Presented as late breaking results contribution
21st IFAC World Congress (Virtual)
Berlin, Germany, July 12-17, 2020

Control System Cyberattack Detection using Lyapunov-Based Economic Model Predictive Control ⋆

Henrique Oyama * **Helen Durand** **

* *Wayne State University, Detroit, MI 48202 USA (e-mail: helen.durand@wayne.edu).*
** *Wayne State University, Detroit, MI 48202 USA (e-mail: helen.durand@wayne.edu).*

Abstract: Cybersecurity of control systems is a topic of increasing concern for chemical processes. In this work, we develop two techniques for detecting cyberattacks involving false state measurements being provided to a specific control formulation known as Lyapunov-based economic model predictive control (LEMPC) that take advantage of the closed-loop stability properties of the control formulation to seek to detect attacks after they occur. The first approach utilizes an integrated detection, control, and state estimation framework to flag deviations of the state estimates from "normal" process behavior as problematic cyberattacks, and the second control framework uses randomized modifications to an LEMPC formulation online, with reference to a baseline LEMPC design, to potentially detect cyberattacks.

Keywords: control system cybersecurity, model predictive control, chemical process control, nonlinear systems, state estimation.

1. INTRODUCTION

Cybersecurity is receiving increasing attention in the chemical process control literature (e.g., Cárdenas et al. (2011)). One of the challenges for cybersecurity of chemical processes is that such processes are often described by nonlinear dynamic models. While a number of results exist on securing linear systems via state estimation, detection, and/or control strategies (e.g., Fawzi et al. (2014)), there are significant gaps with respect to our ability to secure nonlinear systems under cyberattacks. Our recent work has focused on defining cyberattacks on nonlinear systems in a dynamic systems context Durand (2018). In the present manuscript, we explore two false sensor measurement attack detection techniques that are based on the closed-loop stability properties of a control design known as Lyapunov-based economic model predictive control (LEMPC) Heidarinejad et al. (2012). The first utilizes auxiliary state estimators, combined with control, to develop an attack detection methodology under the assumption that a single state estimator can be compromised by the attack but that auxiliary estimates are available for attack detection purposes. The second is based on the random generation of new control laws on-line with guaranteed stability properties for which the effect of the controller (in the sense of a guarantee on a decrease in the Lyapunov function over a sampling period) is characterizable in the absence of an attack, allowing failures of the expected effect to signal a potential cyberattack.

⋆ Financial support from the National Science Foundation CBET-1839675 and CNS-1932026 and Wayne State University is gratefully acknowledged.

2. PRELIMINARIES

2.1 Notation

The notation $|\cdot|$ signifies the Euclidean norm of a vector. $\alpha : [0, a) \to [0, \infty)$ is a class $\mathcal{K}$ function if $\alpha(0) = 0$ and the function is continuous and strictly increasing. Ω_ρ denotes a level set of a scalar-valued function V (i.e., $\Omega_\rho := \{x \in R^n : V(x) \leq \rho\}$). Set subtraction is signified by $'/'$ (i.e., $A/B := \{x \in R^n : x \in A,\ x \notin B\}$). x^T is the transpose of the vector x. A sampling time is denoted by $t_k := k\Delta$, $k = 0, 1, \ldots$, where Δ is a sampling period. A function $f(x)$ is locally Lipschitz in x with Lipschitz constant L_f if $|f(x') - f(x'')| \leq L_f|x' - x''|$ for all x' and x'' in a set.

2.2 Class of Systems

This work considers the following class of systems:

$$\dot{x}(t) = f(x(t)) + g(x(t))u(t) + l(x(t))w(t) \qquad (1)$$

where $x \in X \subset R^n$, $u \in U \subset R^m$, and $w \in W \subset R^z$ are the state, input, and disturbance vectors, respectively, and f, g and l are sufficiently smooth vector or matrix-valued functions with $f(0) = g(0) = 0$. We define $W := \{w \in R^z \mid |w| \leq \theta_w,\ \theta_w > 0\}$ and $U := \{u \in R^m \mid |u| \leq u^{\max}\}$. We consider that the "nominal" system of Eq. 1 ($w \equiv 0$) is stabilizable such that there exists an asymptotically stabilizing feedback control law $h(x)$, a sufficiently smooth Lyapunov function V, and class $\mathcal{K}$ functions $\alpha_i(\cdot)$, $i = 1, 2, 3, 4$, where:

$$\alpha_1(|x|) \leq V(x) \leq \alpha_2(|x|) \qquad (2)$$

$$\frac{\partial V(x)}{\partial x}\left(f(x(t)) + g(x(t))h(x)\right) \leq -\alpha_3(|x|)) \qquad (3)$$

Presented as late breaking results contribution
21st IFAC World Congress (Virtual)
Berlin, Germany, July 12-17, 2020

$$\left|\frac{\partial V(x)}{\partial x}\right| \leq \alpha_4(|x|) \tag{4}$$

$$h(x) \in U \tag{5}$$

$\forall\, x \in D \subset R^n$ (D is an open neighborhood of the origin). A level set of V contained within D and X is denoted by Ω_ρ and is termed the stability region. We assume that there are M sets of measurements $y_i \in R^{q_i}$, $i = 1, \ldots, M$, available continuously as follows:

$$y_i(t) = k_i(x(t)) + v_i(t) \tag{6}$$

where k_i is a vector-valued function, and v_i represents the measurement noise associated with the measurements y_i. We assume that the measurement noise is bounded (i.e., $v_i \in V_i := \{v_i \in R^{q_i} \mid |v_i| \leq \theta_{v,i},\ \theta_{v,i} > 0\}$). It is considered that for each of the M sets of measurements, a deterministic observer exists defined as follows:

$$\dot{z}_i = F_i(\epsilon_i, z_i, y_i) \tag{7}$$

where z_i is the estimate of the process state from the i-th observer, $i = 1, \ldots, M$, F_i is a vector-valued function, and $\epsilon_i > 0$. When a controller $h(z_i)$ with Eq. 7 is used to control the closed-loop system of Eq. 1, we make the following assumptions.

Assumption 1. Ellis et al. (2014); Lao et al. (2015) There exist positive constants θ_w^*, $\theta_{v,i}^*$, such that for each pair $\{\theta_w, \theta_{v,i}\}$ with $\theta_w \leq \theta_w^*$, $\theta_{v,i} \leq \theta_{v,i}^*$, there exist $0 < \rho_{1,i} < \rho$, $e_{m0i} > 0$ and $\epsilon_{Li}^* > 0$, $\epsilon_{Ui}^* > 0$ such that if $x(0) \in \Omega_{\rho_{1,i}}$, $|z_i(0) - x(0)| \leq e_{m0i}$ and $\epsilon_i \in (\epsilon_{Li}^*, \epsilon_{Ui}^*)$, the trajectories of the closed-loop system are bounded in Ω_ρ, $\forall t \geq 0$.

Assumption 2. Ellis et al. (2014); Lao et al. (2015) There exists $e_{mi}^* > 0$ such that for each $e_{mi} \geq e_{mi}^*$, there exist $t_{bi}(\epsilon_i)$ such that $|z_i(t) - x(t)| \leq e_{mi}$, $\forall\ t \geq t_{bi}(\epsilon_i)$.

3. CYBERATTACK-RESILIENT OUTPUT FEEDBACK LEMPC

This section develops an implementation strategy for a combined detection, control, and state estimation framework for cyberattack detection (in the case of state measurement falsification cyberattacks) for nonlinear systems under an LEMPC receiving state estimates (rather than full state feedback) when it is assumed that only the $i = 1$ state estimate is impacted by a cyberattack (i.e., the other z_j, $j = 2, \ldots, M$, are not impacted by the faulty state measurements). The results are initial theoretical advances toward characterizing a state estimation-based cyberattack framework for nonlinear systems under LEMPC; future work will extend these results to consider more than one state estimator impacted by a cyberattack or that the attack is not on the estimator used by the LEMPC, and will explore the concepts via simulation.

3.1 Output Feedback-Based LEMPC

This section uses LEMPC Heidarinejad et al. (2012) as part of a cyberattack detection strategy. LEMPC is formulated as follows Ellis et al. (2014); Lao et al. (2015):

$$\min_{u(t)\in S(\Delta)} \int_{t_k}^{t_{k+N}} L_e(\tilde{x}(\tau), u(\tau))\, d\tau \tag{8a}$$

$$\text{s.t.}\quad \dot{\tilde{x}}(t) = f(\tilde{x}(t)) + g(\tilde{x}(t))u(t) \tag{8b}$$

$$\tilde{x}(t_k) = z_i(t_k) \tag{8c}$$

$$\tilde{x}(t) \in X,\ \forall\, t \in [t_k, t_{k+N}) \tag{8d}$$

$$u(t) \in U,\ \forall\, t \in [t_k, t_{k+N}) \tag{8e}$$

$$V(\tilde{x}(t)) \leq \rho_{e,i},\ \ \forall\, t \in [t_k, t_{k+N}),\ \text{if } \tilde{x}(t_k) \in \Omega_{\rho_{e,i}} \tag{8f}$$

$$\frac{\partial V(\tilde{x}(t_k))}{\partial x}(g(\tilde{x}(t_k))u(t_k)) \leq \frac{\partial V(\tilde{x}(t_k))}{\partial x}(g(\tilde{x}(t_k))h(\tilde{x}(t_k))) \tag{8g}$$

$$\text{if } \tilde{x}(t_k) \in \Omega_\rho / \Omega_{\rho_{e,i}} \tag{8h}$$

In Eq. 8, the stage cost L_e is optimized (Eq. 8a) subject to the state predictions coming from the dynamic model (Eq. 8b) when the state measurement from the i-th estimator is used at t_k (Eq. 8c). The notation $u(t) \in S(\Delta)$ signifies that $u(t)$ (bounded by Eq. 8e) is a piecewise-constant input vector with N pieces (N is the prediction horizon) to be held for Δ. Eq. 8d represents a state constraint, and Eqs. 8f-8h are Lyapunov-based stability constraints that enforce boundedness of the closed-loop state within Ω_ρ. $\Omega_{\rho_{e,i}}$ is a subset of Ω_ρ. Cyberattacks considered are false measurements being provided to the state estimators used for Eq. 8b, where the false estimates are assumed to lie within Ω_ρ to avoid detection on the basis of the state being outside of the region in which it should be maintained when the controller is properly functioning, but are not required to take any specific trajectory.

3.2 Guaranteed Detection Strategy

In this section, we define a detection strategy which guarantees that any cyberattacks on the ($i = 1$) state estimator used in designing the LEMPC of Eq. 8 which would drive the closed-loop state out of Ω_ρ will be detected before this occurs. It recognizes cyberattacks by flagging deviations of the state estimates from "normal" behavior; however, as "normal" behavior includes both measurement noise and disturbances (Eqs. 1 and 6), care must be taken in setting the threshold on the state estimate deviation from a "normal" value to avoid false detections. To determine a threshold, we note that the bounds in Assumption 2 imply that the following holds:

$$\begin{aligned} |z_i(t) - z_j(t)| &= |z_i(t) - x(t) + x(t) - z_j(t)| \\ &\leq |z_i(t) - x(t)| + |z_j(t) - x(t)| \\ &\leq \epsilon_{ij} := (e_{mi}^* + e_{mj}^*) \leq \epsilon_{\max} := \max\{\epsilon_{ij}\} \end{aligned} \tag{9}$$

for all $i \neq j$, $i = 1, \ldots, M$, $j = 1, \ldots, M$, as long as $t \geq t_q = \max\{t_{b1}, \ldots, t_{bM}\}$. Therefore, abnormal behavior can be detected if $|z_i(t_k) - z_j(t_k)| > \epsilon_{\max}$ if $t_k > t_q$. In practice, it would not be possible to know the numbers e_{mi}^* and e_{mj}^*, as they can only be known by knowing an upper bound on how far off each $z_i(t)$ is from $x(t)$, which cannot be known since full state feedback may not be available. In the following, we will assume that an upper bound $\epsilon_{\max}$ can be estimated.

Presented as late breaking results contribution
21st IFAC World Congress (Virtual)
Berlin, Germany, July 12-17, 2020

3.3 Implementation Strategy

This implementation strategy assumes that the process has already been run successfully in the absence of attacks under the LEMPC of Eq. 8 for some time such that $|z_j(t) - x(t)| \leq \epsilon^*_{mj}$ for all $j = 1, \ldots, M$ before an attack:

(1) At sampling time t_k, if $|z_1(t_k) - z_j(t_k)| > \epsilon_{\max}$ for $j = 2, \ldots, M$, or $z_1(t_k) \notin \Omega_\rho$, detect that a cyberattack is occurring and go to Step 1a. Else, go to Step 1b.
   (a) Enter an emergency shut-down mode (e.g., pre-specified control actions like cutting feeds) that no longer operates the process under the LEMPC of Eq. 8 with $i = 1$ but instead under an emergency shut-down sequence.
   (b) Operate the process under the LEMPC of Eq. 8. $t_k \leftarrow t_{k+1}$. Go to Step 1.

3.4 Stability and Feasibility Analysis

The proposed detection strategy avoids false detections and will only detect an attack if a non-ordinary state estimate occurs. Next, we address whether there are any attacks that would be within the detection limits that could cause the closed-loop state to leave Ω_ρ before the attack is detected at t_k. We seek therefore to develop the conditions defining Ω_ρ that allow the closed-loop state to be maintained within Ω_ρ under the proposed strategy even if an undetected attack impacting the $i = 1$ estimate only occurs at t_k. To analyze this, we first define the worst-case deviation of $z_1(t_k)$ from $x(t_k)$ under the proposed detection policy in the following proposition.

Proposition 1. Under the implementation strategy of Section 3.3 and the assumption that multiple state estimates are available with a cyberattack impacting only the $i = 1$ estimate, the worst-case difference between $z_1(t_k)$ and $x(t_k)$, for all $j = 2, \ldots, M$, when no attack is detected at t_k is given by:

$$|z_1(t_k) - x(t_k)| \leq \epsilon^*_{1j} := \epsilon_{\max} + e^*_{mj} \leq \epsilon^* := \max \epsilon^*_{1j} \quad (10)$$

Proof. The bound of Eq. 10 is derived from the following:

$$\begin{aligned} |z_1(t_k) - x(t_k)| &= |z_1(t_k) - z_j(t_k) + z_j(t_k) - x(t_k)| \\ &\leq |z_1(t_k) - z_j(t_k)| + |z_j(t_k) - x(t_k)| \leq \epsilon_{\max} + e^*_{mj} \end{aligned} \quad (11)$$

where the last inequality follows from the fact that the detection algorithm was not activated (such that therefore $|z_1(t_k) - z_j(t_k)| < \epsilon_{\max}$ and the assumption that only the $i = 1$ state estimator used by the LEMPC is affected by the false state measurements (i.e., the other state measurements with $j \neq 1$ continue to have $|z_j(t_k) - x(t_k)| \leq e^*_{mj}$, according to Assumption 2).

We now introduce a theorem that re-purposes a bound on the allowable error in an estimate supplied to output feedback LEMPC which can be tolerated without loss of closed-loop stability (derived in Lao et al. (2015); Ellis et al. (2014) for the case that the error between the state estimate and actual state is due to a combination of disturbances and measurement noise) to the cyberattack problem. Specifically, the proposed detection method allows the bound in Eq. 10 to be developed, allowing cyberattacks to be treated in the framework previously analyzed in Lao et al. (2015); Ellis et al. (2014) for guaranteeing closed-loop stability of output feedback LEMPC in the presence of measurement noise and disturbances, and thereby allowing the combined detection-control framework to guarantee closed-loop stability when a cyberattack is not flagged according to the proposed methodology.

Theorem 1. Consider the system of Eq. 1 in closed-loop under the LEMPC of Eq. 8 based on an observer and controller pair satisfying Assumptions 1-2 and formulated with respect to the $i = 1$ measurement vector, and formulated with respect to a controller $h(\cdot)$ that meets Eqs. 2-5. Let $\theta_w \leq \theta^*_w$, $\theta_{v,i} \leq \theta^*_{v,i}$, $\epsilon_i \in (\epsilon^*_{Li}, \epsilon^*_{Ui})$, and $|z_i(0) - x(0)| \leq e_{m0i}$, for $i = 1, \ldots, M$. Also, let $\epsilon_{W,1} > 0$, $\Delta > 0$, and $\rho > \rho_{1,1} > \rho_{e,1} > \rho_{\min,1} > \rho_{s,1} > 0$, satisfy:

$$\begin{aligned} \rho_{e,1} \leq \rho - \max\{&f_V(f_W(\epsilon^*, \Delta)) \\ &+ f_V(\epsilon^*), M \max\{t_{z1}, \Delta\}\alpha_4(\alpha_1^{-1}(\rho))\} \end{aligned} \quad (12)$$

$$\begin{aligned} -\alpha_3(\alpha_2^{-1}(\rho_{s,1})) + (L^f_V + L^g_V u^{\max})(M\Delta + \epsilon^*) \\ + M^l_V \theta_w \leq -\epsilon_{W,1}/\Delta \end{aligned} \quad (13)$$

$$\rho_{\min,1} = \max\{V(x(t + \Delta)) | V(x(t)) \leq \rho_{s,1}\} \quad (14)$$

where L^f_V and L^g_V are Lipschitz constants for $\frac{\partial V}{\partial x} f$ and $\frac{\partial V}{\partial x} g$, $|\dot{x}| \leq M$, M^l_V bounds $|\frac{\partial V}{\partial x} l|$ for $x \in \Omega_\rho$, t_{z1} is the first sampling time after t_{b1},

$$\begin{aligned} f_W(s, \tau) := (s + \frac{M_l \theta_w}{L_f + L_g u^{\max}}) e^{(L_f + L_g u^{\max})\tau} \\ - \frac{M_l \theta_w}{L_f + L_g u^{\max}} \end{aligned} \quad (15)$$

where L_f and L_g are positive Lipschitz constants for the functions f and g as defined in Section 2.1, and $M_l > 0$ satisfies $|l(x)| \leq M_l$, $\forall\ x \in \Omega_\rho$, and

$$f_V(s) := \alpha_4(\alpha_1^{-1}(\rho))s + M_v s^2 \quad (16)$$

where M_v is a positive constant. Then, if $x(0) \in \Omega_{\rho_{e,1}}$, $x(t) \in \Omega_\rho$, $\forall\ t \geq 0$ until a cyberattack is detected if the attack occurs after t_q.

4. RANDOMIZED LEMPC CHANGES TO PROBE FOR CYBERATTACKS

An alternative to probing for cyberattacks is an LEMPC design that incorporates full state feedback but allows the steady-state of operation to be adjusted. With slight abuse of notation, in this section, we refer to the LEMPC design around the operating steady-state as the baseline or 1-LEMPC, which has stability region Ω_{ρ_1}, stability region subset $\Omega_{\rho_{e,1}}$, Lyapunov function V_1, and controller h_1 used in its design. The model of Eq. 1 with origin at the operating steady-state will be defined with a subscript 1 in the following. We will consider brief periods of use of LEMPC's formulated around other steady-states within Ω_{ρ_1}. These will be referred to as j-LEMPC designs (for $j > 1$) with stability region Ω_{ρ_j}, stability region subset $\Omega_{\rho_{e,j}}$, Lyapunov function V_j, and controller h_j used in design. The model of Eq. 1 rewritten in deviation variable form from the j-th steady-state will utilize a subscript j.

The proposed strategy uses random generation of steady-states within $\Omega_{\rho_{e,1}}$ of the (baseline) 1-LEMPC that have steady-state inputs within U to develop new j-LEMPC ($j > 1$) designs online which can drive the closed-loop state toward the new steady-state in the absence of a cyberattack. Specifically, the LEMPC of Eq. 8 with full state feedback is used until a random time $t_{s,j}$, $j = 2, 3 \ldots$,

when $x(t_k) \in \Omega_{\rho_e,1}$, at which it is desired to run a check to determine whether a cyberattack is occurring. At this random time, a steady-state is selected that has a stability region around it, contained within Ω_{ρ_1}, that includes $x(t_k)$. Then, at $t_{s,j}$, an LEMPC of the form of Eq. 8 with state feedback, but formulated with respect to this new steady-state and with Eq. 8h activated regardless of the position of the initial state, is selected to control the system for the next sampling period. This should cause the value of the Lyapunov function associated with this new LEMPC to decrease over the subsequent sampling period if controller and system parameters, such as the sampling period and disturbance bound, are sufficiently small and the closed-loop state is not in a neighborhood of the new steady-state Heidarinejad et al. (2012); if it does not, a cyberattack may be occurring. There is no guarantee that this method detects an attack (it is possible that under an attack, the Lyapunov function around the random steady-state could also decrease under a false state measurement trajectory provided by an attacker). The guarantee that can be provided is that if the Lyapunov function for the random steady-state does not decrease over a sampling period following the use of the modified LEMPC, abnormal behavior is occurring (which may correspond to a cyberattack). It is noted that this method does not ensure that the closed-loop state will not exit Ω_{ρ_1} under the cyberattack in the sampling period following the probing (if that occurs, this would not be helpful).

5. APPLICATION TO A CHEMICAL PROCESS EXAMPLE

The switching LEMPC cyberattack detection method is illustrated through a simulation of a continuous stirred tank reactor (CSTR) with the dynamic model and process parameters in Alanqar et al. (2015). The states are the reactant concentration and temperature in the reactor (C_A and T, respectively). The manipulated inputs are C_{A0} (the reactant feed concentration) and the rate of heat input Q. The vectors of deviation variables for the states and inputs from their steady-state values, $C_{As} = 1.22$ kmol/m^3, $T_s = 438.2$ K, $C_{A0s} = 4.0$ kmol/m^3, and $Q_s = 0$ kJ/h, respectively, are $x = [x_1\ x_2]^T = [C_A - C_{As}\ T - T_s]^T$ and $u = [u_1\ u_2]^T = [C_{A0} - C_{A0s}\ Q - Q_s]^T$. The objective function $L_e = k_0 e^{-E/(RT)} C_A^2$ was used. Lyapunov-based stability constraints of the form in Eq. 8 were designed using a Lyapunov function $V_1 = x^T P x$, where $P = [1200\ 5; 5\ 0.1]$. The Lyapunov-based controller utilized $h_1(x) = 0$ kmol/m^3 for simplicity and designed $h_2(x)$ via Sontag's control law Lin and Sontag (1991). The stability region was set to $\rho_1 = 300$ (i.e., $\Omega_{\rho_1} = \{x \in R^2 : V_1(x) \leq \rho_1\}$) and $\rho_{e,1} = 225$.

The process state was initialized off steady-state at $x_{init} = [-0.21$ kmol/m^3 28.89 K$]^T$, and for the purpose of illustrating the proposed method we consider that both the cyberattack probing method and the cyberattack (which consists of the false state measurement $x_1 = 0.1$ kmol/m^3, $x_2 = 21.75$ K being provided to the LEMPC at each sampling time) were initiated at t_0. The cyberattack probing mechanism required the creation of a new steady-state (selected to be $x_d = [0$ kmol/m^3 11.75 K$]^T$) that has a stability region in $\Omega_{\rho_{e,1}}$. It was defined using $V_2(x) = (x - x_d)^T P_2 (x - x_d)$ with $P_2 = [2100\ 10; 10\ 0.25]$, $\rho_2 = 100$ (i.e., $\Omega_{\rho_2} = \{(x - x_d) \in R^2 : V_2(x - x_d) \leq \rho_2\}$) and $\rho_{e,2} = 75$, and was selected to include x_{init}. When no attack is performed at t_0, the detection mechanism causes V_2 to decrease after t_0. However, in the presence of the attack, though the closed-loop state continues to evolve (Fig. 1), V_2 remains fixed at a value not equal to 0. This indicates abnormal behavior and here is indicative of a cyberattack which, if no action is taken to combat it, drives the closed-loop state out of the stability region.

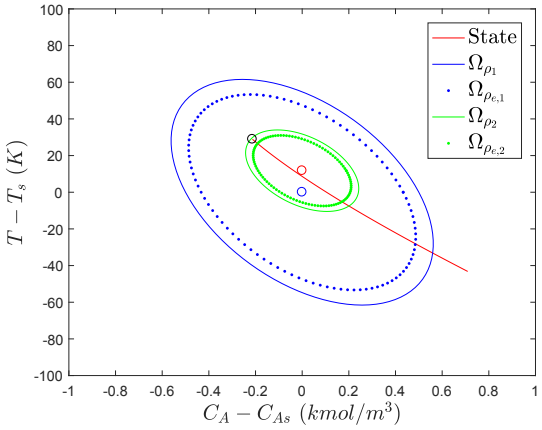

Fig. 1. State-space trajectories for the CSTR process under the switching LEMPC cyberattack detection methodology with an attack for 0.08 h.

REFERENCES

Alanqar, A., Ellis, M., and Christofides, P.D. (2015). Economic model predictive control of nonlinear process systems using empirical models. *AIChE Journal*, 61, 816–830.

Cárdenas, A.A., Amin, S., Lin, Z.S., Huang, Y.L., Huang, C.Y., and Sastry, S. (2011). Attacks against process control systems: Risk assessment, detection, and response. In *Proceedings of the ACM Asia Conference on Computer & Communications Security*. Hong Kong, China.

Durand, H. (2018). A nonlinear systems framework for cyberattack prevention for chemical process control systems. *Mathematics*, 6, 44 pages.

Ellis, M., Zhang, J., Liu, J., and Christofides, P.D. (2014). Robust moving horizon estimation based output feedback economic model predictive control. *Systems & Control Letters*, 68, 101–109.

Fawzi, H., Tabuada, P., and Diggavi, S. (2014). Secure estimation and control for cyber-physical systems under adversarial attacks. *IEEE Transactions on Automatic Control*, 59, 1454–1467.

Heidarinejad, M., Liu, J., and Christofides, P.D. (2012). Economic model predictive control of nonlinear process systems using Lyapunov techniques. *AIChE Journal*, 58, 855–870.

Lao, L., Ellis, M., Durand, H., and Christofides, P.D. (2015). Real-time preventive sensor maintenance using robust moving horizon estimation and economic model predictive control. *AIChE Journal*, 61, 3374–3389.

Lin, Y. and Sontag, E.D. (1991). A universal formula for stabilization with bounded controls. *Systems & Control Letters*, 16, 393–397.