis, Fac. Elect. Eng. Inf. Technol., Technische Universität München, Munich, Germany, 2010.

[24] C. Combastel, "A state bounding observer for uncertain non-linear continuous-time systems based on zonotopes," in *Proc. 44th IEEE Conf. Decis. Control*, Dec. 2005, pp. 7228–7234.

[25] A. Alanwar et al., "Privacy-preserving set-based estimation using partially homomorphic encryption," *Eur. J. Control*, vol. 71, May 2023, Art. no. 100786.

[26] A. Alanwar, F. J. Jiang, M. Sharifi, D. V. Dimarogonas, and K. H. Johansson, "Enhancing data-driven reachability analysis using temporal logic side information," in *Proc. Int. Conf. Robot. Autom. (ICRA)*, May 2022, pp. 6793–6799.

[27] C. Dwork and A. Roth, "The algorithmic foundations of differential privacy," *Found. Trends Theor. Comput. Sci.*, vol. 9, nos. 3–4, pp. 211–407, 2014.

[28] C. Dwork, "Differential privacy," in *Automata, Languages and Programming*. Berlin, Germany: Springer, 2006, pp. 1–12.

[29] B. Jiang, J. Li, G. Yue, and H. Song, "Differential privacy for industrial Internet of Things: Opportunities, applications, and challenges," *IEEE Internet Things J.*, vol. 8, no. 13, pp. 10430–10451, Jul. 2021.

[30] F. Liu, "Generalized Gaussian mechanism for differential privacy," *IEEE Trans. Knowl. Data Eng.*, vol. 31, no. 4, pp. 747–756, Apr. 2019.

[31] W. Croft, J.-R. Sack, and W. Shi, "Differential privacy via a truncated and normalized Laplace mechanism," *J. Comput. Sci. Technol.*, vol. 37, no. 2, pp. 369–388, Apr. 2022.

[32] J. Le Ny and G. J. Pappas, "Differentially private Kalman filtering," in *Proc. 50th Annu. Allerton Conf. Commun., Control, Comput. (Allerton)*, Oct. 2012, pp. 1618–1625.

[33] M. Abadi et al., "Deep learning with differential privacy," in *Proc. ACM SIGSAC Conf. Comput. Commun. Secur.*, New York, NY, USA, 2016, pp. 308–318.

[34] G. Muthukrishnan and S. Kalyani, "Grafting Laplace and Gaussian distributions: A new noise mechanism for differential privacy," *IEEE Trans. Inf. Forensics Security*, vol. 18, pp. 5359–5374, 2023.

[35] A. Alanwar, J. J. Rath, H. Said, K. H. Johansson, and M. Althoff, "Distributed set-based observers using diffusion strategies," *J. Franklin Inst.*, vol. 360, no. 10, pp. 6976–6993, Jul. 2023.

[36] Q. Geng, W. Ding, R. Guo, and S. Kumar, "Tight analysis of privacy and utility tradeoff in approximate differential privacy," in *Proc. 23rd Int. Conf. Artif. Intell. Statist.* (Proceedings of Machine Learning Research), vol. 108, S. Chiappa and R. Calandra, Eds., Palermo, Italy, Aug. 2020, pp. 89–99.

[37] M. U. B. Niazi, M. S. Chong, A. Alanwar, and K. H. Johansson, "Secure set-based state estimation for linear systems under adversarial attacks on sensors," 2024, *arXiv:2309.05075v1*.

[38] J. Zhang, X. Liang, Z. Zhang, S. He, and Z. Shi, "Re-DPoctor: Real-time health data releasing with W-day differential privacy," in *Proc. IEEE Global Commun. Conf. (GLOBECOM)*, Dec. 2017, pp. 1–6.

[39] A.-R. Sadeghi, C. Wachsmann, and M. Waidner, "Security and privacy challenges in industrial Internet of Things," in *Proc. 52nd ACM/EDAC/IEEE Design Autom. Conf. (DAC)*, New York, NY, USA, Jun. 2015, pp. 1–6.

[40] M. Kumar et al., "A smart privacy preserving framework for industrial IoT using hybrid meta-heuristic algorithm," *Sci. Rep.*, vol. 13, no. 1, p. 5372, Apr. 2023.

[41] R. Babanezhad, *Lecture Notes on Optimization Methods and Algorithms*. Stanford, CA, USA: Stanford Univ., 2015.

[42] M. Althoff, "An introduction to CORA," in *Proc. Workshop Appl. Verification Continuous Hybrid Syst.*, 2015, pp. 120–151.

[43] A. Alanwar et al., "D-SLATS: Distributed simultaneous localization and time synchronization," in *Proc. 18th ACM Int. Symp. Mobile Ad Hoc Netw. Comput.*, Jul. 2017, pp. 1–10.

[44] K. Fujii, "Extended Kalman filter," *Reference Manual*, vol. 14, p. 41, 2013.