Secure Set-based State Estimation for Safety-Critical Applications under Adversarial Attacks on Sensors *

M. Umar B. Niazi[1] Michelle S. Chong[2] Amr Alanwar[3]
Karl H. Johansson[1]

Abstract

Set-based state estimation provides guaranteed state inclusion certificates that are crucial for the safety verification of dynamical systems. However, when system sensors are subject to cyberattacks, maintaining both safety and security guarantees becomes a fundamental challenge that existing point-based secure state estimation methods cannot adequately address due to their inherent inability to provide state inclusion certificates. This paper introduces a novel approach that simultaneously ensures safety guarantees through guaranteed state inclusion and security guarantees against sensor attacks, without imposing conservative restrictions on system operation. We propose a Secure Set-based State Estimation (S3E) algorithm that maintains the true system state within the estimated set under sensor attacks, provided the initialization set contains the initial state and the system remains observable from the uncompromised sensor subset. The algorithm gives the estimated set as a collection of constrained zonotopes (agreement sets), which can be employed as robust certificates for verifying whether the system adheres to safety constraints. Furthermore, we demonstrate that the estimated set remains unaffected by attack signals of sufficiently large magnitude and also establish sufficient conditions for attack detection, identification, and filtering. This compels the attacker to only inject signals of small magnitudes to evade detection, thus preserving the accuracy of the estimated set. To address the computational complexity of the algorithm, we offer several strategies for complexity-performance trade-offs. The efficacy of the proposed algorithm is illustrated through several examples, including its application to a three-story building model.

1 Introduction

Secure state estimation is critical for the reliable operation of cyber-physical systems, where control performance and safety depend directly on accurate state knowledge. Adversarial manipulation of sensor networks through false data injection can compromise system stability and lead to operational failures [35]. This vulnerability presents significant challenges for critical infrastructure systems, including power grids, water distribution networks, autonomous vehicles, and industrial automation [18].

*This work is supported by the Swedish Research Council and the Knut and Alice Wallenberg Foundation, Sweden. It has also received funding from the European Union's Horizon Research and Innovation Programme under grant agreement No. 830927 and Marie Skłodowska-Curie grant agreement No. 101062523.
[1] Division of Decision and Control Systems, Digital Futures, KTH Royal Institute of Technology, SE-100 44 Stockholm, Sweden. Emails: mubniazi@kth.se, kallej@kth.se
[2] Department of Mechanical Engineering, Eindhoven University of Technology, the Netherlands. Email: m.s.t.chong@tue.nl
[3] School of Computation, Information and Technology, Technical University of Munich, Germany. Email: alanwar@tum.de

Modern control systems exhibit inherent vulnerabilities due to their distributed architectures and dependence on networked communication infrastructure. These characteristics enable adversaries to manipulate sensor measurements through false data injection attacks, creating divergence between reported and actual system states. Research demonstrates that sophisticated attacks can maintain stealth characteristics while progressively degrading system performance [37, 42]. The evolution of such attack strategies, coupled with increased system connectivity, drives the need for resilient estimation algorithms capable of maintaining state accuracy under sensor compromise.

Traditional approaches to secure state estimation have focused on point-based methods that produce single estimates of the system state at each time instant. These algorithms typically employ redundant sensors and leverage techniques such as majority voting or optimization-based outlier detection to identify and exclude corrupted measurements [14, 11, 20, 31, 16, 3, 9]. While these methods have shown promise in specific scenarios, they suffer from fundamental limitations that restrict their applicability. Most notably, point-based secure estimators require that strictly fewer than half of the sensors remain uncompromised. This constraint may be violated by adversaries with sufficient resources to coordinate large-scale attacks. Furthermore, these methods produce conservative error bounds that limit their practical applicability.

Recent efforts to address these limitations have explored the integration of cryptographic authentication mechanisms to periodically verify the integrity of sensor data [23, 19]. However, the computational overhead associated with cryptographic operations creates challenging tradeoffs between security guarantees and the real-time performance requirements of control systems. This overhead becomes particularly problematic in resource-constrained embedded systems where computational capacity must be carefully allocated between control tasks and security operations.

1.1 Set-based Estimation Framework

Beyond merely detecting attacks or providing state estimates, safety-critical systems require formal verification that the system remains within safe operating bounds even under sensor attacks. Point-based methods fail to address this verification challenge as they cannot represent the complete space of feasible system states given measurement data and attack constraints without excessive conservatism. This limitation becomes critical in industrial processes operating with minimal safety margins, where minor estimation uncertainties may result in constraint violations.

Set-based estimation techniques, particularly those employing zonotopic representations, offer a fundamentally different approach to secure state estimation that addresses many limitations of point-based methods while naturally supporting safety verification [1, 21, 34, 6, 5, 13, 2, 32]. These methods compute bounded sets containing all feasible system states rather than single estimates, yielding inclusion guarantees necessary for constraint verification. Moreover, set-based estimation methods naturally embed uncertainties and noise in cases where their bounds are known. The zonotopic representation balances computational efficiency with tight uncertainty bounds, making real-time implementation feasible.

The key advantage of set-based methods for safety verification lies in their ability to propagate uncertainty through system dynamics while maintaining guaranteed bounds. These methods compute reachable sets encompassing all trajectories consistent with measurements and bounded disturbances, enabling verification of safety constraint satisfaction. This capability proves essential against stealthy attacks that evade immediate detection while gradually compromising system safety. Applications spanning fault diagnosis in industrial systems [7], underwater robotics [17], vehicle localization [8], and water network leak detection [33] demonstrate the practical value of guaranteed uncertainty bounds for operational safety and

performance verification.

1.2 Current Limitations in Secure Set-Based Estimation

Despite their theoretical advantages for safety verification, existing set-based methods for secure state estimation face practical challenges. The reachability-based approach of [36] imposes that any subset of uncompromised sensors can measure the full state vector. Without such strong observability conditions, the estimation bounds become excessively conservative, potentially encompassing regions too large for effective control or meaningful safety verification. Similarly, our previous work [30] on secure state estimation considered restrictive assumptions requiring observability from every individual sensor and was limited to systems that were bounded input-bounded state stable, which significantly constrains the class of systems that can be analyzed.

Alternative approaches that attempt to identify and filter compromised sensors before applying standard zonotopic filters [28, 10, 24] fail to address stealthy attacks. These attacks inject strategically designed signals that remain within expected bounds at each time step while causing significant cumulative deviation in state estimates. Although recent work has begun to address stealthy attacks [26, 25, 39, 43, 27, 45], existing methods either restrict their analysis to specific attack strategies or assume known bounds on attack signal magnitudes. Sophisticated adversaries can readily violate these assumptions.

Furthermore, the fundamental assumption of requiring a majority of uncompromised sensors remains unaddressed in most set-based approaches. This assumption becomes increasingly problematic as systems scale and adversaries gain greater capabilities to compromise distributed sensor networks. These limitations necessitate novel algorithms that leverage the geometric properties of set-based representations to maintain state estimation accuracy and safety guarantees under extensive sensor compromise.

1.3 Contributions of This Work

This paper presents a Secure Set-Based State Estimation (S3E) algorithm that provides guaranteed state inclusion under adversarial sensor attacks while fundamentally relaxing the assumptions required by existing methods. Our approach leverages the redundancy inherent in zonotopic representations to maintain accurate state estimates and verify safety constraints even when all but one sensor can be compromised, requiring only that the system remains observable through any combination of remaining safe sensors.

The S3E algorithm executes four sequential operations per time step. The time update computes feasible states based on system dynamics and bounded process noise. Measurement updates generate state sets consistent with sensor observations across different sensor subsets. Agreement sets form through strategic intersections of measurement updates, leveraging redundancy to exclude compromised sensor data. The final estimated set comprises the union of valid agreement sets, ensuring true state inclusion and enabling safety verification.

Our theoretical analysis establishes several key properties of the S3E algorithm with direct implications for safety verification. We derive explicit lower bounds on attack signal magnitudes that ensure detection at different stages of the estimation process, demonstrating that adversaries face fundamental constraints in designing attacks that compromise safety without detection. We characterize the conditions under which attacked sensors can be identified and filtered, providing system designers with quantitative metrics for evaluating security and safety margins. Additionally, we develop computational methods to verify estimated set containment within safety constraints, supporting real-time monitoring under active attacks.

1.4 Outline of the Paper

The rest of the paper is organized as follows. Section 2 provides the required preliminaries and Section 3 states the main assumptions and the problem. Section 4 presents the S3E algorithm, provides the inclusion guarantees, and discusses methods to reduce the algorithm's complexity. Section 5 provides lower bounds on the attacks that can be detected, identified, and/or filtered. Section 6 demonstrates the proposed algorithm through simulation examples. Finally, Section 7 provides the concluding remarks.

2 Notation and Preliminaries

The set of real numbers and integers are denoted by $\mathbb{R}$ and $\mathbb{Z}$, respectively. We let $\mathbb{Z}_{\geq i} := \{i, i+1, i+2, \dots\}$ and $\mathbb{Z}_{[i,j]} := \{i, i+1, i+2, \dots, j\}$ for $j \geq i$. The Euclidean and maximum norms of a vector $x \in \mathbb{R}^n$ are denoted as $\|x\|$ and $\|x\|_\infty$, respectively. The vector and matrix of zeros are denoted as $0_n \in \mathbb{R}^n$ and $0_{n\times k} \in \mathbb{R}^{n\times k}$. The identity matrix is $I_n \in \mathbb{R}^{n\times n}$. For a finite set $\mathcal{S}$, $|\mathcal{S}|$ denotes its cardinality. The Cartesian product between two sets $\mathcal{S}_1$ and $\mathcal{S}_2$ is denoted by $\mathcal{S}_1 \times \mathcal{S}_2$.

Given a center $c_z \in \mathbb{R}^n$ and a generator $G_z \in \mathbb{R}^{n\times \xi_z}$, a *zonotope* $\mathcal{Z} = \langle c_z, G_z\rangle$ is a set

$$\mathcal{Z} := \{c_z + G_z\beta_z : \beta_z \in [-1, 1]^{\xi_z}\}$$

where ξ_z is the number of generators of $\mathcal{Z}$. That is, a zonotope is an affine transformation of a unit hypercube $\mathcal{H}(0_n, 1) = [-1, 1]^{\xi_z}$ centered at 0_n and with radius 1, where $\xi_z \in \mathbb{Z}_{\geq 1}$ is the dimension of the hypercube.

A matrix $L \in \mathbb{R}^{n'\times n}$ multiplied with a zonotope $\mathcal{Z}$ yields a linearly transformed zonotope $L\mathcal{Z} = \langle Lc_z, LG_z\rangle$. Given two zonotopes $\mathcal{Z}_1 = \langle c_{z_1}, G_{z_1}\rangle$ and $\mathcal{Z}_2 = \langle c_{z_2}, G_{z_2}\rangle$ in $\mathbb{R}^n$, their Minkowski sum is given by

$$\mathcal{Z}_1 \oplus \mathcal{Z}_2 = \langle c_{z_1} + c_{z_2}, [\, G_{z_1} \;\; G_{z_2} \,]\rangle.$$

The Cartesian product of two zonotopes is given by

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

A *constrained zonotope* $\mathcal{Z} = \langle c_z, G_z, A_z, b_z\rangle$ is a set

$$\mathcal{Z} := \{c_z + G_z\beta_z : \beta_z \in [-1, 1]^{\xi_z}, A_z\beta_z = b_z\}$$

where $c_z \in \mathbb{R}^n$, $G_z \in \mathbb{R}^{n\times\xi_z}$, $A_z \in \mathbb{R}^{m\times\xi_z}$ and $b_z \in \mathbb{R}^m$ with $m, n \in \mathbb{Z}_{\geq 1}$. That is, a constrained zonotope is an affine transformation of the linearly constrained unit hypercube $\{\beta_z \in [-1, 1]^{\xi_z} : A_z\beta_z = b_z\}$.

Given two constrained zonotopes $\mathcal{Z} \subset \mathbb{R}^{n_z}$ and $\mathcal{Y} \subset \mathbb{R}^{n_y}$, and a matrix $M \in \mathbb{R}^{n_y\times n_z}$, the *generalized intersection* is defined as

$$\mathcal{Z} \cap_M \mathcal{Y} := \{z \in \mathcal{Z} : Mz \in \mathcal{Y}\}. \tag{1}$$

Let $\mathcal{Z} = \langle c_z, G_z, A_z, b_z\rangle \subset \mathbb{R}^{n_z}$ with $c_z \in \mathbb{R}^{n_z}$, $G_z \in \mathbb{R}^{n_z\times\xi_z}$, $A_z \in \mathbb{R}^{m_z\times\xi_z}$, and $b_z \in \mathbb{R}^{m_z}$, and let $\mathcal{Y} = \langle c_y, G_y, A_y, b_y\rangle \subset \mathbb{R}^{n_y}$ with $c_y \in \mathbb{R}^{n_y}$, $G_y \in \mathbb{R}^{n_y\times\xi_y}$, $A_y \in \mathbb{R}^{m_y\times\xi_y}$, and $b_y \in \mathbb{R}^{m_y}$, then the generalized intersection can be computed as, see [34],

$$\mathcal{Z} \cap_M \mathcal{Y} = \left\langle c_z, \begin{bmatrix} G_z & 0_{n_z\times\xi_y} \end{bmatrix}, \begin{bmatrix} A_z & 0_{m_z\times\xi_y} \\ 0_{m_y\times\xi_z} & A_y \\ MG_z & -G_y \end{bmatrix}, \begin{bmatrix} b_z \\ b_y \\ c_y - Mc_z \end{bmatrix} \right\rangle. \tag{2}$$

The radius of a zonotope, or a constrained zonotope, $\mathcal{Z} \subset \mathbb{R}^n$ is defined as

$$\mathrm{rad}(\mathcal{Z}) := \min \Delta \text{ subject to } \mathcal{Z} \subseteq \mathcal{H}(c_z, \Delta)$$

where Δ is the radius of a minimal hypercube of dimension n denoted by $\mathcal{H}(c_z, \Delta) := \langle c_z, \Delta I_n \rangle$, which is centered at c_z and inscribes $\mathcal{Z}$. Notice that for any point $p \in \mathcal{Z}$, it holds that

$$\|c_z - p\| \leq \sqrt{n}\,\mathrm{rad}(\mathcal{Z})$$

where $\mathrm{rad}(\mathcal{Z}) = \sup_{p \in \mathcal{Z}} \|c_z - p\|_\infty$.

3 Problem Definition

Consider a linear time-invariant (LTI) system

$$x(k+1) = Ax(k) + Bu(k) + w(k) \tag{3a}$$

$$y_i(k) = C_i x(k) + v_i(k) + a_i(k), \quad i \in \mathbb{Z}_{[1,p]} \tag{3b}$$

where $x(k) \in \mathbb{R}^{n_x}$ is the state, $u(k) \in \mathbb{R}^{n_u}$ is a known input, and $y_i(k) \in \mathbb{R}^{m_i}$ is the measured output of the i-th sensor with $i \in \mathbb{Z}_{[1,p]}$. The vector $w(k) \in \mathcal{W}$ represents the process noise, which is assumed to be contained in the zonotope $\mathcal{W} = \langle c_w, G_w \rangle$, and the vector $v_i(k) \in \mathcal{V}_i$ represents the measurement noise of the i-th sensor, which is assumed to be contained in the zonotope $\mathcal{V}_i = \langle c_{v_i}, G_{v_i} \rangle$, for every $i \in \mathbb{Z}_{[1,p]}$. Sensor i's measurement at time k may be corrupted by an arbitrary and unbounded attack signal $a_i(k) \in \mathbb{R}^{m_i}$, which is assumed to be designed by a malicious attacker.

Assumption 1. *We assume the following:*

(i) ***Maximum number of attacked sensors****: The attacker can attack up to $q \leq p-1$ number of sensors. Although the upper bound q is known, the exact number and the set of attacked sensors do not need to be known.*

(ii) ***Redundant observability****: There exists $c_{\mathsf{J}} \leq p - q$ such that, for every subset $\mathsf{J} \subset \mathbb{Z}_{[1,p]}$ of sensors with cardinality $|\mathsf{J}| = c_{\mathsf{J}}$, the pair (A, C_{J}) is observable, i.e., $\mathrm{rank}[\, C_{\mathsf{J}}^\top \; (C_{\mathsf{J}}A)^\top \; \ldots \; (C_{\mathsf{J}}A^{n_x-1})^\top \,]^\top = n_x$, where C_{J} is obtained by stacking all C_i, $i \in \mathsf{J}$, in row blocks.*

(iii) ***Knowledge of an initial set****: At time $k = 0$, we know a bounded set $\mathcal{X}_0 \subset \mathbb{R}^{n_x}$ containing the initial state $x(0) \in \mathcal{X}_0$.*

(iv) ***Bounded noise****: The process and measurement noise zonotopes $\mathcal{W}, \mathcal{V}_1, \ldots, \mathcal{V}_p$ are known.*

Assumption 1(i) restricts the maximum number of attacked sensors at each time instant so that at least one sensor is uncompromised. This assumption is fundamental in this paper, ensuring that, at every time $k \in \mathbb{Z}_{\geq 0}$, there exists a set of *uncompromised* (or *safe*) sensors

$$\mathsf{S}_k \subset \mathbb{Z}_{[1,p]} \text{ with } |\mathsf{S}_k| \geq p - q$$

such that $a_i(k) = 0_{m_i}$ for every $i \in \mathsf{S}_k$. This, along with the redundant observability (Assumption 1(ii)), allows us to ensure that the true state can be theoretically reconstructed from the set of uncompromised sensors in the absence of noise. In addition, the assumption entails that the attacker, even though omniscient about the system dynamics and noise bounds, has limited resources at hand and cannot attack all the sensors. We remark that

this assumption is not restrictive because it neither restricts the set of attacked sensors to be static with respect to time nor requires that q is less than half the number of sensors p, which is the fundamental assumption in the point-based secure state estimation literature. On the contrary, at any time instant, our problem setup allows the attacker to inject arbitrary signals to any subset of sensors with cardinality less than or equal to q, where q is only required to be strictly less than p.

Assumption 1(iii) is required to initialize the set-based state estimation algorithm with $\hat{\mathcal{X}}_0 = \mathcal{X}_0$ such that the inclusion of the state $x(k)$ in the estimated set $\hat{\mathcal{X}}_k \subset \mathbb{R}^{n_x}$ can be guaranteed. Note that Assumption (iii) does not require any bound on the radius of $\mathcal{X}_0$. Finally, Assumption 1(iv) is a standard assumption in robust estimation and control [44, 15].

Problem Statement

Given the uncertain system (3) subject to Assumption 1, we study the following problems:

i. Estimate a set $\hat{\mathcal{X}}_k$ guaranteeing the inclusion of the true state, i.e., $x(k) \in \hat{\mathcal{X}}_k$ for every $k \in \mathbb{Z}_{\geq 1}$.

ii. Provide conditions such that the attack signals $a_i(k)$ can be detected, identified, and/or filtered[1].

4 Secure Set-based State Estimation Algorithm

This section presents the secure set-based state estimation (S3E) algorithm, which is summarized below in Algorithm 1. In the following subsections, we describe all the steps of the algorithm in detail, provide a comparison with the point-based estimators, and present strategies to manage the complexity of the algorithm.

Algorithm 1 Secure set-based state estimation

Initialize: $\hat{\mathcal{X}}_0 = \mathcal{X}_0$
for $k = 1, 2, 3, \ldots$ **do**
  Time update (4): $\hat{\mathcal{X}}_{k|k-1}$
  Measurement update (6): $\hat{\mathcal{Z}}_k^i$ for every $i \in \mathbb{Z}_{[1,p]}$.
  Agreement protocol (7): $\mathcal{I}_k^h$ for $h \in \mathbb{Z}_{[1,n_J]}$.
  Estimate (8): $\hat{\mathcal{X}}_k$
end for

4.1 Time Update

The first step in the S3E algorithm is the time update

$$\hat{\mathcal{X}}_{k|k-1} = A\hat{\mathcal{X}}_{k-1} \oplus Bu(k-1) \oplus \mathcal{W} \tag{4}$$

where $\mathcal{W}$ is the zonotope that bounds the process noise $w(k)$, and $\hat{\mathcal{X}}_{k-1}$ is the estimated set computed in the previous time step. The computation of $\hat{\mathcal{X}}_k$ is described in Section 4.4 with the initialization $\hat{\mathcal{X}}_0 = \mathcal{X}_0$, where $\mathcal{X}_0$ is the known initial set from Assumption 1(iii).

Given $x(k-1) \in \hat{\mathcal{X}}_{k-1}$, the time update set $\hat{\mathcal{X}}_{k|k-1} \ni x(k)$ is the set of states to which the system can evolve at time $k \in \mathbb{Z}_{\geq 0}$ subject to the model (A, B), the input $u(k-1)$, and

[1] These terms are defined later in Section 5.

the noise zonotope $\mathcal{W}$. Although the attacker cannot directly influence the time update set as seen from (4), it can do so indirectly through the previous estimate $\hat{\mathcal{X}}_{k-1}$. Thus, the resilience against sensor attacks is achieved by carefully devising the estimated set $\hat{\mathcal{X}}_k$.

4.2 Measurement Update

The time update (4) uses model-based information to compute a set $\hat{\mathcal{X}}_{k|k-1}$ containing all points that could be reached by the system state in one time step if they are initialized at the previous estimated set $\hat{\mathcal{X}}_{k-1}$. The measurement update, on the other hand, corrects the conservative estimate of the model-based time update by restricting it according to the sensor measurements.

For every $i \in \mathbb{Z}_{[1,p]}$, let

$$\mathcal{Y}_k^i := y_i(k) - \mathcal{V}_i = \langle y_i(k) - c_{v_i}, G_{v_i} \rangle \tag{5}$$

be sensor i's output measurement set. Then, the measurement update of sensor i is given by the following generalized intersection of the time update set $\hat{\mathcal{X}}_{k|k-1}$ with the output measurement set $\mathcal{Y}_k^i$ of sensor i

$$\hat{\mathcal{Z}}_k^i := \hat{\mathcal{X}}_{k|k-1} \cap_{C_i} \mathcal{Y}_k^i = \{x \in \hat{\mathcal{X}}_{k|k-1} : C_i \hat{\mathcal{X}}_{k|k-1} \in \mathcal{Y}_k^i\} \tag{6}$$

which can be computed using (2).

The measurement update $\hat{\mathcal{Z}}_k^i$ is a subset of the time update set $\hat{\mathcal{X}}_{k|k-1}$ that is consistent with the measurement $y_i(k)$ of sensor i. That is, given the measurement noise zonotope $\mathcal{V}_i$, $\hat{\mathcal{Z}}_k^i$ is the set of states $x(k) \in \hat{\mathcal{X}}_{k|k-1}$ that could have produced the measurement $y_i(k)$. Notice that $\hat{\mathcal{Z}}_k^i$ can be an empty set if the sensor i's measurement is inconsistent with the time update.

4.3 Set-based Agreement Protocol

Let $n_{\mathsf{J}} := \binom{p}{c_{\mathsf{J}}}$, where $c_{\mathsf{J}} \leq p - q$ is the redundant observability parameter. Let $\mathsf{J}_1, \ldots, \mathsf{J}_{n_{\mathsf{J}}} \subset \mathbb{Z}_{[1,p]}$ be the disjoint subsets with $|\mathsf{J}_h| = c_{\mathsf{J}}$, for $h \in \mathbb{Z}_{[1,n_{\mathsf{J}}]}$. The agreement protocol between measurement updates in each J_h is given by

$$\mathcal{I}_k^h := \bigcap_{j \in \mathsf{J}_h} \hat{\mathcal{Z}}_k^j. \tag{7}$$

Theorem 1. *Let Assumption 1 hold and assume $x(k) \in \hat{\mathcal{X}}_{k|k-1}$. Then, at every time k, there exists $h \in \mathbb{Z}_{[1,n_{\mathsf{J}}]}$ such that the agreement set $\mathcal{I}_k^h$ is non-empty and contains the true state $x(k)$.*

Proof. We first show that the measurement update $\hat{\mathcal{Z}}_k^i$ contains the true state $x(k)$ if sensor i is in the uncompromised set of sensors S_k at time k. If i is uncompromised at time k, then $a_i(k) = 0$ and $y_i(k) = C_i x(k) + v_i(k)$. From (5), we have $\mathcal{Y}_k^i = \langle c_{y_i(k)}, G_{v_i} \rangle$ with $c_{y_i(k)} = y_i(k) - c_{v_i} = C_i x(k) + v_i(k) - c_{v_i}$. As $\mathcal{Y}_k^i$ and the noise zonotope $\mathcal{V}_i$ have the same generator matrix G_{v_i}, we have $\mathrm{rad}(\mathcal{Y}_k^i) = \mathrm{rad}(\mathcal{V}_i)$. Moreover, it holds that

$$\mathcal{Y}_k^i \supseteq \{c_{y_i(k)} - \xi : \|\xi\|_\infty \leq \mathrm{rad}(\mathcal{Y}_k^i)\}.$$

Thus, $c_{y_i(k)} - v_i(k) + c_{v_i} = C_i x(k) \in \mathcal{Y}_k^i$ because $\|v_i(k) - c_{v_i}\|_\infty \leq \mathrm{rad}(\mathcal{V}_i) = \mathrm{rad}(\mathcal{Y}_k^i)$. Since $x(k) \in \hat{\mathcal{X}}_{k|k-1}$ and $C_i x(k) \in \mathcal{Y}_k^i$, it implies that $\hat{\mathcal{Z}}_k^i$ in (6) is non-empty and contains $x(k)$.

We then show that, at every time k, there is a subset J_h of sensors with $|\mathsf{J}_h| = c_{\mathsf{J}}$ such that $\mathsf{J}_h \subseteq \mathsf{S}_k$. This is straightforward because q out of p sensors are compromised

and $c_{\mathsf{J}} \leq p - q$, so there must be at least one subset J_h with $|\mathsf{J}_h| = c_{\mathsf{J}}$ that contains the uncompromised sensors from the guaranteed $p - q$ uncompromised sensors. In other words, at every time $k \in \mathbb{Z}_{\geq 0}$, there exists $h \in \mathbb{Z}_{[1,n_{\mathsf{J}}]}$ such that $\mathsf{J}_h \subseteq \mathsf{S}_k$. Since $x(k) \in \hat{\mathcal{Z}}_k^i$ for every $i \in \mathsf{J}_h \subseteq \mathsf{S}_k$, the intersection in (7) is non-empty and $x(k) \in \mathcal{I}_k^h$. □

We have shown that the measurement update $\hat{\mathcal{Z}}_k^i$ of sensor i contains the true state $x(k)$ if it is uncompromised by the attacker at time k. Thus, the intersection between the measurement updates of any subset of uncompromised sensors must also contain $x(k)$. In the proof of Theorem 1, we show that by removing the number of attacked sensors q, we guarantee the existence of at least one subset J_h with cardinality $|\mathsf{J}_h| = c_{\mathsf{J}}$ which is attack-free. Therefore, the corresponding agreement set $\mathcal{I}_k^h$ must be non-empty and contain $x(k)$.

It is equivalently true to say that if the agreement set $\mathcal{I}_k^h$ is empty, then there must be at least one sensor in J_h that is compromised by the attacker. Notice that this is only a sufficient condition for detecting an attack because a stealthy attacker can design an attack that ensures that $\mathcal{I}_k^h$ is non-empty at any time k. However, this comes at a cost to the attacker, which is a discussion reserved for the next section in the paper.

4.4 Estimated Set

The estimated set is obtained by taking the union of all agreement sets:

$$\hat{\mathcal{X}}_k := \bigcup_{h \in \mathbb{Z}_{[1,n_{\mathsf{J}}]}} \mathcal{I}_k^h, \quad k \in \mathbb{Z}_{\geq 1} \tag{8}$$

where $\hat{\mathcal{X}}_0 = \mathcal{X}_0$.

Theorem 2. *Let Assumption 1 hold. Then, for every time $k \in \mathbb{Z}_{\geq 1}$, the inclusion $x(k) \in \hat{\mathcal{X}}_k$ is guaranteed, where the estimate $\hat{\mathcal{X}}_k$ is computed according to* (8).

Proof. We prove this result by induction. By Assumption 1(iii), we have $x(0) \in \mathcal{X}_0 = \hat{\mathcal{X}}_0$. Because $w(0) \in \mathcal{W}$, the inclusion holds for the time update at $k = 1$, $x(1) \in \hat{\mathcal{X}}_{1|0} = A\mathcal{X} \oplus Bu(0) \oplus \mathcal{W}$. Therefore, Theorem 1 can be applied, implying the existence of $h \in \mathbb{Z}_{[1,n_{\mathsf{J}}]}$ such that $x(1) \in \mathcal{I}_1^h$, which gives $x(1) \in \hat{\mathcal{X}}_1 = \bigcup_{h \in \mathbb{Z}_{[1,n_{\mathsf{J}}]}} \mathcal{I}_1^h$. This, in turn, implies that the time update at $k = 2$ contains the true state, $x(2) \in \hat{\mathcal{X}}_{2|1} = A\hat{\mathcal{X}}_1 \oplus Bu(1) \oplus \mathcal{W}$.

Now, let us assume $x(k') \in \hat{\mathcal{X}}_{k'|k'-1}$ for some $k' \in \mathbb{Z}_{\geq 2}$. Then, by Theorem 1, there exists $h \in \mathbb{Z}_{[1,n_c]}$ such that $x(k') \in \mathcal{I}_{k'}^h$, implying $x(k') \in \hat{\mathcal{X}}_{k'} = \bigcup_{h \in \mathbb{Z}_{[1,n_{\mathsf{J}}]}} \mathcal{I}_{k'}^h$. Therefore, we have the inclusion at the next time update, $x(k'+1) \in \hat{\mathcal{X}}_{k'+1|k'} = A\hat{\mathcal{X}}_{k'} \oplus Bu(k') \oplus \mathcal{W}$. Thus, the proof is completed because we showed that, for every $k \in \mathbb{Z}_{\geq 1}$, $x(k-1) \in \hat{\mathcal{X}}_{k-1|k-2}$ implies $x(k-1) \in \hat{\mathcal{X}}_{k-1}$, which, in turn, implies $x(k) \in \hat{\mathcal{X}}_{k|k-1}$. Hence, $x(k) \in \hat{\mathcal{X}}_k$. □

Although the above theorem guarantees the inclusion of the true state, it is important to remark that the number of (constrained) zonotopes in the estimated set $\hat{\mathcal{X}}_k$ may increase with respect to time if the attack is stealthy. We address this issue in Section 4.8 by proposing several techniques that manage the computational efficiency of the algorithm. However, it is worth mentioning that the proposed estimation algorithm is resilient because the attacker cannot deteriorate the estimation accuracy arbitrarily over time. If a sensor is injected with a large-valued attack signal, it will be automatically discarded by either yielding an empty measurement update or an empty agreement set (see Section 5).

4.5 Illustrative and numerical examples

We illustrate Algorithm 1 on a simple example of $p = 3$ sensors, where $q = 1$ sensor has been compromised. For the purposes of this illustration in Fig. 1, we chose all the sensors to have the same dimension $\mathbb{R}^{n_y}$ and the first sensor to be compromised. For each sensor i, we obtain its corresponding output measurement set $\mathcal{Y}_k^i$ in the output space as depicted in Fig. 1, which do not necessarily intersect as each sensor i can be measuring different components of the state $x(k) \in \mathbb{R}^{n_x}$.

Next, the time update set $\hat{\mathcal{X}}_{k|k-1}$ is constructed according to (4) and shown in the second box from the top in Fig. 1. Each output measurement set $\mathcal{Y}_k^i$ is mapped back to the state space $\mathbb{R}^{n_x}$ and constrained to be in the time update set $\hat{\mathcal{X}}_{k|k-1}$ as shown in the third box from the top in Fig. 1. The intersection of the measurement update sets $\hat{\mathcal{Z}}_k^i$ is informative as the intersections of the attack-free $\hat{\mathcal{Z}}_k^i$ is non-empty. However, the measurement update set $\hat{\mathcal{Z}}_k^i$ corresponding to a sensor i that is corrupted could also intersect with the other attack-free measurement update sets. This scenario is depicted in Fig. 1 where the measurement update set $\hat{\mathcal{Z}}_k^1$ of the compromised sensor 1 intersects with the attack-free measurement update set $\hat{\mathcal{Z}}_k^2$ due to an intelligently designed attack signal a_1. We discuss classes of attack signals that can be detected by our algorithm in Section 5.

Since the identity of the compromised sensor(s) is unknown at any given time step k, but the number of compromised sensors q is known, we have to check for non-empty intersections of all $p - q$ combinations of the measurement update sets $\hat{\mathcal{Z}}_k^i$, which totals to $\binom{p}{p-q}$ checks. Its union then forms our estimated set $\hat{\mathcal{X}}_k$ as depicted in the bottom box of Fig. 1.

To further elucidate Algorithm 1, we consider a simple numerical example. Simulations for more realistic systems are provided later in Section 6.

Example 1. *Consider the following system*

$$\begin{aligned} x(k+1) &= 0.1x(k) + w(k), \\ y_i(k) &= x(k) + v_i(k) + a_i(k), \qquad i \in \{1, 2, 3\} \end{aligned}$$

where for all $k \in \mathbb{Z}_{\geq 0}$, the disturbance and noise are $w(k) \in [-10, 10]$ and $v_i(k) \in [-0.1, 0.1]$, respectively. Notice that this system is observable via each sensor i. For the purposes of this illustration, suppose $x(0) = 0$, $\hat{\mathcal{X}}_0 = [-100, 100]$ and the disturbance and noise at time step $k = 1$ are $w(1) = 4$, $v_1(1) = 0.03$, $v_2(1) = -0.02$, $v_3(1) = -0.07$, respectively. Notice also that $x(1) = 4$.

Suppose sensor 1 and 2 are under attack and sensor 3 is uncompromised, i.e., $p = 3$, $q = 2$, $c_{\mathsf{J}} = 1$, the parameter of Algorithm 1 is $n_{\mathsf{J}} = 3$, and $a_3(k) = 0$ for all $k \in \mathbb{Z}_{\geq 0}$. Further, suppose the attacks on sensors 1 and 2 at time step $k = 1$ are $a_1(1) = -10.03$ and $a_2(1) = -10$. We construct the following sets according to Algorithm 1 for time step $k = 1$:

- *The measurement output set for each sensor according to (5) is $\mathcal{Y}_k^1 = [-6.1, -5.9]$, $\mathcal{Y}_k^2 = [-6.12, -5.92]$ and $\mathcal{Y}_k^3 = [3.83, 4.03]$, respectively.*
- *The time update set according to (4): $\hat{\mathcal{X}}_{1|0} = 0.1\hat{\mathcal{X}}_0 \oplus \mathcal{W} = [-10, 10] \oplus [-10, 10] = [-20, 20]$.*
- *Measurement update sets according to (6): $\hat{\mathcal{Z}}_1^1 = [-6.1, -5.9]$, $\hat{\mathcal{Z}}_1^2 = [-6.12, -5.92]$, $\hat{\mathcal{Z}}_1^3 = [3.83, 4.03]$.*
- *Agreement sets using (7): for $\mathsf{J}_1 = \{1\}$, $\mathsf{J}_2 = \{2\}$ and $\mathsf{J}_3 = \{3\}$, we get $\mathcal{I}_1^1 = \hat{\mathcal{Z}}_1 = [-6.1, -5.9]$, $\mathcal{I}_1^2 = \hat{\mathcal{Z}}_1^2 = [-6.12, -5.92]$, and $\mathcal{I}_1^3 = \hat{\mathcal{Z}}_1^3 = [3.83, 4.03]$.*
- *Finally, the estimated set using (8) is $\hat{\mathcal{X}}_1 = \mathcal{I}_1^1 \cup \mathcal{I}_1^2 \cup \mathcal{I}_1^3 = [-6.12, -5.9] \cup [3.83, 4.03]$.*

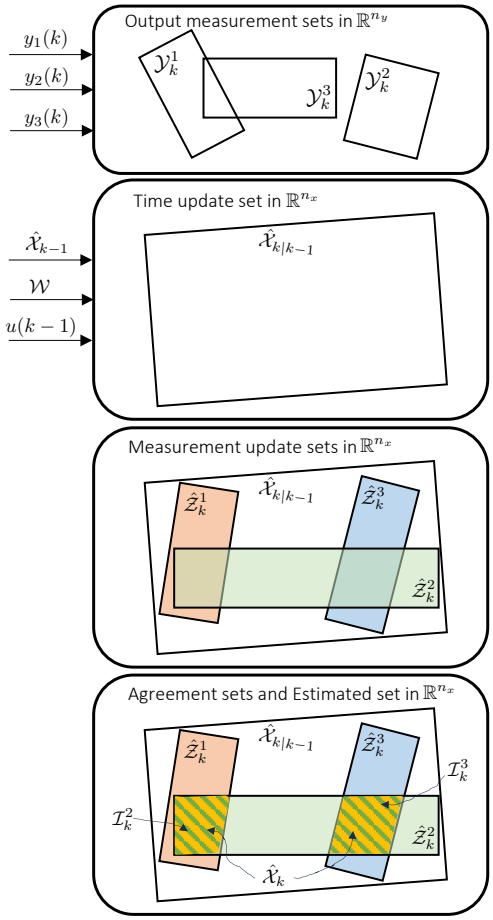

Figure 1: Illustration of Algorithm 1 with $p = 3$ and $q = 1$. The algorithm deduces that either sensor 1 or 3 has been compromised since their corresponding measurement update sets $\hat{\mathcal{Z}}_k^1$ and $\hat{\mathcal{Z}}_k^3$ do not intersect.

By Theorem 2, the true state $x(1)$ at time step $k = 1$ can lie either in the interval $[-6.12, -5.9]$ or $[3.83, 4.03]$. The disparity between the two estimated intervals highlights the limitation of the S3E algorithm: is the true state contained in one interval or the other? Nonetheless, having multiple sets of reasonable sizes that can potentially contain the true state provide utility in safety-critical applications, where the estimated set serves to verify whether the system's trajectory could have entered regions of the state space that are designated as unsafe.

4.6 S3E Algorithm vs. Point-based Methods

Point-based secure state estimators provide a point-based estimate $\hat{x}(k) \in \mathbb{R}^{n_x}$. The security guarantee for these estimators against sensor attacks is an estimation error bound that only depends on the noise and is independent of the attack signal; see [12, 31, 38, 20, 29, 11, 22, 16]. However, for providing such a security guarantee, the fundamental assumption is that the number of attacked sensors is strictly less than half of the total number of sensors. The

estimation error bound for point-based estimators has the following form

$$\|\hat{x}(k) - x(k)\| \leq \beta(\|\hat{x}(0) - x(0)\|, k) + \gamma_1(\|w\|_{\ell^\infty}) + \gamma_2(\|v\|_{\ell^\infty}) + \gamma_3(\|a\|_{\ell^\infty})\mathbb{1}_{q \geq \frac{p}{2}} \tag{9}$$

where β is a $\mathcal{KL}$ function, $\gamma_1, \gamma_2, \gamma_3$ are $\mathcal{K}_\infty$ functions[2], and $\mathbb{1}_{q \geq \frac{p}{2}}$ is the indicator function, which is 1 when $q \geq \frac{p}{2}$ and 0 otherwise. Here, $v = [\, v_1^\top \; \dots \; v_p^\top \,]^\top$ and $a = [\, a_1^\top \; \dots \; a_p^\top \,]^\top$ concatenate the measurement noise and attack signals in a column vector, respectively.

It is easy to see that if the number of sensors p is not strictly twice the number of compromised sensors q (i.e., the assumption of $q < p/2$ is violated), the estimation error bound depends on the magnitude of the attack signal. In this case, the attacker can arbitrarily degrade the accuracy of point-based estimators by injecting attack signals of arbitrarily large magnitude. On the other hand, our set-based approach avoids this situation because it restricts the attack space by allowing only the signals of small magnitude to go undetected. Large attack signals are automatically detected and discarded by either the measurement update (6) or the agreement protocol (7).

Another important advantage of our set-based approach is that the estimated set $\hat{\mathcal{X}}_k$ contains a collection of agreement sets, which serve as multiple hypotheses for the true state $x(k)$ in the state space[3]. Point-based estimators, on the other hand, have an estimation error bound (9) that indicates a single ball inside which the true state might lie. In the presence of attacks on more than half the number of sensors, the size of this ball depends on the magnitude of the attacks. Thus, as described in Example 1, our set-based method can at least identify plausible regions in the state space containing the true state. Moreover, even when less than half the number of sensors are compromised, bounds like (9) are very conservative in practice because the comparison functions $\beta \in \mathcal{KL}$ and $\gamma \in \mathcal{K}_\infty$ are hard to estimate.

To illustrate this point, we revisit Example 1 for the case where fewer than half the number of sensors are attacked.

Example 2. *Reconsider Example 1, but suppose sensor* 1 *is under attack while sensors* 2 *and* 3 *are uncompromised. We have the number of sensors* $p = 3$ *and the number of attacked sensors* $q = 1$*, and, for the sake of this illustration, assume the attack signals* $a_2(k) = a_3(k) = 0$*, for all* $k \in \mathbb{Z}_{\geq 0}$*. Furthermore, suppose the attack on sensor* 1 *is* $a_1(1) = -10.03$*. We obtain the estimated set according to Algorithm 1 at time step* $k = 1$ *as follows:*

- *The measurement output set for each sensor is* $\mathcal{Y}_k^1 = [-6.1, -5.9]$ *and* $\mathcal{Y}_k^3 = [3.83, 4.03]$ *as calculated before, and* $\mathcal{Y}_k^2 = [3.88, 4.08]$ *according to* (5)*.*
- *The time update set* $\hat{\mathcal{X}}_{1|0} = [-20, 20]$*.*
- *Measurement update sets according to* (6)*:* $\hat{\mathcal{Z}}_1^1 = [-6.1, -5.9]$*,* $\hat{\mathcal{Z}}_1^2 = [3.88, 4.08]$*, and* $\hat{\mathcal{Z}}_1^3 = [3.83, 4.03]$*.*
- *Agreement sets using* (7)*: for* $\mathsf{J}_1 = \{1, 2\}$*,* $\mathsf{J}_2 = \{1, 3\}$ *and* $\mathsf{J}_3 = \{2, 3\}$*, we get* $\mathcal{I}_1^1 = \hat{\mathcal{Z}}_1^1 \cap \hat{\mathcal{Z}}_1^2 = \emptyset$*,* $\mathcal{I}_1^2 = \hat{\mathcal{Z}}_1^1 \cap \hat{\mathcal{Z}}_1^3 = \emptyset$*, and* $\mathcal{I}_1^3 = \hat{\mathcal{Z}}_1^2 \cap \hat{\mathcal{Z}}_1^3 = [3.88, 4.03]$*.*
- *Finally, the estimated set using* (8) *is* $\hat{\mathcal{X}}_1 = \mathcal{I}_1^1 \cup \mathcal{I}_1^2 \cup \mathcal{I}_1^3 = [3.88, 4.03]$*.*

[2] A continuous function $\gamma : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ is a class $\mathcal{K}$ function, if it is strictly increasing and $\gamma(0) = 0$. It is a class $\mathcal{K}_\infty$ function if it is also unbounded. A continuous function $\beta : \mathbb{R}_{\geq 0} \times \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$ is a class $\mathcal{KL}$ function, if: (i) $\beta(., s)$ is a class $\mathcal{K}$ function for any $s \geq 0$; (ii) $\beta(r, .)$ is non-increasing and (iii) $\beta(r, s) \to 0$ as $s \to \infty$ for any $r \geq 0$.

[3] Although beyond the scope of this paper, an interesting research direction is developing a hypothesis testing technique to figure out the true set from the sets influenced by the attack signals.

We consider the center 3.955 *of* $\hat{\mathcal{X}}_1$ *as an estimate of* $x(1) = 4$*, which gives the estimation error* 0.045*. On the other hand, using the secure state estimation algorithm given in [12], the point-based estimate is obtained as* $\hat{x}(1) = 3.164$*, which gives a much higher estimation error equal to* 0.836*.*

4.7 Bound on the Set-based Estimation Error

Since Theorem 2 guarantees that the true state $x(k)$ of system (3) lies in at least one of the agreement sets in the estimated collection of sets $\hat{\mathcal{X}}_k$ at each $k \in \mathbb{Z}_{\geq 0}$, it must also lie in a zonotope that overbounds $\hat{\mathcal{X}}_k$. That is, let $\hat{\mathcal{X}}_k$ be overbounded by a zonotope

$$\bar{\mathcal{X}}_k = \langle \hat{c}_x(k), \hat{G}_x(k) \rangle \tag{10}$$

which is obtained by solving

$$\min \operatorname{rad}(\bar{\mathcal{X}}_k) \text{ subject to } \hat{\mathcal{X}}_k \subseteq \bar{\mathcal{X}}_k. \tag{11}$$

Then, considering $\hat{c}_x(k)$ to be an estimate of $x(k)$, the estimation error can be bounded by

$$\|\hat{c}_x(k) - x(k)\| \leq \operatorname{rad}(\bar{\mathcal{X}}_k).$$

We remark that this error bound is significantly smaller in practice than the error bounds obtained by point-based secure estimators.

4.8 Strategies to Handle Computational Complexity

A major drawback of Algorithm 1 is that the attacker can design intelligent and stealthy attacks such that the number of agreement sets in the estimate $\hat{\mathcal{X}}_k$ keep growing with time. Although generating such attacks is a difficult problem as we argue in Section 5, the computational complexity of the algorithm may increase exponentially with time k in the worst case scenario. Therefore, it is important to reduce the complexity by adopting several pruning methods that are described below.

An obvious step is to remove any empty or subsets of other sets in the estimated set $\hat{\mathcal{X}}_k$. It is also possible to obtain a single overbounding zonotope of $\hat{\mathcal{X}}_k$ as in (11), and use it instead of $\hat{\mathcal{X}}_k$ in the next time update. However, a better trade-off between accuracy and complexity is to not overbound the whole collection but only the intersecting zonotopes in the collection $\hat{\mathcal{X}}_k$. This may not make the cardinality of $\hat{\mathcal{X}}_k$ equal to one, but it reduces it significantly by allowing minimal loss of accuracy. Another method is to employ zonotope reduction methods [41] to reduce the number of generators in the zonotopes, which are often increased when performing the Minkowski sum operations.

5 Attack Detection, Identification, and Filtering

In this section, we provide sufficient conditions that guarantee the security of the S3E algorithm.

5.1 Attack Detection

We say that the attack is *detected* if the system operator (defender) knows for certain that at least one sensor is compromised at time k, i.e., $|\mathsf{S}_k| < p$. While attack detection does not necessarily reveal the specific compromised sensor(s), it allows the system operator to raise the alarm and take necessary steps to secure the system. We provide a sufficient condition for attack detection below.

Proposition 3. *Consider the agreement sets $\mathcal{I}_k^h$, for $h \in \mathbb{Z}_{[1,n_\mathsf{J}]}$, given by* (7)*. Then, if*

$$\bigcap_{h \in \mathbb{Z}_{[1,n_\mathsf{J}]}} \mathcal{I}_k^h = \emptyset \tag{12}$$

there is at least one sensor that is compromised at time k.

Proof. We prove the proposition by showing its logical equivalence: 'no sensor is compromised' implies '(12) does not hold.' Since for every $h \in \mathbb{Z}_{[1,n_\mathsf{J}]}$, the system is redundantly observable through every subset J_h of sensors with $|\mathsf{J}| = c_\mathsf{J}$. If no sensor is compromised, then, by Theorem 1, we have that, for every $h \in \mathbb{Z}_{[1,n_\mathsf{J}]}$, the agreement set $\mathcal{I}_k^h$ contains the true state $x(k)$. Therefore, $\mathcal{I}_k^1 \cap \cdots \cap \mathcal{I}_k^{n_\mathsf{J}} \neq \emptyset$. □

The above proposition shows that whenever the estimated set contains more than one agreement sets that are disjoint, there has to be at least one sensor that is compromised. Once the attack is detected, the system operator can take necessary measures to secure the system and bring it to safety. If the attacker wants to stay stealthy, they must design attack signals that ensure that all the agreement sets yield a non-empty intersection. In the next subsections, we show that the magnitude of such a stealthy attack has to be sufficiently small.

5.2 Attack Identification

We say that the attack $a_i(k)$ is *identified* if the system operator knows for certain that sensor i is compromised at time k, i.e., $i \notin \mathsf{S}_k$. The attack identification allows us to automatically discard the measurements coming from the compromised sensors in the S3E algorithm. We achieve attack identification through a measurement update (6), which involves a generalized intersection between the time update $\hat{\mathcal{X}}_{k|k-1}$ and the output measurement set $\mathcal{Y}_k^i$.

If sensor i is attacked at time k, we derive a sufficient condition on the magnitude of the attack signal $a_i(k)$ that would yield an empty measurement update set $\hat{\mathcal{Z}}_k^i$. This will automatically identify the attack on sensor i and discard its effect from the estimation algorithm. In the next theorem, we use the notation $\bar{\mathcal{X}}_{k|k-1}$ for representing the overbounding zonotope of the time update $\hat{\mathcal{X}}_{k|k-1}$, i.e., $\hat{\mathcal{X}}_{k|k-1} \subseteq \bar{\mathcal{X}}_{k|k-1}$, where

$$\bar{\mathcal{X}}_{k|k-1} \coloneqq \langle c_x(k|k-1), G_x(k|k-1) \rangle.$$

Theorem 4. *Suppose sensor i is attacked at time k. Then, if*

$$\|a_i(k)\| > \sqrt{m_i}[\mathrm{rad}(C_i \bar{\mathcal{X}}_{k|k-1}) + \mathrm{rad}(\mathcal{V}_i)] - \|C_i[x(k) - c_x(k|k-1)] + v_i(k) - c_{v_i}\| \tag{13}$$

then the measurement update $\hat{\mathcal{Z}}_k^i = \emptyset$. Recall that m_i is the dimension of the output y_i of system (3)*.*

Proof. Given a zonotope $C_i \bar{\mathcal{X}}_{k|k-1}$ with center $C_i c_x(k|k-1)$ and radius $\mathrm{rad}(C_i \bar{\mathcal{X}}_{k|k-1})$, we have that any point $p \in \mathbb{R}^{m_i}$ is outside the $C_i \bar{\mathcal{X}}_{k|k-1}$ if $\|C_i c_x(k|k-1) - p\| > \sqrt{m_i}\mathrm{rad}(C_i \bar{\mathcal{X}}_{k|k-1})$. Similarly, for $\mathcal{Y}_k^i$ whose center is $C_i x(k) + v_i(k) - c_{v_i} + a_i(k)$ and radius is $\mathrm{rad}(\mathcal{Y}_k^i)$, any $p \in \mathbb{R}^{m_i}$ is outside $\mathcal{Y}_k^i$ if $\|C_i x(k) + v_i(k) - c_{v_i} + a_i(k) - p\| > \sqrt{m_i}\mathrm{rad}(\mathcal{Y}_k^i)$. Here, the square root term $\sqrt{m_i}$ appears because the radius of a zonotope is a maximum norm $\|\cdot\|_\infty$, which relates to the Euclidean norm $\|\cdot\|$ as follows: for any vector $p \in \mathbb{R}^{m_i}$, $\|p\| \leq \sqrt{m_i}\|p\|_\infty$.

If the distance between the centers of $C_i \bar{\mathcal{X}}_{k|k-1}$ and $\mathcal{Y}_k^i$ is more than $\sqrt{m_i}[\mathrm{rad}(\bar{\mathcal{X}}_{k|k-1}) + \mathrm{rad}(\mathcal{Y}_k^i)]$, then any point $x \in \bar{\mathcal{X}}_{k|k-1}$ will be such that $C_i x \notin \mathcal{Y}_k^i$. Therefore, in this case, the measurement update (6) is empty. In other words, we have $\hat{\mathcal{Z}}_k^i = \emptyset$ if

$$\|C_i c_x(k|k-1) - [C_i x(k) + v_i(k) - c_{v_i} + a_i(k)]\| > \sqrt{m_i}[\mathrm{rad}(\bar{\mathcal{X}}_{k|k-1}) + \mathrm{rad}(\mathcal{Y}_k^i)].$$

Moreover, we obtain the following due to the triangle inequality

$$\|C_i[x(k)-c_x(k|k-1)]+v_i(k)-c_{v_i}\|+\|a_i(k)\| \geq \|C_i c_x(k|k-1)-[C_i x(k)+v_i(k)-c_{v_i}+a_i(k)]\|.$$

Thus, (13) is established by combining the above two inequalities and, from (5), observing that $\mathrm{rad}(\mathcal{Y}_k^i) = \mathrm{rad}(\mathcal{V}_i)$. □

The sufficient condition in Theorem 4 provides a class of attack signals that can be identified at the measurement update step of the S3E algorithm. If the measurement update set $\hat{\mathcal{Z}}_k^i$ is empty, then it reveals that sensor i is compromised at time k. However, the implication does not hold in the other direction. Nonetheless, while (13) establishes a threshold above which an attack will lead to an empty generalized intersection, it does not guarantee that smaller attacks will be undetected. The overapproximations used in the proof of Theorem 4 can introduce conservatism, meaning that attacks slightly below the threshold might still be detectable. In other words, the measurement update $\hat{\mathcal{Z}}_k^i$ could still be empty even if the attacker violates (13). This can happen if the attack signal is such that $C_i\hat{\mathcal{X}}_{k|k-1}$ and $\mathcal{Y}_k^i$ are inconsistent with each other at time k, resulting in an empty generalized intersection.

Notice that the bound (13) depends on the specific realization of measurement noise $v_i(k)$ at time k, which is unknown to both the attacker and the defender. One can further lower bound (13) and obtain a deterministic bound as follows:

$$\begin{aligned}
\|a_i(k)\| &> \sqrt{m_i}[\mathrm{rad}(C_i\bar{\mathcal{X}}_{k|k-1}) + \mathrm{rad}(\mathcal{V}_i)] - \|C_i[x(k) - c_x(k|k-1)] + v_i(k) - c_{v_i}\| \\
&\geq \sqrt{m_i}[\mathrm{rad}(C_i\bar{\mathcal{X}}_{k|k-1}) + \mathrm{rad}(\mathcal{V}_i)] - \|C_i[x(k) - c_x(k|k-1)]\| - \|v_i(k) - c_{v_i}\| \\
&\geq \sqrt{m_i}[\mathrm{rad}(C_i\bar{\mathcal{X}}_{k|k-1}) + \mathrm{rad}(\mathcal{V}_i)] - \|C_i[x(k) - c_x(k|k-1)]\| - \sqrt{m_i}\mathrm{rad}(\mathcal{V}_i) \\
&= \sqrt{m_i}\mathrm{rad}(C_i\bar{\mathcal{X}}_{k|k-1}) - \|C_i[x(k) - c_x(k|k-1)]\|.
\end{aligned}$$

If the attacker wants to stay stealthy and effective, it can choose the attack signal $a_i(k)$, for every $i \in \mathbb{Z}_{[1,p]}$, satisfying

$$\|a_i(k)\| \leq \sqrt{m_i}\mathrm{rad}(C_i\bar{\mathcal{X}}_{k|k-1}) - \|C_i[x(k) - c_x(k|k-1)]\|.$$

Although this still does not ensure a non-empty $\hat{\mathcal{Z}}_k^i$, it may serve as a rough guideline for the attacker to generate a stealthy attack signal.

5.3 Attack Filtering

We say that the attack $a_i(k)$ is *filtered out* if it is detected and discarded from the agreement protocol. To elucidate, let $i \in \mathsf{J}_h$, for some $h \in \mathbb{Z}_{[1,n_\mathsf{J}]}$. Then, if the attack signal $a_i(k)$ is filtered out, then the corresponding agreement set $\mathcal{I}_k^h$ is empty and does not affect the estimated set $\hat{\mathcal{X}}_k$. Although the attack is filtered out whenever it is identified, filtering does not imply identification in general.

In addition to the measurement update, the attacks could also be filtered out via the intersections in the agreement protocol (7), which could yield an empty agreement set $\mathcal{I}_k^h$ in case of inconsistencies between the measurement updates of the sensors in J_h. However, unlike the measurement update, the attacks may not always be identified whenever some of the agreement sets are empty. This is because a compromised sensor i could be a part of multiple agreement sets $\mathcal{I}_k^h$, for $h \in \mathbb{Z}_{[1,n_\mathsf{J}]}$, some of which may turn out to be non-empty. One can identify the attack $a_i(k)$ if, for every $h \in \mathbb{Z}_{[1,n_\mathsf{J}]}$ such that $i \in \mathsf{J}_h$, we have $\mathcal{I}_k^h = \emptyset$.

To present the next proposition, we introduce some notation. Recall that $\mathsf{S}_k \subseteq \mathbb{Z}_{[1,p]}$ is the subset of uncompromised sensors at time k. Similarly, let $\mathsf{A}_k = \mathbb{Z}_{[1,p]} \setminus \mathsf{S}_k$ be the subset

of compromised sensors at time k. Define

$$\mathcal{I}_k^{h,\mathsf{A}_k} = \bigcap_{j \in \mathsf{J}_h \setminus \mathsf{S}_k} \hat{\mathcal{Z}}_k^j$$

to be the agreement set of attacked sensors in J_h and let its overbounding zonotope be

$$\bar{\mathcal{I}}_k^{h,\mathsf{A}_k} = \langle c_{\mathcal{I}}^{h,\mathsf{A}_k}(k), G_{\mathcal{I}}^{h,\mathsf{A}_k}(k) \rangle.$$

Similarly, define

$$\mathcal{I}_k^{h,\mathsf{S}_k} = \bigcap_{j \in \mathsf{J}_h \setminus \mathsf{A}_k} \hat{\mathcal{Z}}_k^j$$

be the agreement set of uncompromised sensors in J_h and let its overbounding zonotope be

$$\bar{\mathcal{I}}_k^{h,\mathsf{S}_k} = \langle c_{\mathcal{I}}^{h,\mathsf{S}_k}(k), G_{\mathcal{I}}^{h,\mathsf{S}_k}(k) \rangle.$$

Proposition 5. *Let $\mathsf{J}_h \subset \mathbb{Z}_{[1,p]}$ be a subset of sensors with $|\mathsf{J}_h| = c_{\mathsf{J}}$. Then, if*

$$\|c_{\mathcal{I}}^{h,\mathsf{A}_k}(k)\| > \sqrt{n_x}[\mathrm{rad}(\bar{\mathcal{I}}_k^{h,\mathsf{S}_k}) + \mathrm{rad}(\bar{\mathcal{I}}_k^{h,\mathsf{A}_k})] - \|c_{\mathcal{I}}^{h,\mathsf{S}_k}(k)\| \tag{14}$$

then the agreement set $\mathcal{I}_k^h = \mathcal{I}_k^{h,\mathsf{S}_k} \cap \mathcal{I}_k^{h,\mathsf{A}_k} = \emptyset$. Recall that n_x is the dimension of the state $x(k)$ of system (3).

Proof. The proof follows from a similar line of argument as the proof of Theorem 4. □

The interpretation of the inequality in Proposition 5 is not as straightforward as the one in Theorem 4. If the attacker wants to be effective and remain unfiltered, it is necessary for them to devise an attack strategy that violates (14) at every time instant. Otherwise, the attack is discarded due to an empty agreement set. To violate (14), the attacker must take into account the distance between the intersections of the measurement updates of the compromised and uncompromised sensors.

Notice that $c_{\mathcal{I}}^{h,\mathsf{A}_k}(k)$, which is the center of the overbounding zonotope $\bar{\mathcal{I}}_k^{h,\mathsf{A}_k}$, is actually the Chebyshev center of the agreement set $\mathcal{I}_k^{h,\mathsf{A}_k}$. This agreement set is formed by the intersections of the measurement update sets of the attacked sensors A_k. The centers of these measurement update sets are influenced by the corresponding attack signals $a_j(k)$. Therefore, to achieve $\mathcal{I}_k^h \neq \emptyset$, it is necessary that the attacker coordinates the attacks on sensors in such a way that the condition of Proposition 5 is violated.

From the attacker's perspective, the condition of Proposition 5 is necessary for $\mathcal{I}_k^h \neq \emptyset$, and is not sufficient. This points to several interesting open problems for designing effective attacks. Beyond Theorem 4 and Propositions 3 and 5, our future work will investigate whether a sufficient condition could be derived for the attacker to ensure that both the measurement updates and the agreement sets are non-empty. Moreover, the complexity of generating such attacks is also one of the important questions that will be deferred for our future work.

5.4 Algorithm for Attack Filtering and Identification

A notable contribution of the set-based state estimation scheme in this paper over other secure schemes is Assumption 1(i), which allows the attacker to compromise not only up to $p-1$ sensors at each time instant but also different subsets of sensors at different times. To this end, we remark that we can detect only those compromised sensors that are injected with non-stealthy attack signals, whereby non-stealthiness we mean the attack signals that

Algorithm 2 Attack filtering and identification

Require: Measurement updates $\hat{\mathcal{Z}}_k^i$ for $i \in \mathbb{Z}_{[1,p]}$, sensor subsets J_h, for $h \in \mathbb{Z}_{[1,n_\jmath]}$, and their corresponding agreement sets $\mathcal{I}_k^h$.

Initialize $\hat{\mathsf{S}}_k = \emptyset$
$\hat{\mathsf{A}}_k \leftarrow \{i \in \mathbb{Z}_{[1,p]} : \hat{\mathcal{Z}}_k^i = \emptyset\}$
for $h = 1, \ldots, n_\jmath$ **do**
    if $\mathcal{I}_k^h \neq \emptyset$ **then**
        Estimated safe subset $\hat{\mathsf{S}}_k \leftarrow \hat{\mathsf{S}}_k \cup \mathsf{J}_h$
    end if
end for
Identified attacked sensors $\hat{\mathsf{A}}_k \leftarrow \hat{\mathsf{A}}_k \cup \{\mathbb{Z}_{[1,p]} \setminus \hat{\mathsf{S}}_k\}$.

satisfy the conditions of Theorem 4 and Propositions 3 and 5. The attack identification algorithm is fairly simple and can be summarized as in Algorithm 2.

At time k, Algorithm 2 first identifies some of the attacked sensors by checking whether the corresponding measurement update is empty. Then, for each $h \in \mathbb{Z}_{[1,n_\jmath]}$, if the agreement set $\mathcal{I}_k^h$ is non-empty, then the sensors with indices in J_h are either safe or compromised with a stealthy attack signal. Otherwise, the set J_h contains at least one attacked sensor. By checking all the combinations J_h and storing a 'potentially' safe subset of sensors in $\hat{\mathsf{S}}_k$ at every iteration, a subset of attacked sensors $\hat{\mathsf{A}}_k$ are updated by adding those sensors that are not in $\hat{\mathsf{S}}_k$.

Remark 1. *The estimated safe subset $\hat{\mathsf{S}}_k$ contains the true safe subset S_k at time k, i.e., $\mathsf{S}_k \subseteq \hat{\mathsf{S}}_k$. If there are sensors that are injected with small-valued stealthy attack signals, Algorithm 2 may not identify those attacks and considers those sensors to be uncompromised. Therefore, the estimated set of identified attacked sensors $\hat{\mathsf{A}}_k$ is only a subset of the true set of attacked sensors $\mathsf{A}_k = \mathbb{Z}_{[1,p]} \setminus \mathsf{S}_k$.*

Remark 2. *In Algorithm 2, we identify a subset of attacked sensors at every time k. However, in the time-invariant attack setting where the attacker does not change the subset of sensors to compromise at every time k, Algorithm 2 can be adapted to cumulatively detect and remove the attacked sensors over time. This has potential applications in sensor fault detection and isolation, as faults can be considered as naive, time-invariant attacks.*

6 Numerical Simulation

We evaluated our proposed algorithms using two examples. In the first illustrative example, we assume that half of the total number of sensors can be targeted by an attacker at every time instant. We illustrate the estimation algorithm and use a strategy to reduce the complexity with a minimal loss of estimation accuracy. In the second example of a three-story building structure, we assume redundant observability from every subset of two sensors. We use the CORA toolbox [4] to generate our simulations.

6.1 Two Dimensional Linear System

In this example, we consider a simple two-dimensional linear system for illustrative purposes. We consider a zero control input $u(k) \equiv 0$, and the system matrices are given by

$$A = \begin{bmatrix} 1 & 0 \\ 1 & 1 \end{bmatrix}, \quad C_1 = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}, \quad C_2 = \begin{bmatrix} 1 & 1 \\ 1 & 0 \end{bmatrix},$$
$$C_3 = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}, \quad C_4 = \begin{bmatrix} 1 & 2 \\ 2 & 1 \end{bmatrix}.$$

The number of sensors is $p = 4$, and we suppose that an attacker can target any combination of $q = 2$ sensors at different time instants. In this case, since $q = p/2$, the point-based state estimators requiring $q < p/2$ cannot be employed.

Let the process noise bound be $\mathcal{W} = \langle 0, \sigma_{\mathcal{W}} I_2 \rangle$ with $\sigma_{\mathcal{W}} = 0.02$ and the measurement noise bounds of the four sensors be $\mathcal{V}_1 = \mathcal{V}_2 = \mathcal{V}_3 = \mathcal{V}_4 = \langle 0, \sigma_{\mathcal{V}} I_2 \rangle$ with $\sigma_{\mathcal{V}} = 1$, where I_2 is the identity matrix of dimension 2×2. The attack is generated according to the following

$$a_i(k) \sim U(-\sigma_{\mathcal{V}}\phi(k), \sigma_{\mathcal{V}}\phi(k)) \tag{15}$$

where U is a uniform distribution over an interval $(-\sigma_{\mathcal{V}}\phi(k), \sigma_{\mathcal{V}}\phi(k))$ and $\phi(k)$ is strictly increasing sequence with $\phi(1) = 1$. The index of the two attacked sensors rotates among the available sensor indices.

The time update, measurement update, agreement, and estimated sets can be computed using Algorithm 1. Fig. 2 illustrates the time update set $\hat{\mathcal{X}}_{k|k-1}$ (green), the measurement sets $\hat{\mathcal{Z}}_k^i$ (pink (unattacked) and red (attacked)), the estimated set $\hat{\mathcal{X}}_k$ (black), and the true state $x(k)$ (blue) for different sets of sensors attacked at different times. The attacker chooses a random combination of two sensors to compromise at each time instant. At time $k = 1$, we have in Fig. 2a sensors 2 and 3 under attack. Then, sensors 3 and 4 are attacked in Fig. 2b at time step $k = 2$. Notice that the true state stays included in the estimated set $\hat{\mathcal{X}}_k$ for every k.

The generated attack values lead to an increase in the number of generated sets, as depicted in Fig. 2c. To manage this complexity, we employ a reduction technique that involves taking the union of intersecting estimated sets. The impact of this technique is illustrated in Fig. 3, where we use the same seed for generating random attacks from (15). This approach can be seen by comparing Fig. 2a with Fig. 3a. The reduced estimated sets are then carried forward to subsequent steps, as shown in Fig. 3b and Fig. 3c. This is why the sets in these figures show slight differences when compared to those in Fig. 2b and Fig. 2c.

6.2 Three-story Building Structure

We now consider a three-story building structure of [40] described by a mechanical system

$$M\ddot{q}(t) + D\dot{q}(t) + Sq(t) = Gu(t), \tag{16}$$

where $q(t) \in \mathbb{R}^3$ is the vector of relative horizontal displacements of the floors and $u(t) \in \mathbb{R}$ is the ground acceleration due to earthquake which is a measured input signal. Also, $M \in \mathbb{R}^{3\times 3}$ is the mass matrix, $D \in \mathbb{R}^{3\times 3}$ is the damping matrix, $S \in \mathbb{R}^{3\times 3}$ is the stiffness matrix, and $G \in \mathbb{R}^3$ is the loading vector. The parameter values of the system (16) are provided by [40]

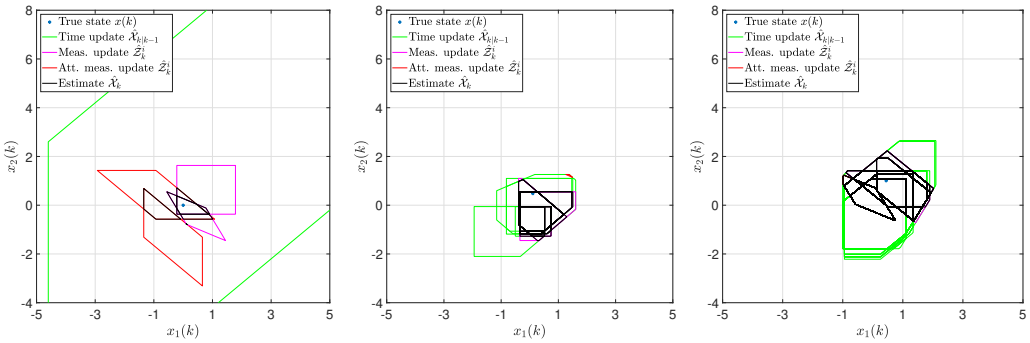

(a) Time $k = 1$. Sensors $\{2, 3\}$ are attacked. (b) Time $k = 2$. Sensors $\{3, 4\}$ are attacked. (c) Time $k = 3$. Sensors $\{1, 4\}$ are attacked.

Figure 2: Snapshots of estimated sets using Algorithm 1.

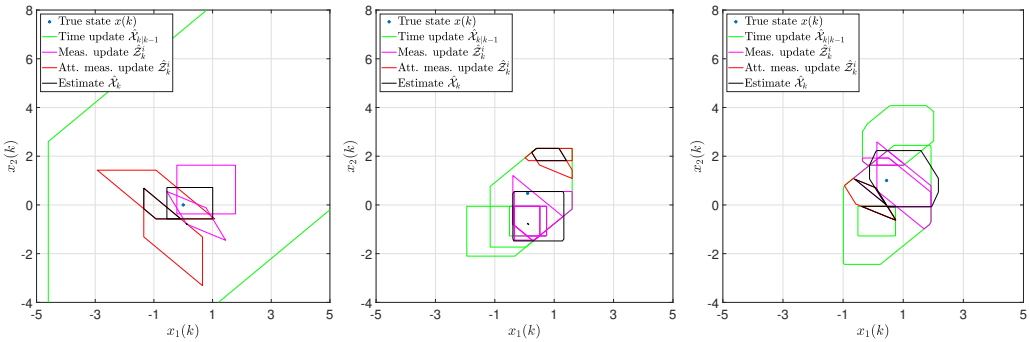

(a) Time $k = 1$. Sensors $\{2, 3\}$ are attacked. (b) Time $k = 2$. Sensors $\{3, 4\}$ are attacked. (c) Time $k = 3$. Sensors $\{1, 4\}$ are attacked.

Figure 3: Snapshots of estimated sets using Algorithm 1 while applying a complexity reduction technique to overapproximate the estimated intersecting sets by one set.

as:

$$
\begin{aligned}
M &= \mathrm{diag}([\ 478350 \quad 478350 \quad 517790\ ]) \quad (\mathrm{kg}) \\
D &= 10^5 \times \begin{bmatrix} 7.7626 & -3.7304 & 0.6514 \\ -3.7304 & 5.8284 & -2.0266 \\ 0.6514 & -2.0266 & 2.4458 \end{bmatrix} \quad (\mathrm{Ns/m}) \\
S &= 10^8 \times \begin{bmatrix} 4.3651 & -2.3730 & 0.4144 \\ -2.3730 & 3.1347 & -1.2892 \\ 0.4144 & -1.2892 & 0.9358 \end{bmatrix} \quad (\mathrm{N/m}) \\
G &= [\ 478350 \quad 478350 \quad 517790\ ]^\top \quad (\mathrm{kg}).
\end{aligned}
$$

By considering the state $x(t) = [\ q(t)^\top \quad \dot{q}(t)^\top\ ]^\top$, we can obtain the state-space representation in continuous time

$$\dot{x} = A_c x + B_c u$$

where

$$A_c = \begin{bmatrix} 0_{3\times 3} & I_3 \\ -M^{-1}S & -M^{-1}D \end{bmatrix}, \quad B_c = \begin{bmatrix} 0_{3\times 1} \\ -M^{-1}G \end{bmatrix}.$$

After discretization with sample time $\delta = 10^{-3}$, we obtain the system in the form (3a), where

$$A = \exp(A_c\delta), \quad B = A_c^{-1}(A - I_6)B_c.$$

Here, our goal is to monitor the building dynamics under an earthquake using secure set-based state estimation algorithm.

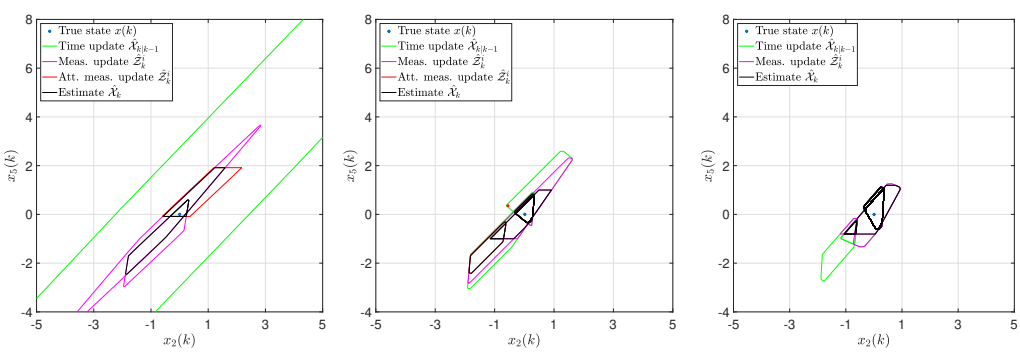

(a) Time $k = 1$, Sensor 2 attacked.

(b) Time $k = 2$, Sensor 3 attacked.

(c) Time $k = 3$, Sensor 1 attacked.

Figure 4: Snapshots of estimated sets using Algorithm 1.

We assume that each floor of the building is equipped with a sensor, i.e., $p = 3$, that measures the relative displacement and the velocity of that floor, which can be collected in the output vector $y_i(k) \in \mathbb{R}^3$ as given by (3b), for $i \in \mathbb{Z}_{[1,3]}$, where

$$C_1 = \begin{bmatrix} 1 & -1 & 0 & 0 & 0 & 0 \\ 1 & 0 & -1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \end{bmatrix}, \quad C_2 = \begin{bmatrix} -1 & 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & -1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \end{bmatrix}$$

$$C_3 = \begin{bmatrix} -1 & 0 & 1 & 0 & 0 & 0 \\ 0 & -1 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix}.$$

We suppose that the attacker can compromise only one sensor at each time, i.e., $q = 1$, whose rationale is explained below. Although the pairs (A, C_1) ,(A, C_2), and (A, C_3) are 'theoretically' observable, they are not 'practically' observable because the corresponding observability matrices have some singular values very close to zero. Therefore, in this example, we assume redundant observability from every pair of two sensors. Our goal is to monitor the building's floor displacements and velocities irrespective of the compromised sensor.

Suppose the dynamics are corrupted by the process noise $w(k)$, which is bounded by $\mathcal{W} = \langle 0, \sigma_{\mathcal{W}} I_6 \rangle$ with $\sigma_{\mathcal{W}} = 0.02$. Similarly, the sensor measurements are corrupted by noise $v_1(k)$, $v_2(k)$, and $v_3(k)$, respectively, which are bounded by $\mathcal{V}_i = \langle 0, \sigma_{\mathcal{V}} I_3 \rangle$ with $\sigma_{\mathcal{V}} = 1$.

To illustrate the efficacy of our algorithm, we apply a similar attack as in (15) to the sensors in which the attacker randomly chooses a sensor i at every time step k and injects false data into its measurement $y_i(k)$. In Fig. 4a, sensor 2 is compromised. We compute the estimated set (black), which contains the true state. Then, sensor 3 is attacked in Fig. 4b in which the time update set $\hat{\mathcal{X}}_{k|k-1}$ shrinks due to the progressive intersection. Finally, sensor 1 is attacked in Fig. 4c. The true state $x(k)$ remains enclosed by the estimated measurement update sets at every time step. Also, the estimation error remains bounded, and the attacker cannot destroy the accuracy of the set-based state estimate. Finally, we remark that the intersections depicted in the figure may not appear accurate because the figure shows a projection of six-dimensional sets on a two-dimensional plane.

7 Conclusion and Future Work

We presented a secure set-based state estimation algorithm that can ensure the inclusion of the true state of an LTI system in the estimated set even when all but one sensor could be compromised by an attacker. We achieved this by constructing agreement sets from the intersection of various observable combinations of measurement update sets. We showed that our algorithm guarantees the inclusion of the true state in the estimated set. Moreover, we proposed sufficient conditions for detecting, identifying, and filtering the attack signals and presented a simple algorithm to identify the set of compromised sensors at every time instant. The proposed algorithm may find applications in safety verification of safety-critical systems whose multiple sensors could be compromised by an attacker.

While our algorithm's worst-case complexity may increase with time under intelligently-designed stealthy attacks, we argued that it is challenging for attackers to execute such attacks due to the requirement of a complete understanding of the system and estimation algorithm and substantial computational resources. Nonetheless, we suggested various strategies to reduce the complexity of our algorithm to facilitate its implementation.

Our future work will focus on the set-based secure state estimation of nonlinear systems and developing a data-driven approach for secure estimation when the system model is unknown. We also highlight several open questions that remain unaddressed in this paper. For instance, devising better agreement protocols that can improve the lower bounds on the detectable, identifiable, and filterable attack signals will further strengthen the set-based estimation approach presented in this paper. The question of computational complexity under stealthy attacks remains to be addressed in a rigorous manner. The stability of the estimated set is only addressed for the case when strictly less than half of the sensors are attacked. When more than half of the sensors are attacked, it is challenging to upper-bound the estimated set by a contracting set. Finally, from the attacker's perspective, analyzing the complexity of generating stealthy attacks that result in the increased complexity of the set-based estimation is an interesting research problem.

Acknowledgment

This work has received funding from the European Union's Horizon Research and Innovation Programme under grant agreement No. 830927 and Marie Skłodowska-Curie grant agreement No. 101062523. It is also supported by the Swedish Research Council and the Knut and Alice Wallenberg Foundation, Sweden.

References

[1] Teodoro Alamo, José Manuel Bravo, and Eduardo F Camacho. Guaranteed state estimation by zonotopes. *Automatica*, 41(6):1035–1043, 2005.

[2] Amr Alanwar, Jagat Jyoti Rath, Hazem Said, Karl Henrik Johansson, and Matthias Althoff. Distributed set-based observers using diffusion strategies. *Journal of the Franklin Institute*, 360(10):6976–6993, 2023.

[3] Amr Alanwar, Hazem Said, and Matthias Althoff. Distributed secure state estimation using diffusion Kalman filters and reachability analysis. In *58th IEEE Conference on Decision and Control (CDC)*, pages 4133–4139, 2019.

[4] Matthias Althoff. An introduction to CORA 2015. In Goran Frehse and Matthias Althoff, editors, *ARCH14-15: 1st and 2nd International Workshop on Applied veRification for Continuous and Hybrid Systems*, volume 34 of *EPiC Series in Computing*, pages 120–151. EasyChair, 2015.

[5] Matthias Althoff, Goran Frehse, and Antoine Girard. Set propagation techniques for reachability analysis. *Annual Review of Control, Robotics, and Autonomous Systems*, 4(1):369–395, 2021.

[6] Matthias Althoff and Jagat Jyoti Rath. Comparison of guaranteed state estimators for linear time-invariant systems. *Automatica*, 130:109662, 2021.

[7] Joaquim Blesa, Vicenç Puig, and Jordi Saludes. Robust fault detection using polytope-based set-membership consistency test. *IET Control Theory & Applications*, 6(12):1767–1777, 2012.

[8] Pascal Bouron, Dominique Meizel, and Ph Bonnifait. Set-membership non-linear observers with application to vehicle localisation. In *European Control Conference (ECC)*, pages 1255–1260, 2001.

[9] Young Hwan Chang, Qie Hu, and Claire J Tomlin. Secure estimation based Kalman filter for cyber–physical systems against sensor attacks. *Automatica*, 95:399–412, 2018.

[10] Aijun Chen, Thach Ngoc Dinh, Tarek Raissi, and Yi Shen. Outlier-robust set-membership estimation for discrete-time linear systems. *International Journal of Robust and Nonlinear Control*, 32(4):2313–2329, 2022.

[11] Michelle S Chong, Henrik Sandberg, and Joao P Hespanha. A secure state estimation algorithm for nonlinear systems under sensor attacks. In *59th IEEE Conference on Decision and Control (CDC)*, pages 5743–5748, 2020.

[12] Michelle S Chong, Masashi Wakaiki, and Joao P Hespanha. Observability of linear systems under adversarial attacks. In *American Control Conference (ACC)*, pages 2439–2444, 2015.

[13] Alesi A de Paula, Guilherme V Raffo, and Bruno OS Teixeira. Zonotopic filtering for uncertain nonlinear systems: Fundamentals, implementation aspects, and extensions [applications of control]. *IEEE Control Systems Magazine*, 42(1):19–51, 2022.

[14] Hamza Fawzi, Paulo Tabuada, and Suhas Diggavi. Secure estimation and control for cyber-physical systems under adversarial attacks. *IEEE Transactions on Automatic Control*, 59(6):1454–1467, 2014.

[15] Randy Freeman and Petar V Kokotovic. *Robust nonlinear control design: state-space and Lyapunov techniques*. Springer Birkhäuser: Boston, MA, 2008.

[16] Xingkang He, Xiaoqiang Ren, Henrik Sandberg, and Karl Henrik Johansson. How to secure distributed filters under sensor attacks. *IEEE Transactions on Automatic Control*, 67(6):2843–2856, 2021.

[17] Luc Jaulin. Robust set-membership state estimation: Application to underwater robotics. *Automatica*, 45(1):202–206, 2009.

[18] Hakan Kayan, Matthew Nunes, Omer Rana, Pete Burnap, and Charith Perera. Cybersecurity of industrial cyber-physical systems: A review. *ACM Computing Surveys*, 54(11s):1–35, 2022.

[19] Amir Khazraei and Miroslav Pajic. Attack-resilient state estimation with intermittent data authentication. *Automatica*, 138:110035, 2022.

[20] Junsoo Kim, Chanhwa Lee, Hyungbo Shim, Yongsoon Eun, and Jin H Seo. Detection of sensor attack and resilient state estimation for uniformly observable nonlinear systems having redundant sensors. *IEEE Transactions on Automatic Control*, 64(3):1162–1169, 2018.

[21] Vu Tuan Hieu Le, Cristina Stoica, Teodoro Alamo, Eduardo F Camacho, and Didier Dumur. Zonotopic guaranteed state estimation for uncertain systems. *Automatica*, 49(11):3418–3424, 2013.

[22] Jin Gyu Lee, Junsoo Kim, and Hyungbo Shim. Fully distributed resilient state estimation based on distributed median solver. *IEEE Transactions on Automatic Control*, 65(9):3935–3942, 2020.

[23] Vuk Lesi, Ilija Jovanov, and Miroslav Pajic. Security-aware scheduling of embedded control tasks. *ACM Transactions on Embedded Computing Systems*, 16(5s):1–21, 2017.

[24] Jitao Li, Zhenhua Wang, Yi Shen, and Lihua Xie. Attack detection for cyber-physical systems: A zonotopic approach. *IEEE Transactions on Automatic Control*, 68(11):6828–6835, 2023.

[25] Xin Li, Guoliang Wei, and Licheng Wang. Distributed set-membership filtering for discrete-time systems subject to denial-of-service attacks and fading measurements: A zonotopic approach. *Information Sciences*, 547:49–67, 2021.

[26] Lei Liu, Lifeng Ma, Yiwen Wang, Jie Zhang, and Yuming Bo. Distributed set-membership filtering for time-varying systems under constrained measurements and replay attacks. *Journal of the Franklin Institute*, 357(8):4983–5003, 2020.

[27] Lei Liu, Lifeng Ma, Jie Zhang, and Yuming Bo. Distributed non-fragile set-membership filtering for nonlinear systems under fading channels and bias injection attacks. *International Journal of Systems Science*, 52(6):1192–1205, 2021.

[28] Nacim Meslem and Ahmad Hably. Robust set-membership state estimator against outliers in data. *IET Control Theory & Applications*, 14(13):1752–1761, 2020.

[29] Aritra Mitra and Shreyas Sundaram. Byzantine-resilient distributed observers for LTI systems. *Automatica*, 108:108487, 2019.

[30] M Umar B Niazi, Amr Alanwar, Michelle S Chong, and Karl H Johansson. Resilient set-based state estimation for linear time-invariant systems using zonotopes. *European Journal of Control*, 74:100837, 2023.

[31] Miroslav Pajic, Insup Lee, and George J Pappas. Attack-resilient state estimation for noisy dynamical systems. *IEEE Transactions on Control of Network Systems*, 4(1):82–92, 2016.

[32] Brenner S Rego, Guilherme V Raffo, Joseph K Scott, and Davide M Raimondo. Guaranteed methods based on constrained zonotopes for set-valued state estimation of nonlinear discrete-time systems. *Automatica*, 111:108614, 2020.

[33] Brenner S Rego, Stelios G Vrachimis, Marios M Polycarpou, Guilherme V Raffo, and Davide M Raimondo. State estimation and leakage detection in water distribution networks using constrained zonotopes. *IEEE Transactions on Control Systems Technology*, 30(5):1920–1933, 2021.

[34] Joseph K Scott, Davide M Raimondo, Giuseppe Roberto Marseglia, and Richard D Braatz. Constrained zonotopes: A new tool for set-based estimation and fault detection. *Automatica*, 69:126–136, 2016.

[35] Mariana Segovia-Ferreira, Jose Rubio-Hernan, Ana Cavalli, and Joaquin Garcia-Alfaro. A survey on cyber-resilience approaches for cyber-physical systems. *ACM Computing Surveys*, 56(8):1–37, 2024.

[36] Takumi Shinohara and Toru Namerikawa. Reach set-based secure state estimation against sensor attacks with interval hull approximation. *SICE Journal of Control, Measurement, and System Integration*, 11(5):399–408, 2018.

[37] Yasser Shoukry, Michelle Chong, Masashi Wakaiki, Pierluigi Nuzzo, Alberto Sangiovanni-Vincentelli, Sanjit A Seshia, Joao P Hespanha, and Paulo Tabuada. SMT-based observer design for cyber-physical systems under sensor attacks. *ACM Transactions on Cyber-Physical Systems*, 2(1):1–27, 2018.

[38] Yasser Shoukry, Pierluigi Nuzzo, Alberto Puggelli, Alberto L Sangiovanni-Vincentelli, Sanjit A Seshia, and Paulo Tabuada. Secure state estimation for cyber-physical systems under sensor attacks: A satisfiability modulo theory approach. *IEEE Transactions on Automatic Control*, 62(10):4917–4932, 2017.

[39] Haiyu Song, Peng Shi, Cheng-Chew Lim, Wen-An Zhang, and Li Yu. Set-membership estimation for complex networks subject to linear and nonlinear bounded attacks. *IEEE Transactions on Neural Networks and Learning Systems*, 31(1):163–173, 2019.

[40] Thao HT Truong, Peter Seiler, and Lauren E Linderman. Analysis of networked structural control with packet loss. *IEEE Transactions on Control Systems Technology*, 30(1):344–351, 2021.

[41] Xuejiao Yang and Joseph K Scott. A comparison of zonotope order reduction techniques. *Automatica*, 95:378–384, 2018.

[42] Sze Zheng Yong, Minghui Zhu, and Emilio Frazzoli. Switching and data injection attacks on stochastic cyber-physical systems: Modeling, resilient estimation, and attack mitigation. *ACM Transactions on Cyber-Physical Systems*, 2(2):1–2, 2018.

[43] Yilian Zhang, Yanfei Zhu, and Qinqin Fan. A novel set-membership estimation approach for preserving security in networked control systems under deception attacks. *Neurocomputing*, 400:440–449, 2020.

[44] Kemin Zhou and John Comstock Doyle. *Essentials of Robust Control*, volume 104. Prentice Hall: Upper Saddle River, NJ, 1998.

[45] Yanfei Zhu, Hang Liu, Chuanjiang Li, and Jiahao Yu. Consensus and security control of multi-agent systems based on set-membership estimation with time-varying topology under deception attacks. *International Journal of Control, Automation and Systems*, 20(11):3624–3636, 2022.