Efficient Predictive Monitoring of Linear Time-Invariant Systems Under Stealthy Attacks

Mazen Azzam¹, Liliana Pasquale², Gregory Provan³, and Bashar Nuseibeh¹

Abstract-Attacks on Industrial Control Systems (ICS) can lead to significant physical damage. While offline safety and security assessments can provide insight into vulnerable system components, they may not account for stealthy attacks designed to evade anomaly detectors during long operational transients. In this paper, we propose a predictive online monitoring approach to check the safety of the system under potential stealthy false data injection attacks on sensors. Specifically, we adapt previous results in reachability analysis for attack impact assessment to provide an efficient algorithm for online safety monitoring for Linear Time-Invariant (LTI) systems. The proposed approach relies on an offline computation of symbolic reachable sets in terms of the estimated physical state of the system. These sets are then instantiated online, and safety checks are performed by leveraging ideas from ellipsoidal calculus. We illustrate and evaluate our approach using the Tennessee-Eastman process. We also compare our approach with the baseline monitoring approaches proposed in previous work and assess its efficiency and scalability. Our evaluation results demonstrate that our approach can predict in a timely manner if a false data injection attack will be able to cause damage, while remaining undetected. Thus, our approach can be used to provide operators with real-time early warnings about stealthy attacks.

I. INTRODUCTION

Industrial Control Systems (ICS) denote systems where safety-critical physical processes are augmented with computation and communication capabilities, e.g. transportation systems, manufacturing, and chemical processes. Recently, the security of ICS has received increasing attention, especially with the rise in the number of attacks against these systems, e.g. Ukrainian power grid blackout [30]. Differently from attacks targeting IT systems, attacks against ICS can also cause physical damage, rather than only harming digital assets, e.g. sensitive data. In particular, stealthy attacks, where resourceful attackers exploit noise [4] or control theoretic properties [38] to avoid detection, can cause significant damage. Although a variety of techniques [15] consider the behaviour of the physical process to detect attacks on ICS, the detection of stealthy attacks still presents several limitations [16].

Assessing the risk of stealthy attacks involves performing offline impact assessment [36, 19, 31, 45, 33, 44], which may provide operators with more insight into potential vulnerabilities, such as the inability of a residual-based anomaly detector to detect certain sensor attacks before they cause damage. However, offline impact assessment cannot account for potential transients and variations in operating modes that a physical system may experience. In particular, chemical plants often experience long transients and frequent changes in operating conditions due to potential unforeseen disturbances, real-time optimisation modules, or high-level control decisions [34, 28]. Safety analysis consists of checking whether, from a current state, the system can enter an unsafe state given the current control settings. Conventional control methods cannot guarantee the safety of the system given the possibility of stealthy attacks to exploit noise or controltheoretic properties to avoid detection [25]. Therefore, there is a need for safety monitoring techniques that evaluate the safety of the system in real-time given such threats.

The objective of this paper is to develop an efficient online safety monitoring algorithm, specifically under stealthy False Data Injection Attacks (FDIA's) on sensors. We particularly consider a practical stealthy attack model where the attacker evades detection by ensuring that the false alarm rate of the anomaly detector is maintained [36]. To the best of our knowledge, only a few works [25, 11, 8, 13] fall into this line of research, and they either do not consider intelligently crafted stealthy attacks or are resource-intensive. Conversely, our approach to online safety monitoring provides a computationally-efficient online mechanism to detect the potential impact of a range of stealthy attacks, that would be undetectable using traditional monitoring approaches. In terms of efficiency and scalability, the main feature of our approach is to perform the most computationally intensive operations offline, and reduce the online safety checks to a computation of a distance measure. Namely, the intensive offline computations consist of approximating symbolic reachable sets of states. When deployed online, these sets only need to be instantiated depending on the current state estimate and a prediction of the state over a certain time horizon. We then take advantage of the geometric representation of the reachable sets to perform efficient safety checks. This general approach is inspired by the work by Chen et al. [9] in the context of real-time monitoring for simplex control architectures.

The main contribution of the current work is an efficient and scalable online safety monitoring algorithm for Linear, Time-Invariant (LTI) systems presented with the threat of stealthy attacks. The efficiency of the algorithm stems first from the offline pre-computation of a symbolic ellipsoid approximation of the reachable set using a Linear Matrix Inequalities (LMI) programme proposed by Murguia et

¹Mazen Azzam (mazen.azzam@ul.ie) and Bashar Nuseibeh (bashar.nuseibeh@ul.ie) are with Lero, The Irish Software Research Centre, University of Limerick, Limerick, Ireland.

²Liliana Pasquale (liliana.pasquale@ucd.ie) is with Lero, University College Dublin, Dublin, Ireland.

³Gregory Provan (g.provan@cs.ucc.ie) is with Lero, University College Cork, Cork, Ireland.

al. [36]. Second, we take advantage of results in ellipsoidal calculus [6] and the half-space geometrical representation of unsafe sets to perform fast real-time safety checks given an instantiation of the pre-computed reachable set at a predicted system state. When the algorithm is deployed online, it performs a prediction of the system state given the current state of the system and a specified time horizon. At each predicted state, the pre-computed reachable set is instantiated and emptiness checks of its intersection with the unsafe set are performed. The algorithm halts when a non-empty intersection is found or when the prediction time horizon is exhausted.

As a secondary contribution, we propose two online security metrics that can be computed by leveraging ellipsoidal calculus. The potential impact metric quantifies the potential impact of a stealthy attack. When the emptiness check returns a negative result, the intersection between the reachable set and the set of unsafe states can be approximated using an ellipsoid. We use the size of this ellipsoid to quantify the potential impact. When the intersection between the reachable set and the set of unsafe states is non-empty, it is also possible to compute the time-to-unsafe metric. This metric estimates the shortest time that an attacker would need to cause damage before being detected. This timeto-unsafe metric is fundamentally different from proximitybased metrics previously proposed in the literature [8, 7]. These metrics rely on the raw estimate of the state of the system to compute an euclidean distance to unsafe states, and are used to perform safety monitoring. Instead, our time-tounsafe metric relies on reachable sets induced by a potential stealthy attack. As such, we account for the fact that the given estimate may not represent the real state of the system.

Finally, we evaluate the proposed algorithm using the Tennessee-Eastman process (TEP) as a case study. We first validate our algorithm through extensive simulations aimed at assessing its ability to warn about potential damage due to a stealthy attack. Second, we compare it to existing online safety monitoring techniques for attacks, namely those that only rely on proximity-based metrics using raw state estimates. We show through simulation scenarios that under "low-and-slow" stealthy attacks, existing techniques will not convey the security and safety situation accurately. Conversely, our reliance on reachable sets in our approach allows for early warnings to be provided to operators before a stealthy attack can cause damage. Finally, we demonstrate the suitability of the algorithm to real-time applications. Specifically, we show that safety checking takes place in a time frame that is shorter than the system's sampling period, and that the algorithm scales well with the complexity of safety constraints and the desired length of time horizon for online prediction. We have applied our monitoring approach within a framework for physics-based early warnings for stealthy attacks [3].

The rest of this paper is organised as follows: Section II discusses related work, Section III provides an overview of our approach, Section IV describes the adopted modelling framework, Section V details the proposed algorithm, and

Section VI presents numerical simulation results. Finally, Section VII concludes the paper.

II. RELATED WORK

While there exists a large body of work on model-based attack detection in control systems [15], to the best of our knowledge, approaches tackling the problem of online safety monitoring for ICS under stealthy attacks are scarce. In model-based attack detection, we ask whether the current observations of the system are consistent with a mathematical model, up to a certain degree of uncertainty. Model-based attack detection is a reactive approach to security where the occurrence of an anomaly triggers an alert. With online safety monitoring under potential attacks, we ask whether the current (estimated) state of the system can be taken by an attack to a target state that violates at least one safety constraint. Online safety monitoring is a proactive and a predictive approach to safety/security that may help in guiding the selection of preemptive safety measures, such as switching to a safe and secure redundant controller [9].

Kwon et al. [25, 26] have proposed a recursive method to compute exact reachable sets under stealthy attacks online. While this method is computationally efficient, it uses large recursive matrices, which can make an extensive use of resources. Furthermore, safety checking in this work relies on the characterisation of a time-varying safe set as an ellipsoid centred at the current state. Although this is suitable for the Unmanned Aerial Vehicle (UAV) application used by the authors, it may not be applicable in chemical process control, where unsafe operating levels are usually fixed limits imposed on physical state variables. In contrast, the bulk of the computation required for our method is performed offline, resulting in symbolic sets with a lightweight characterisation when instantiated online. Furthermore, we consider more practical time-invariant unsafe sets which can be interpreted geometrically as a union of half-spaces.

Existing online monitoring schemes [13, 7, 11, 8] rely on a notion of proximity to a predefined set of unsafe/critical states. This line of work does not consider formal safety guarantees, but it uses metrics reflecting the proximity of the system to unsafe states as a way to either determine the level of safety or to detect attacks. For example, Carcano and Coletta [7, 11] compute the minimum Euclidean distance from current states to the unsafe operating region. Castellanos and Zhou [8] extend this notion further by computing an approximate "time-to-critical-states" metric. However, these approaches rely only on raw sensor values and do not consider the effect of stealthy attacks. For example, intelligently crafted sensor attacks introduce "low-and-slow" modifications to sensor values, which may eventually not reflect the real state of the system. In our work, instead of using raw sensor values, we rely on reachable sets under stealthy attacks which bound — with a certain confidence level — the actual state of the system.

Other related work [19, 31, 45, 33, 44] has proposed techniques to quantify the worst-case impact of potential stealthy attacks. To the best of our knowledge, these techniques are developed with the objective of performing risk assessment offline. For example, Milosevic et al. [32] propose a framework for security measure allocation given certain impact and attack complexity metrics. Murguia et al. [37] use the volume of ellipsoidal approximations of reachable sets under stealthy attacks as a measure of impact. In our work, we quantify in real-time the potential impact of a stealthy attack based on the size of the intersection of the reachable set with the set of unsafe states. Using this intersection, instead of the entire reachable set, gives a more precise estimate of potential impact. This is made possible by using the geometric properties of the sets' representations.

Finally, our approach is inspired by recent work on realtime reachability analysis [2, 9, 43], notably in the context of real-time monitoring for simplex control architectures [21]. Similarly to the algorithms proposed in this paper, a few works in this area consider pre-computing reachable sets offline before instantiating them online, in order to perform safety checking efficiently. However, to the best of our knowledge, previous work in this context did not consider safety checking in the presence of stealthy attacks seeking to cause damage to a safety-critical system.

III. PREDICTIVE ONLINE MONITORING APPROACH

In this section, we describe the online monitoring problem tackled in this paper and outline the proposed approach. We also describe the main idea behind existing work on online safety monitoring under attacks. We consider this body of work as a baseline against which we compare our approach.

A. Online Monitoring Problem

This paper considers the online monitoring framework depicted in Figure 1. The objective of our approach is to check in real-time whether a potential undetected attack can cause damage to the system before being detected. In other words, given the current physical state estimate $\hat{x}(t)$, the online monitoring problem asks whether there exists a stealthy false data injection attack on sensors that can bring the system into an unsafe state over the next T time instants.

If this check returns a negative result, then operators can be reassured that even if an intrusion is present, the alleged attacker may not be able to cause any damage without being detected. Otherwise, the check can serve as an early warning and can prompt operators to take preemptive safety or security measures. Such additional measures are however beyond the scope of the current work.

B. Outline of the Proposed Approach

The proposed approach, shown in Figure 1, is composed of two main steps:

• *Offline Initialisation:* this step consists of computing symbolic reachable sets under a stealthy false data injection attack in terms of the state estimate based on a model of the system and attack. This is possible by considering the evolution of the state estimation error under a stealthy attack instead of the physical state itself. As a result, we express the reachable set of errors

in terms of the actual state estimate. This allows us to perform the bulk of the computation offline, leading to more efficient real-time safety checking.

Online Monitoring: at runtime, the proposed monitor takes as input the current physical state estimate and the state of the controller and predicts the value of the state up to T time steps into the future. We assume that this prediction can be done using an identified physical model of the system, with T chosen to maintain an acceptable degradation in the confidence level of the predicted state. We then instantiate the precomputed symbolic reachable set at each predicted state value, and perform an emptiness check of its intersection with a predefined set of unsafe states. The prediction stops when a non-empty intersection is encountered, or when T is exhausted. Furthermore, we compute two security metrics when the intersection is non-empty: (i) potential impact of the attack, and (ii) time-to-unsafe states reflecting the approximate time that a potential attack would need to cause damage. These metrics are aimed at providing operators with a better assessment of the current safety/security situation.

In this paper, we apply this general approach to LTI systems, where we propose the use of ellipsoids as overapproximations of the symbolic reachable sets. Ellipsoids have been extensively used for safety verification for control systems [23, 20]. They feature an efficient quadratic representation in terms of the dimension of the state of the system [29], which presents an advantage in real-time monitoring. Furthermore, in most practical applications of process control, unsafe operating regions can be represented as unions of half-spaces. With reachable sets represented as ellipsoids, safety checking reduces to checking the sign of the distance from the ellipsoid to each of the hyperplanes composing the unsafe set [6, 23]. As a result, real-time safety monitoring is enabled with minimal resource utilisation.

In our application to LTI systems, we use results by Murguia [36] in reachability analysis under stealthy attacks to precompute a symbolic reachable set in the offline initialisation phase. We then use results in ellipsoidal calculus [23, 6] to design an efficient and scalable online safety monitoring algorithm for stealthy attacks.

C. Existing Monitoring Techniques

To the best of our knowledge, existing online safety monitoring techniques [13, 7, 11, 8] rely mainly on a proximity metric to assess the safety of the system against attacks targeting physical processes. In this paper, we compare our approach to these techniques which all feature the same idea detailed in the following. We provide a comparison based on intuition in this section.

Given a set of unsafe states S_u , and the current estimated state $\hat{x}(t)$, existing online monitoring techniques [13, 7, 11, 8] compute a distance metric $d_u = \min d(\hat{x}(t), S_u)$. The most commonly used distance metric is the Euclidean distance, which is mainly suitable for continuous or hybrid systems where state variables of interest assume real values.



Fig. 1. Outline of the proposed online monitoring approach.

Given a set of thresholds $\tau_1 > \tau_2 > \tau_3 \dots$, alarms of different levels of criticality can be raised based on the value of d_u . The proximity metric can further account for the dynamics of the system by computing a *time-to-unsafe/critical states* metric, $t_u = d_u/r$; where r is the approximate rate of change of the system state over a given time period [8]. In the rest of the paper, the term "traditional time-to-unsafe metric" refers to t_u computed using the aforementioned formula.

However, under a stealthy attack which slowly drives the system into unsafe states to avoid detection, the real state of the system x(t) will diverge from the estimated state. Thus, the metric d_u may not provide an accurate measure of the proximity of the system to unsafe states. Instead of relying merely on a proximity measure based on $\hat{x}(t)$, our online monitoring approach accounts for the possibility of stealthy attacks by considering reachable sets under such attacks. Namely, if $\mathcal{R}_{r}(t)$ is the reachable set of states under a stealthy attack given the used anomaly detector at time t, then $x(t) \in \mathcal{R}_x(t)$ even if $x(t) \neq \hat{x}(t)$. Otherwise, by definition, the attack would be detected by the anomaly detector. If over the next T time period, $\mathcal{R}_x(t') \cap \mathcal{S}_u \neq \emptyset$ for some $t' \in [t; t+T]$, then it is possible for an attacker to drive the system into the unsafe operating region within at least t'-ttime units. As such, our safety monitoring approach relies on emptiness checking of this intersection, instead of a mere proximity measure. Furthermore, we consider the proximity metric t' - t as an additional security metric that may assist operators in assessing the current safety/security situation.

IV. MODELLING FRAMEWORK

We apply our approach in this paper to Linear Time-Invariant (LTI) systems. Before detailing the proposed approach, we briefly describe the layout of the system and the attack model in this section.

A. System Layout

Consider the control system architecture in Figure 1. We assume that the physical system can be approximated using

a Linear Time-Invariant (LTI) model:

$$\begin{cases} x(k+1) = Ax(k) + B\bar{u}(k) + w(k) \\ y(k) = Cx(k) + v(k) \end{cases}$$
(1)

Where $x(k) \in \mathbb{R}^n$ is the state vector, $\bar{u}(k) \in \mathbb{R}^l$ denotes control signals received by the system, and $y(k) \in \mathbb{R}^m$ is a vector of sensor measurements. $w(k) \in \mathbb{R}^n$ and $v(k) \in \mathbb{R}^m$ denote process and measurement noise, respectively, and are assumed to follow a zero-mean Gaussian distribution with respective covariance matrices Σ_1 and Σ_2 . $k = t/\Delta_t \in \mathbb{N}$ denotes discrete time instants, where Δ_t is the sampling period and t is continuous time. A, B, and C are real, timeinvariant matrices of appropriate dimensions. The system is assumed to be equipped with an output feedback control loop, such that $u(k) = \mathcal{K}(\bar{y} - y_r(k))$, where u(k) denotes control signals originally sent to the process, \mathcal{K} is the control law, $y_r(k)$ denotes reference output, and $\bar{y}(k)$ denotes measurements received by the controller. In this work, we focus on sensor attacks, such that $u(k) = \bar{u}(k) \forall k$.

Furthermore, we assume that a subset of state variables, denoted as critical, are grouped in the vector $x_c = C_c x$, $x_c \in \mathbb{R}^{n_c}$, $C_c \in \mathbb{R}^{n_c \times n}$, and are required to remain within a certain safe set to ensure safe operation. Let S_u be the set of unsafe states; we assume that unsafe conditions are given as a linear combination of the critical state variables, such that the unsafe set becomes a union of half-spaces¹:

$$\mathcal{S}_u = \left\{ x(k) \in \mathbb{R}^n \mid \bigcup_{i=0}^{n_c} C_{c,i} x_i(k) \ge b_i \right\}$$
(2)

Where $b_i \in \mathbb{R}$ is the *i*th half-space scalar, $C_{c,i}$ denotes the *i*th row of the matrix C_c , and $x_i(k)$ the *i*th element of x(k).

At a time k, given previous sensor measurements and control actions, a Kalman filter generates an estimate of the

¹This assumption is typical in several process control applications.

physical state and expected sensor measurements as follows:

$$\begin{cases} \hat{x}(k) = A\hat{x}(k-1) + Bu(k-1) \\ + L(\bar{y}(k-1) - C\hat{x}(k-1)) \\ \hat{y}(k) = C\hat{x}(k) \end{cases}$$
(3)

Where L denotes the Kalman gain, and \hat{x} and \hat{y} denote estimated state and measurements, respectively.

In addition, a chi-squared anomaly detector compares received measurements with the generated estimate through a residual $r(k) := \bar{y}(k) - \hat{y}(k)$. Under nominal conditions, the residual metric has a zero-mean and a covariance matrix Σ . To check for this hypothesis, a chi-squared metric, $z(k) = r^T(k)\Sigma^{-1}r(k)$ is computed and compared with a threshold τ , such that exceeding this threshold implies a possible anomaly and raises an alarm. The threshold τ is designed to maintain a certain false alarm rate β such that $\Pr[z(k) \leq \beta] = 1 - \beta$ under nominal operation, and can be set as described by Murgia et al. [36] (Proposition 1).

B. Attacker Model

We consider in this work False Data Injection Attacks (FDIA's) on sensors which are masked by the system noise, in order to drive the latter slowly to the unsafe set (2). While different stealthy attack strategies exist in the literature, we choose to focus on one that is feasible in practice. Hence, we provide the following justification for our particular choice as opposed to other possible stealthy attacks:

- We consider FDIA's specifically on sensors as it has been shown both theoretically [27] and empirically [45] that their stealthiness is easier to maintain than attacks on actuators. FDIA's that consist of corrupting all sensor and actuator channels, also known as "covert" attacks [41], even though undetectable, require a significant amount of resources to be executed [42]. In practice, it may not be possible for an attacker to have simultaneous access to all sensor/actuator communication channels or associated control devices. Each sensor/actuator may require different kinds of software/hardware tools in order to corrupt the data they send or receive [1].
- 2) Replay attacks, similarly to FDIA's, threaten the integrity of sensor measurements. Although they require less resources to remain stealthy [42], the stealthiness of replaying old sensor measurements relies on whether these measurements could be admissible at a given time. Therefore, the system needs to be operating in a steady state for the replayed measurements to be considered nominal by the anomaly detector. If transients are experienced, then replaying old measurements would reveal the attack since they no longer correspond to the current control inputs. In this work, we consider process control systems whereby long transients may be experienced, thus limiting the extent to which stealthy replay attacks may be successful.

In our attack model, we assume that the attacker has sufficient resources and knowledge about the system, including knowledge of the system dynamics and anomaly detector properties. Let $\{k_s, \ldots, k_f\}$ denote the time period of the attack; we model the attack as a bias imposed on sensor measurements:

$$\bar{y}(k) := y(k) + \delta(k) \ \forall \ k \in \{k_s, \dots, k_f\}$$
(4)

The attack $\delta(k)$ remains stealthy by ensuring that the false alarm rate is maintained throughout the attack period. We use this characterisation because anomaly detectors do raise alarms under nominal operation. A sudden disappearance of these alarms in practice may raise suspicion in operators and lead them to uncover the attack before it can cause damage [18]. Under the attack in (4), the residual is given by: $r(k) = \bar{y}(k) - \hat{y}(k) = y(k) - \hat{y}k + \delta(k)$. As such, the chi-squared metric under attack is given by:

$$z(k) = (y(k) - \hat{y}(k) + \delta(k))\Sigma^{-1}(y(k) - \hat{y}(k) + \delta(k))$$
(5)

By selecting $\delta(k)$ to be such that $\Pr[z(k) \leq \tau] = 1 - \beta$, the attacker manages to remain undetected. For example, given K time steps, the attacker may choose to raise alarms for βK steps so that the false alarm rate is mimicked as closely as possible $[18]^2$. This is possible since we assume that the attacker knows the detector's parameters, i.e. Σ , β and τ , in addition to the system and estimator outputs; i.e. y(k) and $\hat{y}(k)$ respectively. In practice, such information could be obtained, for example, through reconnaissance attacks or insider knowledge. In addition, while this attack strategy is specifically designed for chi-squared detectors, most existing work in model-based anomaly detection employs this statistical change detection test [35].

We note that optimal stealthy FDIA strategies that take advantage of system noise are studied in the literature [17, 46, 47]. However, our choice for the previous attack model stems from its relative simplicity and practicality. In works describing optimal stealthy attacks, the attacker requires a deeper knowledge and understanding of the dynamics of the system. With the adversary model that we adopt, the attacker is required to only know the basic parameters of the system model and its anomaly detector, which are relatively easily obtained through a reconnaissance phase. Furthermore, optimising the attack strategy comes at a significant computational cost for the attacker and thus limits the attack's realtime performance, especially if the attacker needs to adapt to an adaptive defender. We also note that our approach requires only the attack to mimic the false alarm rate of the anomaly detector. As such, it can be applied to optimised attack strategies that also take advantage of the false alarm rate to remain stealthy.

V. PROPOSED MONITORING ALGORITHM

The proposed approach relies mainly on the offline computation of the symbolic reachable set of estimation error under the attack described in (4). This set is an ellipsoidal over-approximation of the exact reachable set, parametrised

²Due to space limitations, the reader is referred to [18] for a more detailed description of the distribution of the detector metric under such attack.

by the state estimate. In real-time, given the state estimate at time k, the symbolic set is instantiated at the K-step predicted state. Emptiness checks of its intersection with S_u are then performed. We detail both the offline and online computations in the following sections, and we propose online security metrics based on the computed reachable set.

A. Offline Computation of Symbolic Reachable Set

To compute the reachable set of the estimation error under the attack in (4), we use the method described by Murgia et al. [36]. We define this error as $e(k) := x(k) - \hat{x}(k)$, and assume that at the start of an attack the estimation error is always almost zero. The reachable set of the estimation error under the attack in (4) is independent of the actual physical state at the start of the attack. As such, this set serves as a symbolic reachable set parametrised by the state estimate.

By setting $e(k) = x(k) - \hat{x}(k)$, and performing some algebraic manipulations of Equation (3), the evolution of the estimation error under an attack is given by:

$$e(k+1) = Ae(k) - L(Ce(k) + v(k) + \delta(k)) + w(k)$$

= Ae(k) - L(y(k) - $\hat{y}(k) + \delta(k)) + w(k)$ (6)

Since the error is partially driven by the Gaussian noise w(k) and the attack-dependent sequence $\bar{\delta}(k) = y(k) - \hat{y}(k) + \delta(k)$, using a deterministic approach will yield an unbounded reachable set, as the support of w(k) and $\delta(k)$ (as characterised in (4) and (5)) is infinite. This issue can be overcome by setting a confidence level on the energy of both of these vectors. For the attack, the sequence $\bar{\delta}(k)$ is already constrained to be such that $\Pr[z(k) \leq \tau] = \Pr[\|\Sigma^{-1/2}\bar{\delta}(k)\|^2 \leq \tau] = 1 - \beta$ where $\|.\|$ denotes the L_2 -norm. For the noise, let $p = \Pr[\|w(k)\|^2 \leq \bar{w}]$; since w(k) follows a zero-mean Gaussian distribution, the bound \bar{w} on $\|w(k)\|^2$ can be determined using the gamma distribution for a desired confidence p.

By using this assumption, the resulting reachable set can be interpreted as a level set of the distribution of the reachable error. A larger confidence level would lead to a larger set, at the cost of being overly conservative with the safety checking. A reasonable choice for p would be $1-\beta$, as the false alarm β is designed to be small. This also simplifies the computation of the reachable set, since for $p = 1 - \beta$, we readily have $\Pr[||w(k)||^2 \leq \bar{w}] = \Pr[z(k) \leq \tau]$ under the attack in (4). The following is based on this choice; for a more detailed treatment of this confidence level and a comparison of reachable sets under different choices of p, the reader is referred to [36] and [18] due to space limitations. Let \mathcal{R}_e^p denote the reachable set of error under the attack in (4) and a confidence level $p = 1 - \beta$:

$$\mathcal{R}_{e}^{p} := \{ e(k) \in \mathbb{R}^{n} \mid e(k) \text{ is s.t. (6)}, \\ p = \Pr[\|w(k)\|^{2} \le \bar{w}] = 1 - \beta \}$$
(7)

While computing \mathcal{R}_e^p is intractable, it is possible to overapproximate the set using an ellipsoid in \mathbb{R}^n , given by:

$$\mathcal{R}_e^p \subseteq \mathcal{E}_e^p = \{ e(k) \mid e^T(k) \mathbf{\Pi}^{-1} e(k) \le 1 \}$$
(8)

Algorithm 1 Offline Symbolic Reachable Set Computation INPUTS: $(A, L, \Sigma, \tau, \bar{w}, \Delta h); 0 < \Delta h < 1$

OUTPUT: Reachable set shape matrix Π

- 1: $b \leftarrow \Delta h$;
- 2: SolutionList $\leftarrow EmptyList();$
- 3: while b < 1 do
- $\triangleright \text{ Solve the programme in (9) for the current value of } b$ 4: SolveSemiDefiniteProgramme(A, L, Σ, τ, \bar{w}, b);
- 5: SolutionList.*append*(CurrentSolution);
- 6: $b \leftarrow b + \Delta h;$
- 7: end while
- 8: BestShapeMatrix *← MinObjectiveValue*(SolutionList);
- 9: return BestShapeMatrix;

Where the positive definite matrix Π is the ellipsoid's shape matrix. Letting $\mathcal{P} = \Pi^{-1}$, the minimum volume ellipsoid containing the set \mathcal{R}_e^p can be obtained by solving the following semi-definite programme [36]:

$$\mathcal{P} = \arg \min - \log \det \mathcal{P}$$
s.t.
$$\mathcal{P} > 0 : \mathcal{Q} > 0$$
(9)

Where:

$$Q = \begin{bmatrix} b\mathcal{P} & A^{T}\mathcal{P} & \mathbf{0} & \mathbf{0} \\ \mathcal{P}A & \mathcal{P} & \mathcal{P} & -\mathcal{P}L\Sigma^{1/2} \\ \mathbf{0} & \mathcal{P} & \frac{1-b}{\tau+\bar{w}}I & \mathbf{0} \\ \mathbf{0} & -\Sigma^{1/2}L^{T}\mathcal{P} & \mathbf{0} & \frac{1-b}{\tau+\bar{w}}I \end{bmatrix}, \quad (10)$$
$$b \in (0,1)$$

Note that while b is an optimisation variable, it is necessary to fix it to ensure the convexity of the programme. A grid search can then be performed over the interval (0, 1) to find the optimal shape matrix corresponding to the minimumvolume ellipsoid.

Given the shape matrix $\mathbf{\Pi} = \mathcal{P}^{-1}$, and replacing e(k) by its definition, we obtain a symbolic ellipsoidal approximation $\mathcal{E}_x^p(k)$ of the reachable set $\mathcal{R}_x^p(k)$ of the actual system state x(k), parametrised by the current state estimate $\hat{x}(k)$:

$$\mathcal{R}_x^p(k) \subseteq \mathcal{E}_x^p(k) = \{x(k) \in \mathbb{R}^n \mid (x(k) - \hat{x}(k))^T \mathbf{\Pi}^{-1}(x(k) - \hat{x}(k)) \le 1\}$$
(11)

Algorithm 1 summarizes the offline steps to obtain $\mathcal{E}_x^p(k)$. Given the system matrix A, the Kalman gain L, the residual covariance matrix Σ , the anomaly detector's threshold τ and the confidence bound \bar{w} , the algorithm performs a grid search over (0, 1) by partitioning the interval into segments of length Δh . The choice of Δh will depend on the desired tightness of the ellipsoidal approximation given the computational resources available. Note that this step only needs to be performed offline once, and only the matrix Π needs to be stored to instantiate $\mathcal{E}_x^p(k)$ online given a state estimate $\hat{x}(k)$.

Algorithm 2 Online Safety Checking

INPUTS: $(K, \Pi, \hat{x}(k))$, ControllerState, UnsafeSet) OUTPUT: *true* if the system is safe under a potential stealthy attack; *false* otherwise

1:	$\hat{x}_p \leftarrow \hat{x}(k)$
2:	for all $l \in \{0, 1,, K\}$ do
3:	ReachEll $\leftarrow Ellipsoid(\hat{x}_p, \Pi);$
4:	for all Hyperplane \subset UnsafeSet do
5:	DistToUnsafe \leftarrow <i>dist</i> (ReachEll,Hyperplane);
6:	if DistToUnsafe ≤ 0 then
7:	return <i>false</i> ;
8:	end if
9:	end for
10:	$\hat{x}_p \leftarrow PredictControlFlow(\hat{x}_p, ControlState);$
11:	end for
12:	return true;

B. Online Safety Checks

Algorithm 2 outlines the steps needed to perform online safety checks. Given the current state estimate $\hat{x}(k)$ and the state of the controller, we estimate the state of the system for K time steps into the future using the identified model of the plant. At each time step $l \in \{0, \ldots, K\}$, we instantiate $\mathcal{E}_x^p(k+l)$ and we check whether it intersects the set \mathcal{S}_u . The algorithm halts and reports an unsafe state when a nonempty intersection is encountered. If the prediction horizon is exhausted, the algorithm reports a safe state. In the following, we detail the procedure we use to perform the emptiness checks.

Let $\mathcal{H}_i = \{x \in \mathbb{R}^n \mid C_{c,i}x \geq b_i\}$ be a half-space representing one of the safety conditions composing the set \mathcal{S}_u (Equation (2)). Checking whether $\mathcal{E}_x^p(k_f) \cap \mathcal{S}_u = \emptyset$ involves checking whether $\mathcal{E}_x^p(k_f) \cap \mathcal{H}_i = \emptyset$ for each $i \in \{1, \ldots, n_c\}$. If the latter is true for all i, then the former is also true, since $\mathcal{S}_u = \bigcup_{i=1}^{n_s} \mathcal{H}_i$.

To check whether $\mathcal{E}_x^p(k_f) \cap \mathcal{H}_i = \emptyset$, it suffices to compute the minimum distance from $\mathcal{E}_x^p(k_f)$ to the hyperplane that delimits the half-space \mathcal{H}_i . Let $\mathcal{H}_{p,i} = \{x \mid C_{c,i}x = b_i\}$ be such hyperplane, the minimum distance from $\mathcal{E}_x^p(k_f)$ to $\mathcal{H}_{p,i}$ is given by [24]:

$$d_i(k_f) = \frac{|b_i - C_{c,i} x(k_f)| - \sqrt{x(k_f)^T \mathbf{\Pi} x(k_f)}}{\|C_{c,i}^T\|}$$
(12)

If $d_i(k_f) \leq 0$, then $\mathcal{E}_x^p(k_f) \cap \mathcal{H}_i \neq \emptyset$. Otherwise, if $d_i(k_f) \geq 0$, then the ellipsoid $\mathcal{E}_x^p(k_f)$ is either contained in \mathcal{H}_i or does not intersect the half-space, depending on whether its centre $\hat{x}(k_f)$ belongs to \mathcal{H}_i . However, since the state estimate is within the safe region³, i.e. $\hat{x}(k_f) \notin \mathcal{H}_i$, then in our case, $d_i(k_f) > 0$ always implies that $\mathcal{E}_x^p(k_f) \cap \mathcal{H}_i = \emptyset$.

C. Real-time Security Metrics

In addition to checking the emptiness of the intersection of the reachable set with the set of unsafe states, it is possible to derive two online security metrics. The first metric can help operators get a better idea of the potential impact of a stealthy false data-injection attack, while the second approximates minimum amount of time that would be required for an attacker to cause damage. In this section we show how ellipsoidal methods can be used to compute such metrics efficiently.

1) Real-Time Impact of Stealthy Attack: In the case where $\mathcal{E}_x^p(k_f) \cap \mathcal{S}_u \neq \emptyset$, we can quantify the impact of a potential stealthy false data-injection attack using the approximate size of this intersection. Namely, for each half-space $\mathcal{H}_i \subset \mathcal{S}_u$ it is possible to over-approximate $\mathcal{E}_x^p(k_f) \cap \mathcal{H}_i$ using a minimum-volume ellipsoid $\mathcal{E}_i(k_f)$ of centre $q_i(k_f) \in \mathbb{R}^n$ and shape matrix $\Pi_i(k_f)$, as follows [6]:

$$q_i(k_f) = \hat{x}(k_f) - \frac{1 + \alpha_i n}{n+1} \mathbf{\Pi} \bar{c}_i$$

$$\mathbf{\Pi}_i(k_f) = \frac{n^2 (1 - \alpha_i^2)}{n^2 - 1} \times$$

$$\left(\mathbf{\Pi} - \frac{2(1 + \alpha_i n)}{(n+1)(\alpha_i + 1)} \mathbf{\Pi} \bar{c}_i \bar{c}_i^T \mathbf{\Pi} \right)$$
(13)

Where $\bar{c}_i = C_{c,i}/(C_{c,i}\Pi C_{c,i}^T)^{0.5}$ and $\alpha_i = (C_{c,i}\hat{x}(k_f) - b_i)/(C_{c,i}\Pi C_{c,i}^T)^{0.5}$. As such, we quantify the impact of a potential stealthy false data-injection attack using the volume of $\mathcal{E}_i(k_f)$. The volume of a general ellipsoid in \mathbb{R}^n with a shape matrix Q is given by:

$$\mathbf{vol}(\mathcal{E}) = \mathbf{vol}[\mathcal{B}_n] \sqrt{\det Q} \tag{14}$$

Where $\operatorname{vol}[\mathcal{B}_n]$ and $\det Q$ denote the volume of the unit n-ball and the determinant of the matrix Q, respectively. It is worthwhile to note that different system dimensions may lead to vastly different number ranges for the volume of the intersection ellipsoid. Thus, in order to make the impact metric more meaningful, we propose to use the ratio of the volume of the intersection ellipsoid to that of the ellipsoid approximating the reachable set. This guarantees that the impact metric will fall in the range [0;1], thus becoming more intuitive to interpret. From (14), the impact metric reduces to the following:

$$\mathbf{Im}(k) = [\max_{i=1...n_c} \det \mathbf{\Pi}_i(k_f)] / \det \mathbf{\Pi}$$
(15)

2) Approximate Time to Unsafe States: In the case $\mathcal{E}_x^p(k_f) \cap \mathcal{S}_u \neq \emptyset$ for some $k_f \in \{k, \ldots, k+K\}$, we use the time k_f as the approximate time to unsafe states metric in our approach, namely:

$$\mathbf{Tc}(k) = (k_f - k)\Delta_t,$$

$$\exists k_f \in \{k, \dots, k + K\} \text{ s.t. } \mathcal{E}_x^p(k_f) \cap \mathcal{S}_u \neq \emptyset \quad (16)$$

Where Δ_t is the system's sampling period. The advantage of using k_f instead of the distance from the state estimate $\hat{x}(k)$ itself is that the former approach accounts for the fact that if an undetected attack is present, then the actual

 $^{^{3}}$ Otherwise, it would be clear that the system is evolving to an unsafe state and Algorithm 2 in this case would become obsolete.

 TABLE I

 Safety constraints considered for the TE case study [12].

Output	Low Limit	High Limit	
Reactor Pressure	none	2895 kPa	
Reactor Temperature	none	150 °C	
Reactor Level	$11.8 \ m^3$	$21.3 m^3$	
Product Separator Level	$3.3 m^3$	9.0 m^3	
Stripper Base Level	$3.5 m^3$	$6.6 m^3$	

state $x(k_f) \neq \hat{x}(k_f)$ still lies within $\mathcal{E}_x^p(k_f)$ (otherwise, the attack would be detected). This metric provides operators with an idea of the minimum time they have to react before a potential stealthy false data injection attack manages to bring the system into an unsafe state.

VI. EVALUATION

In this paper, we use the Tennessee-Eastman Process (TEP) with the control architecture designed by Ricker et al. [39] as a case study. This benchmark process is widely regarded as one that reflects to a high degree of accuracy a real-life chemical process [22]. In addition, it has been used widely to test ideas in process control [40] as well as in model-based approaches to ICS security [14].

Namely, we used the simulation written in Simulink by Bathelt et al. [5]. While the controllers are implemented as Simulink blocks, the physical process itself is simulated in continuous-time and is written in C and incorporated into Simulink using MATLAB's S-function blocks. Rate transition blocks are in turn used to simulate the discrete-time sampling of sensor measurements and actuator signals by digital controllers, in a fashion that mimics real-life situations in process control systems. We implemented Algorithms 1 and 2 in MATLAB, with the ellipsoidal techniques based on the Ellipsoidal Toolbox written by Kurzhanskiy et al. [24]. We approximated the TEP process as an LTI system with 50 state variables using MATLAB's n4sid algorithm. We considered the safety constraints discussed by Down and Vogel [12], shown here in Table I. We augmented the TEP Simulink simulation with a Kalman filter and a chi-squared anomaly detector. To initialise the online monitoring tool. we ran Algorithm 1 to determine the reachable ellipsoid's shape matrix. We performed the grid search for parameter bby dividing the interval into segments of length 0.01.

Our evaluation consists of three main parts. First, we validated our approach by measuring true and false positives rates using extensive simulations. Second, we compared our approach with existing online monitoring approaches. Finally, we assessed the performance and scalability of our approach.

A. Validation

The objective of our approach is not to detect attacks, but rather to perform safety checking under *potential* stealthy attacks that seek to cause damage. Namely, Algorithm 2 checks whether the current state of the system can be taken



Fig. 2. True/false positive/negative rates as a function of the length of prediction horizon K.

to an unsafe state by a stealthy attack within the next K time instants and cause damage before the anomaly detector detects the attack. Thus, to evaluate our approach, we ran several simulations of the TEP only considering different stealthy attacks on the sensors that report values of safety critical parameters shown in Table I. We avoided using the true/false positive/negative rate performance metrics as traditionally defined in the attack detection literature. Instead, we considered a notion of true/false positives/negatives similar to the one adopted in previous work on online safety monitoring [10]:

- A *true positive* occurs when Algorithm 2 raises a warning within K time instants before damage occurs due to an attack, and the system reaches an unsafe state before the anomaly detector raises an alarm.
- A *true negative* occurs when Algorithm 2 does not raise any warnings within K time instants before damage occurs, but the attack is detected by the anomaly detector.
- A *false positive* occurs when Algorithm 2 reports an unsafe state within K time instants before damage occurs, but the anomaly detector manages to detect the attack before the system enters the unsafe state.
- A *false negative* occurs when Algorithm 2 does not raise any warnings within K time instants before damage occurs, even if a stealthy attack is taking place, and the anomaly detector does not raise any alarm.

We first tested the effect of the length of the prediction horizon K on these rates, with results presented in Figure 2. For each value of K that we tested, we ran 500 simulations where we picked the attacked sensors at random, and we simulated the attack as a slowly growing bias on sensor measurements.

We can see from Figure 2 that for a small prediction horizon length, Algorithm 2, returns mostly negative checks, with true and false negatives accounting for the vast majority of predictions for K < 500. As K grows, the number of true and false positives increases, with the false positive



Fig. 3. Receiver Operating Characteristic (ROC) curve for Algorithm 2 with the length of prediction horizon K as the third dimension.

rate increasing in a much slower manner. For $K \ge 1500$, although the rate of true positives is high, Algorithm 2 returns a high number of false positives as well. This behaviour is the result of the design of Algorithm 2. First, for small K, the algorithm will likely not be exploring a sufficient number of states where a stealthy attack would cause a violation of safety constraints. Thus, it is expected to observe a high rate of both false and true negatives, with true and false positive rates remaining very low. As K increases, the algorithm is allowed to explore more states, therefore increasing the number of true positives. The slow parallel increase in false positives shows that Algorithm 2 exhibits high accuracy for intermediate values of K. However, at high values of K, the accuracy of the predicted states is expected to decrease, which explains the high false positive rates.

These simulations show that there exists a trade-off between how early we would like to raise warnings about potential safety violations due to a stealthy attack and the accuracy of Algorithm 2. These experiments can also serve as a method to tune the choice of K. To showcase these ideas, we have plotted in Figure 3 the Receiver Operating Curve (ROC) for Algorithm 2, with the length of prediction horizon K as the third dimension. We note that it exists a "cut-off" point at K = 500 time steps where we obtain acceptable values for the true/false positive rates (90.8% for true positives, 4.05% for false positives). This is equivalent to about 15 minutes ahead-of-time prediction, which is a reasonable choice in practice for K.

For K = 500, we tested the accuracy of Algorithm 2 under different numbers of sensors being attacked at the same time. We ran 500 simulations for each different number of sensors being attacked. In each simulation, we picked the attacked sensors at random, and we ran Algorithm 2 while considering the safety constraints associated with the sensor(s) under attack (Table I). The results in Table II show high true positive and low false positive rates in each case.

TABLE II True positive/negative rates vs. the number of sensors attacked at the same time.

Attacked Sensors	1	2	3	4	5
True Positive Rate	0.908	0.905	0.91	0.905	0.92
False Positive Rate	0.0405	0.041	0.04	0.0408	0.04

These experiments demonstrate the accuracy of Algorithm 2 with respect to all safety-critical sensors. Given the large number of random simulations we ran, we can conclude that Algorithm 2 can report potential safety violations due to a stealthy attack with respect to all the safety constraints imposed on the system.

B. Comparison with Existing Monitoring Approaches

In this section, we empirically showcase the usefulness of our approach compared to existing monitoring approaches described in Section III-C. To this end, we implemented an online monitoring tool measuring the time-to-unsafe states metric based on the euclidean distance from the current state estimate to the set of unsafe states. We also measure the average rate of change of the estimated system state. We avoided a comparison based on the accuracy metrics depicted in the previous section. This is due to the fact that the traditional time-to-unsafe states metric relies on the selection of different thresholds to raise alarms of different criticality. With the lack of precise methods to select these thresholds, it is hard to perform a meaningful quantitative comparison between the metric proposed in this paper and the traditional one. Therefore, we used a set of attack scenarios on safety critical sensors to empirically demonstrate the advantages of our approach. We particularly focus on attacks targeting sensors with a slowly growing bias.

We chose three attack scenarios. Scenarios 1 and 2 depict individual attacks targeting the level and pressure sensors, respectively, of the main reactor in the TEP. Scenario 3 depicts an attack performed simultaneously on the level, pressure, and temperature sensors of the reactor. We chose Scenarios 1 and 2 to illustrate the typical kind of attacks targeting safety-critical sensors, individually. We can obtain similar results for individual attacks on other safety-critical sensors. Scenario 3 illustrates a more dangerous coordinated attack where all main reactor sensors in the TEP are manipulated at the same time. Again, similar results can be obtained for other combination of sensors for safety-critical variables listed in Table I. Figures 4, 5 and 6 present the results obtained for each scenarios.

1) Scenario 1: In this scenario, we simulate an attack on the reactor's level sensor, where a growing bias on level measurements is introduced to trick the controller into overflowing the reactor. This simulation is shown in Figure 4 where we can see that the anomaly detector raises an alarm almost at the moment the overflow takes place. Our impact metric increases significantly over the period preceding the physical damage to the reactor, and our time-to-unsafe states



Fig. 4. Results from Simulation Scenario 1.



Fig. 5. Results from simulation Scenario 2.

metric shows that the plant is dangerously close to damage. Conversely, the traditional time-to-unsafe states computed based on the estimated state alone shows a slight increase, indicating that the plant appears to be moving away from the unsafe operating region.

2) Scenario 2: In this scenario, we perform an attack on the reactor's pressure sensor. Namely, the attack is a slowly growing bias on pressure measurements seeking to trick the pressure controller into increasing the pressure inside the reactor while it appears lower than its set-point. We can see from Figure 5 that the anomaly detector fails to raise any alarm before excessive pressure builds up in the reactor. However, our online monitoring algorithm reports that the plant's operation may be unsafe in the presence of an attack throughout this ramp-down operation, which is shown by a non-zero impact metric. While the impact metric shows an increase over the few hours between the start of the attack and the damage taking place. Instead, the traditional timeto-unsafe states metric shows the plant moving away from unsafe states. 3) Scenario 3: In this scenario, we simulate simultaneous attacks on the main reactor's pressure, temperature, and level sensors. All three attacks are slowly growing biases. Figure 6 shows that damage occurs faster in this scenario than in the previous two, with the anomaly detector again failing to raise any alarms. Our impact metric however shows again that the plant's operation may be unsafe under the attack. Instead, the traditional time-to-unsafe states metric depicts the plant moving away from the unsafe operating region.

These scenarios demonstrate the usefulness of our approach in the presence of stealthy attacks when compared to simple distance-to-unsafe metrics. Relying on the traditional distance-to-unsafe metric may relay an inaccurate idea of the current security or safety conditions. This was especially highlighted in Scenarios 2 and 3. While the plant appears to drift away from the unsafe operating region, our monitoring approach can still warn operators that an attacker is able to damage the system without being detected.

Figure 7 shows a comparison between the time-to-unsafe metric computed using our algorithm and the same metric computed based on the raw estimated state. Namely, we plot the difference (error) between the metric in each case and the time-to-unsafe states computed based on the real state of the system. In each scenario, we observe that the metric based on the raw estimated state is relatively accurate before the attack starts (the error is close to zero). However, the error starts to grow as the stealthy attack progresses and the real state diverges from its estimate. Conversely, as the stealthy attack progresses, this error decreases for the timeto-unsafe states metric computed according to our algorithm and reaches almost zero towards the end of the attack.

This demonstrates the usefulness of our algorithm in the worst case, where the actual state of the system significantly diverges from the real estimate under a stealthy attack. While this may be overly conservative when the system is not under attack, the safety criticality of the systems we consider justifies the need to employ a monitoring that can generate early warnings when a dangerous stealthy attack is taking place. Hence, the task of pre-computing reachable sets under attacks and using them for online monitoring is well justified.

Our monitoring approach can be used as part of an early warning system to improve situational awareness and potentially preserve important data relevant to an investigation should a stealthy attack indeed cause damage. A full treatment of this aspect of our approach is however deferred for future work.

C. Performance and Scalability

We ran 100-hour simulations of the TEP to test the performance of our approach. For performance testing purposes, we modified the algorithm so that all the states along the prediction horizon are visited. This represents the worst-case scenario. We tested the ability of Algorithm 2 to scale with respect to (i) the length of the prediction horizon (number of steps K), and (ii) the number of safety constraints. We performed this testing on a machine with an Intel i7-9750H CPU clocked at 2.6 GHz and with 16 GB of RAM.



Fig. 6. Results from simulation Scenario 3.



Fig. 7. Difference (error) between time-to-unsafe metric computed based on state estimate and based on the proposed algorithm vs. the time-to-unsafe states based on the real state of the system; (a) - Scenario 1, (b) - Scenario 2, (c) - Scenario 3.

Figure 8 shows the average time needed to perform state prediction and safety checking (Algorithm 2) for each number of steps, with a fixed number of safety constraints ⁴. Results show that the performance of Algorithm 2 scales linearly with the number of steps required for state prediction. Furthermore, at K = 1000 steps, equivalent to approximately 30 min ahead-of-time prediction, the execution time of the safety checking algorithm is smaller than the sampling period (1.8 sec).

In addition, we tested the ability of our algorithm to scale with the number of safety constraints. To this end, we fixed K at 500 steps (i.e. ≈ 15 min) and we ran 100-hour simulations for each number of safety constraints. For the purposes of testing, we generated random half-spaces representing safety constraints. Results are shown

 $^4 {\rm For}$ a 100-hour simulation, this average is taken over approximately 2×10^5 safety checks for a sampling period of 1.8 sec.

in Figure 9. The execution time scales linearly with the number of safety constraints, and with 500 safety constraints is still less than the sampling time of the system. Hence, the proposed algorithm exhibits excellent real-time performance in the presence of more complex safety constraints.

It is worth noting that the performance of the proposed algorithm can be improved significantly if implemented with a compiled language such as C++ instead of MATLAB. This is the case in most control systems applications.

D. Implementation Challenges

Our evaluation has shown that the proposed monitoring approach exhibits high accuracy and real-time performance. When it comes to its implementation in practice, we note that the approach relies on a model of the system and its anomaly detector, which is normally available during the control design phase. The stealthy attack model that we focus



Fig. 8. Average execution time of the proposed scheme vs. the number of steps for prediction.



Fig. 9. Average execution time of the proposed scheme vs. the number of safety constraints.

on encompasses several optimised stealthy attack strategies proposed in previous work. This is also evident by the ability of Algorithm 2 to exhibit high accuracy with respect to a large number of randomised attacks in our validity tests (Section VI-A). Finally, the attack model's relative practicality with respect to the attacker makes it a reliable heuristic threat model to consider for a practical design of monitoring algorithms.

In our approach, we require some design parameters to be set. For Algorithm 1, a choice for the size of the partitioning interval Δ_h must be made to guide the grid search. This parameter depends entirely on the desired tightness of the reachable set approximation. As this is a design-time activity, the computational cost of choosing an arbitrarily small Δ_h could be negligible when considering the increased tightness of the resulting approximation. Our experiments have shown that even with a choice of $\Delta_h = 0.01$, we can achieve a high accuracy in terms of safety checks. For Algorithm 2, the design choice is to set a proper length for the prediction horizon K. We showed that extensive randomised simulations performed in Section VI-A can help in choosing a value for K that achieves an acceptable trade-off between how "early" we would like to raise warnings vs. the accuracy of Algorithm 2 (Figures 2-3).

VII. CONCLUSION

In this paper we have presented a predictive online safety monitoring approach for LTI systems under potential stealthy sensor attacks. Our approach precomputes offline symbolic reachable sets in terms of the system's state estimate, by considering the evolution of the estimation error under a potential stealthy attack. Given the current state of the system and controllers, we predict in real time the control flow of the system for a certain umber of steps in the future. The precomputed sets are then instantiated at the predicted estimates. We use ellipsoidal calculus techniques to perform emptiness checks of the intersection of the precomputed set with a set of unsafe states. We applied the approach to the large-scale Tennessee-Eastman process (TEP) where we validated our approach and we showed that it can perform safety checks in a timely manner. Furthermore, we demonstrated the improvement over existing online monitoring techniques and we showed that the computation of reachable sets under stealthy attacks is well justified in safety-critical applications. In the future, we will study in more detail the uncertainty propagation caused by the prediction of future states and its effect on the validity of the safety checking. In addition, we will study the possibility of extending the proposed online monitoring approach to other attack models such as the replay or the covert attack, particularly in situations where it is justifiable to consider more resource-intensive stealthy attacks from the attacker's point of view. We also plan to consider systems that exhibit a high degree of non-linearity and may not be modelled in the LTI framework.

ACKNOWLEDGMENT

This work was supported by Science Foundation Ireland grants 13/RC/2094 P2 and 16/RC/3918.

REFERENCES

- Francis Akowuah and Fanxin Kong. "Real-time adaptive sensor attack detection in autonomous cyberphysical systems". In: 2021 IEEE 27th Real-Time and Embedded Technology and Applications Symposium (RTAS). IEEE. 2021, pp. 237–250.
- [2] Daniel Althoff, Matthias Althoff, and Sebastian Scherer. "Online safety verification of trajectories for unmanned flight with offline computed robust invariant sets". In: 2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). IEEE, 2015, pp. 3470–3477. ISBN: 1479999946.

- [3] Mazen Azzam et al. "Grounds for Suspicion: Physicsbased Early Warnings for Stealthy Attacks on Industrial Control Systems". In: (2021). arXiv: 2106. 07980 [cs.CR].
- [4] Cheng-Zong Bai, Fabio Pasqualetti, and Vijay Gupta. "Security in stochastic control systems: Fundamental limitations and performance bounds". In: *American Control Conference (ACC), 2015.* IEEE, 2015, pp. 195–200.
- [5] Andreas Bathelt, N Lawrence Ricker, and Mohieddine Jelali. "Revision of the tennessee eastman process model". In: *IFAC-PapersOnLine* 48.8 (2015), pp. 309– 314.
- [6] Stephen Boyd, Stephen P Boyd, and Lieven Vandenberghe. *Convex optimization*. Cambridge university press, 2004.
- [7] Andrea Carcano et al. "A multidimensional critical state analysis for detecting intrusions in SCADA systems". In: *IEEE Transactions on Industrial Informatics* 7.2 (2011), pp. 179–186.
- [8] John Henry Castellanos and Jianying Zhou. "A modular hybrid learning approach for black-box security testing of CPS". In: *International Conference on Applied Cryptography and Network Security*. Springer, 2019, pp. 196–216.
- [9] Xin Chen and Sriram Sankaranarayanan. "Model predictive real-time monitoring of linear systems". In: 2017 IEEE Real-Time Systems Symposium (RTSS). IEEE, 2017, pp. 297–306. ISBN: 1538614154.
- [10] Yi Chou, Hansol Yoon, and Sriram Sankaranarayanan. "Predictive Runtime Monitoring of Vehicle Models Using Bayesian Estimation and Reachability Analysis". In: Intl. Conference on Intelligent Robots and Systems (IROS). 2020.
- [11] Alessio Coletta. "Predictive Detection of Known Security Criticalities in Cyber Physical Systems with Unobservable Variables". In: 11th international conference on security and its applications (cnsa). 2018, pp. 61–77.
- [12] James J Downs and Ernest F Vogel. "A plant-wide industrial process control problem". In: *Computers and chemical engineering* 17.3 (1993), pp. 245–255. ISSN: 0098-1354.
- [13] Sriharsha Etigowni et al. "Crystal (ball) I Look at Physics and Predict Control Flow! Just-Ahead-Of-Time Controller Recovery". In: *Proceedings of the* 34th Annual Computer Security Applications Conference. 2018, pp. 553–565.
- [14] Yangyang Geng et al. "A survey of industrial control system testbeds". In: *IOP Conference Series: Materials Science and Engineering*. Vol. 569. IOP Publishing, 2019, p. 042030.
- [15] Jairo Giraldo et al. "A Survey of Physics-Based Attack Detection in Cyber-Physical Systems". In: ACM Computing Surveys (CSUR) 51.4 (2018), p. 76.
- [16] Paul Griffioen et al. "A Tutorial on Detecting Security Attacks on Cyber-Physical Systems". In: 2019 18th

European Control Conference (ECC). IEEE, 2019, pp. 979–984. ISBN: 3907144007.

- [17] Ziyang Guo et al. "Worst-case stealthy innovationbased linear attack on remote state estimation". In: *Automatica* 89 (2018), pp. 117–124.
- [18] Navid Hashemi, Carlos Murguia, and Justin Ruths. "A comparison of stealthy sensor attacks on control systems". In: 2018 Annual American Control Conference (ACC). IEEE, 2018, pp. 973–979.
- [19] Navid Hashemi and Justin Ruths. "Gain design via lmis to minimize the impact of stealthy attacks". In: 2020 American Control Conference (ACC). IEEE. 2020, pp. 1274–1279.
- [20] Haimin Hu et al. "Reach-SDP: Reachability Analysis of Closed-Loop Systems with Neural Network Controllers via Semidefinite Programming". In: arXiv preprint arXiv:2004.07876 (2020).
- [21] Taylor T Johnson et al. "Real-time reachability for verified simplex design". In: ACM Transactions on Embedded Computing Systems (TECS) 15.2 (2016), p. 26. ISSN: 1539-9087.
- [22] Marina Krotofil and Jason Larsen. "Rocking the pocket book: Hacking chemical plants". In: *DefCon Conference, DEFCON*. 2015.
- [23] Alexander B Kurzhanski and Pravin Varaiya. "Ellipsoidal techniques for reachability analysis". In: International Workshop on Hybrid Systems: Computation and Control. Springer. 2000, pp. 202–214.
- [24] Alex A Kurzhanskiy and Pravin Varaiya. "Ellipsoidal toolbox (ET)". In: Proceedings of the 45th IEEE Conference on Decision and Control. IEEE. 2006, pp. 1498–1503.
- [25] Cheolhyeon Kwon and Inseok Hwang. "Reachability analysis for safety assurance of cyber-physical systems against cyber attacks". In: *IEEE Transactions on Automatic Control* 63.7 (2018), pp. 2272–2279.
- [26] Cheolhyeon Kwon and Inseok Hwang. "Recursive reachable set computation for on-line safety assessment of the cyber-physical system against stealthy cyber attacks". In: 2016 54th Annual Allerton Conference on Communication, Control, and Computing (Allerton). IEEE, 2016, pp. 1123–1128. ISBN: 1509045503.
- [27] Cheolhyeon Kwon, Weiyi Liu, and Inseok Hwang. "Security analysis for cyber-physical systems against stealthy deception attacks". In: *American Control Conference (ACC)*, 2013. IEEE, 2013, pp. 3344–3349.
- [28] Truls Larsson et al. "Self-optimizing control of a large-scale plant: The Tennessee Eastman process". In: *Industrial and engineering chemistry research* 40.22 (2001), pp. 4889–4901. ISSN: 0888-5885.
- [29] Colas Le Guernic. "Reachability analysis of hybrid systems with linear continuous dynamics". PhD thesis. 2009.
- [30] Robert M Lee, Michael J Assante, and Tim Conway. "Analysis of the cyber attack on the ukrainian power grid: Defense use case". In: SANS ICS (2016).

- [31] Jezdimir Milošević et al. "Quantifying the impact of cyber-attack strategies for control systems equipped with an anomaly detector". In: 2018 European Control Conference (ECC). IEEE, 2018, pp. 331–337. ISBN: 3952426989.
- [32] Jezdimir Milošević et al. "Security measure allocation for industrial control systems: Exploiting systematic search techniques and submodularity". In: *International Journal of Robust and Nonlinear Control* (2018).
- [33] Yilin Mo and Bruno Sinopoli. "On the performance degradation of cyber-physical systems under stealthy integrity attacks". In: *IEEE Transactions on Automatic Control* 61.9 (2016), pp. 2618–2624.
- [34] Saeid Mokhatab and William A Poe. Handbook of natural gas transmission and processing. Gulf professional publishing, 2012. ISBN: 0123869757.
- [35] Carlos Murguia and Justin Ruths. "On model-based detectors for linear time-invariant stochastic systems under sensor attacks". In: *IET Control Theory and Applications* (2019).
- [36] Carlos Murguia and Justin Ruths. "On reachable sets of hidden cps sensor attacks". In: 2018 Annual American Control Conference (ACC). IEEE, 2018, pp. 178–184.
- [37] Carlos Murguia et al. "Security metrics of networked control systems under sensor attacks (extended preprint)". In: *arXiv preprint arXiv:1809.01808* (2018).
- [38] Fabio Pasqualetti, Florian Dorfler, and Francesco Bullo. "Control-theoretic methods for cyberphysical security: Geometric principles for optimal cross-layer resilient control systems". In: *IEEE Control Systems* 35.1 (2015), pp. 110–127.
- [39] N Lawrence Ricker. "Decentralized control of the Tennessee Eastman challenge process". In: *Journal of Process Control* 6.4 (1996), pp. 205–221. ISSN: 0959-1524.
- [40] N. Lawrence Ricker. "Model predictive control of a continuous, nonlinear, two-phase reactor". In: *Journal* of Process Control 3.2 (1993), pp. 109–123.
- [41] Roy S Smith. "Covert misappropriation of networked control systems: Presenting a feedback structure". In: *IEEE Control Systems Magazine* 35.1 (2015), pp. 82– 92. ISSN: 1066-033X.
- [42] André Teixeira et al. "A secure control framework for resource-limited adversaries". In: *Automatica* 51 (2015), pp. 135–148.
- [43] Hoang-Dung Tran et al. "Decentralized real-time safety verification for distributed cyber-physical systems". In: *International Conference on Formal Techniques for Distributed Objects, Components, and Systems*. Springer, 2019, pp. 261–277.
- [44] David Umsonst and Henrik Sandberg. "Anomaly Detector Metrics for Sensor Data Attacks in Control Systems". In: 2018 Annual American Control Conference (ACC). IEEE, 2018, pp. 153–158.

- [45] David I. Urbina et al. "Limiting the impact of stealthy attacks on industrial control systems". In: *Proceedings* of the 2016 ACM SIGSAC Conference on Computer and Communications Security. ACM, 2016, pp. 1092– 1105.
- [46] Xiao-Guang Zhang, Guang-Hong Yang, and Saud Wasly. "Man-in-the-middle attack against cyberphysical systems under random access protocol". In: *Information Sciences* 576 (2021), pp. 708–724.
- [47] Xiao-Guang Zhang and Guanghong Yang. "Kullback-Leibler Divergence-based Attacks against Remote State Estimation in Cyber-physical Systems". In: *IEEE Transactions on Industrial Electronics* (2021).