Resilience of Process Control Systems to Cyber-Physical Attacks

Marina Krotofil[1] and Alvaro A. Cárdenas[2]

[1] Hamburg University of Technology, Hamburg, Germany
marina.krotofil@tuhh.de
[2] University of Texas at Dallas, Richardson, TX 75080, USA
alvaro.cardenas@utdallas.edu

Abstract. In this work we investigate the matter of "secure control" – a novel research direction capturing security objectives specific to Industrial Control Systems (ICS). We provide an empirical analysis of the well known Tennessee Eastman process control challenge problem to gain insights into the behavior of a physical process when confronted with cyber-physical attacks. In particular, we investigate the impact of integrity and DoS attacks on sensors which measure physical phenomena. We also demonstrate how the results of process-aware security analysis can be applied to improve process resilience to cyber-physical attacks.

Keywords: Cyber-physical attacks, Tennessee-Eastman process, simulations, secure control.

1 Introduction

Advances in computing and networking have added new capabilities to physical systems that could not be feasibly added before. This has led to the emergence of engineered systems called cyber-physical systems: systems where the physical world is measured and controlled thanks to modern advances in computation and control. Aircrafts, robots, utilities, chemical and food plants and even modern smartphones are the examples of such systems. In this paper our focus is on process control systems (PCS).

Modernization of control systems has been motivated by plant operators' demands for better performance, easier maintenance, and more uptime. What used to be a panel of relays is now an embedded computer, and what used to be a simple analog sensor is now a smart transmitter [17] with multiple wired and wireless communication modes, self-diagnostic capabilities, and even a web-server with an interactive GUI for device configuration and troubleshooting. While security engineers try to limit the numbers of access points, helpful vendors are giving more options on how to access sensors inserted into physical processes.

While this modernization is necessary for improving the efficiency of a process, over the past decade many concerns have been raised about the vulnerabilities in industrial control systems to both random cyber failures and security attacks.

H. Riis Nielson and D. Gollmann (Eds.): NordSec 2013, LNCS 8208, pp. 166–182, 2013.

The primary focus of academia and industry has been on securing the communication infrastructure and hardening of control systems. There is a large body of literature on how to adapt existing IT-security methods to the characteristic features of the control domain. However modern malware for persistent attacks may now be equipped with "trusted" certificates, travel in USB sticks and laptops of "trusted" parties, carrying zero-days exploits, rootkits and propagate through "trusted" security updates. It is becoming increasingly difficult to prevent, detect and to halt these attacks based solely on technical measures deployed in the cyber-layer.

To address the limitations of defending a system using only IT methods, a new line of research has focused on understanding the adversary's interactions with the physical system. Analyzing the effects of attacks in the process control domain is a growing area of research. Some experimental works [11], [5] were conducted on the basis of a simplified model of the Tennessee Eastman (TE) process [19]. In the closest work to ours, Yu-Lung Huang *et al.* [11] proposed models of cyber attacks in control systems and evaluated physical and economical consequences of proposed attacks. We extend their results by analyzing the *full* model (as opposed to the simplified one) of the TE process with the goal of analyzing more realistic, larger-scale PCS with multiple control loops and physical interdependencies. Another addition to work is scrutiny of timing parameters of the attacks. We also extended the analysis of integrity attacks to less aggressive modifications of sensor readings to slow down process response and to analyze process dynamic. Our simulation results do not coincide closely with their as the simplified model of TE process is only moderately non linear, whereas the full TE model is highly non linear. Moreover the models follow different control strategies. The impact of DoS attacks on network routers is investigated by Chabukswar *et al.* [5]: in this work few sensors and actuators at a time become inaccessible for the controller causing process changes from negligible to drastic. In our work we provide further insights into the impact of DoS attacks and their timing parameters. The effect of network parameters and specific properties of control systems in the example of a Boiling Water Power Plant is evaluated by Genge *et al.* [8]: they identify that speed of control valves and task scheduling play an important role in designing processes resilient to malicious actions.

2 Preliminaries

Addressing the challenges of securing an industrial process requires knowledge about how the process is actually being managed with the help of actuators and control laws, and an understanding of the security requirements specific to process control.

2.1 Process control Fundamentals

In the process industry *process* refers to the methods of changing or refining raw materials to create an end product. Process industries include (petro)chemical,

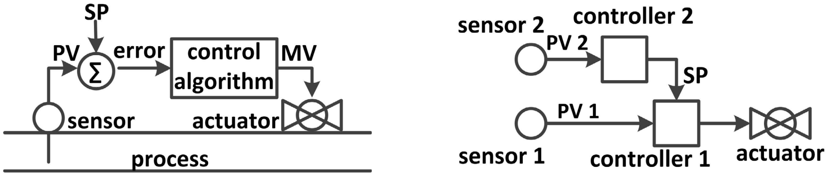

(a) Single feedback control loop

(b) Multi-variable control loop

Fig. 1. Types of control loops

food, water treatment, power and other industries. *Control* refers to the methods that are used to control process variables when manufacturing a product. This is done for three major reasons: (1) reducing variability; (2) increasing efficiency; (3) ensuring safety. The first two points are important for plant economy. Reduced variability lowers operational costs and ensures consistent quality of the end product. Efficiency refers to the accurate maintenance of optimal production conditions to decrease the production bill. Precise control is important for preventing runaway processes and ensuring safe operations.

The starting point in process engineering is deciding on a *setpoint* (SP) – the desired value of a certain process parameter, e.g. a tank level L. Level L is called *measured variable* and must be kept as close as possible to the setpoint by the means of control methods. Level L might be in fact determined indirectly via measuring two *process variables* (PV), in- and out-flows. If a level is measured directly, measured and process variable are the same. Process variables are processed by a controller containing a control algorithm based on a complex set of equations. The controller calculates the offset between SP and PV and outputs an actionable *manipulated value* (MV) to the actuator to bring the process closer to the SP. Such interactions form a basic feedback control loop as shown in Fig. 1(a). In practice, control loops can be complex. More common are multivariable or advanced control loops in which each MV depends on two or more of the measured variables (Fig. 1(b)). The strategies for holding a process at setpoint are not trivial, and the interactions of numerous setpoints in the overall process control plan can be subtle and complex. Process interactions may cause loop interactions via hidden feedback control loops. This makes controller tuning difficult and yields unstable loops.

2.2 Secure Control

The security goal in the traditional IT domain is the protection of information, be it data in storage or in transit. The security goal in the realm of industrial control systems is to protect the operations from intentional assaults so that, in the words of Ross Anderson, “the electricity continues to come out of the wall socket, regardless of the attempts of either Murphy or Satan to interrupt the supply” [2]. In the language of process control it means ensuring *process survivability* or if not possible – its *graceful degradation.*

While preserving device-to-device data integrity is of concern, ensuring that sensors faithfully capture process state (i.e., that the physical measurement is represented faithfully) is even more important. This security requirement is called *veracity* [10]. On one hand this property is crucial for state estimation algorithms, based on which manipulated values are computed. On the other hand, state estimation may help to identify implausible readings and by that improve *process resilience* to sensor manipulation attacks. In ICS with hard real-time requirements denial of service amounts to milliseconds. In some cases, process data are only valid for a short time and become irrelevant if arriving too late. The same applies to the scheduling of needed task actions. Therefore *timeliness* must be protected and "stale" data should be detected.

3 Approach

Conducting analysis of the dynamic behavior of a chemical process under cyber-attacks requires: (1) good knowledge of the process steady-state flow-sheet and its operating conditions; (2) thorough understanding of process control configuration; (3) defined attack models.

3.1 Process Modeling

The Tennessee Eastman (TE) challenge process [7] is a modified model of a real plant-wide industrial process. The process produces two liquid (*liq*) products from four gaseous (*g*) reactants involving two irreversible exothermic reactions:

$$\begin{aligned} &A(g) + C(g) + D(g) \rightarrow G(liq), &&\text{Product 1,}\\ &A(g) + C(g) + E(g) \rightarrow H(liq), &&\text{Product 2.} \end{aligned}$$

The process has five major operation units: the reactor, the product condenser, a vapor-liquid separator, a recycle compressor and a product stripper as shown in Fig. 2. The gaseous reactant and products are not specifically identified. Feed C is not pure and consists of 48.5% A and 51% C. The gas phase reactions are catalyzed by a substance dissolved in the liquid phase in the reactor. The reactor is pressurized and relies on an internal cooling system to remove the heat produced by the reactions. The products and the unreacted ingredients leave the reactor in the vapor phase, pass through a cooler that condenses the products, and from there to a vapor-liquid separator. Non-condensed components cycle back to the reactor feed via a compressor. Condensed components are sent to a stripping column that removes the remaining reactants. Products G and H exit the stripper base and are separated in a downstream refining section, which is not included in the problem statement. The byproducts and inerts are purged from the system in the vapor phase using a vapor-liquid separator.

The system may be operated in six distinctive modes which are determined by G/H mass ratios. Mode 1 is a base case with $G/H = 50/50$. The goal of plant operation is to maintain desired production rate and product composition within $\pm 5 mol\%$ while keeping other variables within specified operational limits. The process control goal is to minimize variability and absorb disturbances.

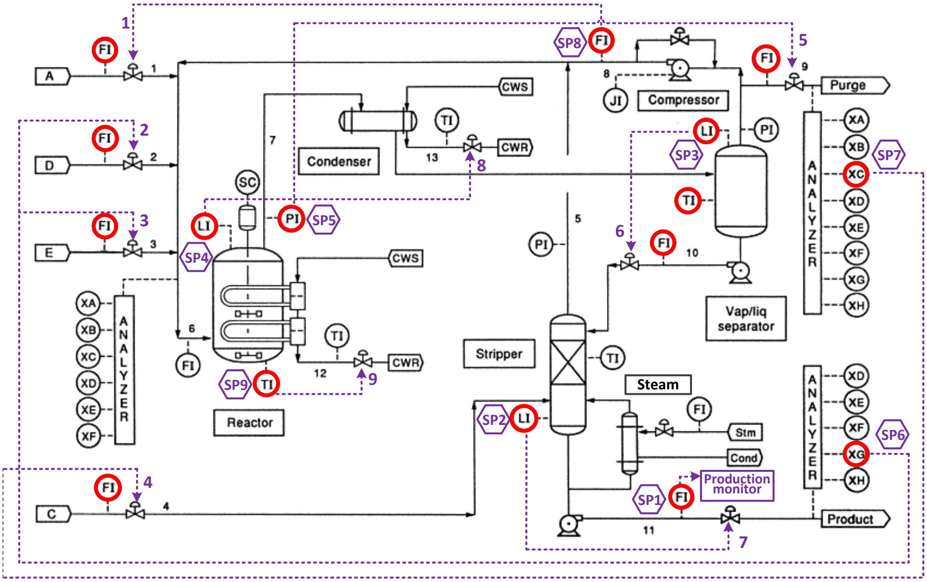

Fig. 2. Tennessee Eastman test problem under control based on [18]

If the process exceeds the safety limits, it automatically shuts down. The initial problem statement does not give recommendations on how the plant should be controlled. Instead the authors specify a process flow sheet, steady state material balance, operating conditions, possible types of plant disturbances and safety constraints. The control engineers are challenged to come up with their own control strategies, e.g. [15], [21], [14]. The original TE process has more degrees of freedom (valves) than necessary for control and the engineers are free to decide which ones to engage. The resulting control structures are usually designed to meet a specific control objective, e.g. optimal steady state, maximum rejection of process disturbances, ensuring on-demand production rate, adapting to an on-supply reactants rate or a combination of few. Optimization of the control strategy is subject to multiple constraints and an optimal solution is not always feasible.

For our empirical analysis we use the Matlab model of the TE process developed by Ricker [18]. It is implemented as a C-based MEX S-function with a Simulink model. The plant operates in mode 1 with a default simulation time of 72 hours and a sampling frequency of 100 measurement samples per hour. The model does not simulate start-up and shutdown procedures, instead its execution starts with the predefined base values. The plant has eleven valves for manipulation, and in total 41 measurements are involved in process monitoring. The proposed control configuration consists of 18 proportional-integral (PI) controllers, 16 process measurements XMEAS{1;2;3;4;5;7;8;9;10;11;12;14;15;17;31;40} and 9 setpoints which form 8 multivariable control loops and 1 single feedback control loop as specified in Table 1. The resulting control structure is depicted in Fig. 2. There are two auxiliary control loops for improved management of the

Table 1. Valves and measured variables

MV	Valve	Variable 1	Variable 2
XMV(1)	A-feed rate	Recycling rate	FI (stream 1)
XMV(2)	D-feed rate	%G in product	FI (stream 2)
XMV(3)	E-feed rate	%G in product	FI (stream 3)
XMV(4)	C-feed rate	%C in purge	FI (stream 4)
XMV(5)	Purge flow rate	Reactor pressure	FI (stream 9)
XMV(6)	Separator underflow	Separator level	FI (stream 10)
XMV(7)	Stripper underflow	Stripper level	FI (stream 11)
XMV(8)	Condenser cooler	Reactor level	TI (stream 13)
XMV(9)	Reactor cooler	Reactor temperature	—
XMV(10)	Compressor (recycle)	Not used	Not used
XMV(11)	Steam feed rate	Not used	Not used

Table 2. Process operating constraints [7]

	Normal operating limits		Shutdown limits	
Process variable	Low limit	High limit	Low limit	High limit
Reactor pressure	none	2895 kPa	none	3000 kPa
Reactor level	50% (11.8m^3)	100% (21.3m^3)	2.0m^3	24.0m^3
Reactor temperature	none	150∘C	none	175∘C
Product separator level	30% (3.3m^3)	100% (9.0m^3)	1.0m^3	12.0m^3
Stripper base level	30% (3.5m^3)	100% (6.6m^3)	1.0m^3	8.0m^3

production rate. The first generates a contributory control value which is used in the calculations of the XMV{1-7}. A second auxiliary loop is used to calculate additional control values for D- and E-feed rates depending on the G *mol*% in the product flow (stream 11).

All process measurements include Gaussian noise with standard deviation typical of the measurement type. The full notation and units of process characteristics can be found in [7]. From the attacker point of view, the most interesting process information are the operation constraints presented in Table 2. All 20 disturbances modes IDV{1-20} from the original problem statement are implemented in the model and can be included in the simulation selectively. Disturbances are an inevitable concern in plant operations. They enlarge variations in the process dynamics, complicating control and increasing operating costs. Modes IDV(6) and IDV(8) are the most difficult to handle. IDV(8) introduces random variations in feed composition of the reactor feed (stream 4). As can be seen in Fig. 3 it causes greater variability in reactor pressure P_{reac}. However, the process control scheme successfully rejects this type of disturbance so it does not affect the production goals. In contrast, the disturbance IDV(6) that shuts off the A feed cannot be absorbed and the process shuts down on high pressure in less than 8 hours. Such control deficiencies are usually compensated by the override controls (e.g. [21], [14]) which are not implemented in the model.

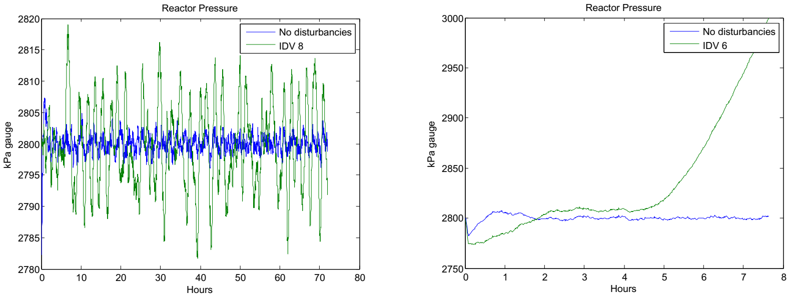

Fig. 3. Reactor pressure with and without disturbances

3.2 Attack Modeling

The adversary's goal is to cause tangible impact on the process, either on its safety or on its economy. In the physical domain, the attacker can either tamper with the sensor readings or modify the manipulated values issued by the controller. In this work we limit our study to the analysis of sensor compromise. We assume an adversary capable of either attacking sensors directly or being able to subvert communication channels and forge messages with respect to the protocol specification. Let $X_i(t)$ be a measurement of sensor i at time t, where $0 \leq t \leq T$, and T the duration of the simulation. The attack interval T_a is arbitrary and is limited to the simulation run time. In our setting, we simulate manipulated sensor readings X_i^a as follows:

$$X_i^a(t) = \begin{cases} X_i(t), & \text{for } t \notin T_a \\ X_i'(t), & \text{for } t \in T_a, \end{cases}$$

where $X_i'(t)$ is the modified reading.

Integrity attacks on process measurements involve forging sensor readings. Multiple strategies can be applied to falsify a sensor reading. We will investigate the case when the attacker claims measured physical phenomena as being too low or too high to deceive the controller and to evoke harmful compensating reaction. For example, claiming pressure in the reactor being too low will make the controller take correcting steps to increase the pressure, which with time can reach an unsafe boundary. To model this attack we first run the model without any attack to determine the span of PV for each sensor in the presence of the greatest disturbance, IDV(8). We then determine the boundary values for each variable. As $X_i^a(t)$ we use correspondingly:

$$X_i^{low}(t) = \min_{t \in T} X_i(t) \quad \text{and} \quad X_i^{high}(t) = \max_{t \in T} X_i(t).$$

The rationale behind using values which are not drastically too low or too high is to avoid rapid process shutdown due to exceeding of safety limits. This would not allow us to observe the process dynamic under attack. However, a sensitivity analysis of each control loop to the magnitude of the manipulation is required for a complete analysis. This includes scrutiny of loops under integrity attacks which consider the full sensing range of a variable as was done in [11].

During a *DoS attack* sensor signals do not reach the controller. If the attack starts at time t_a, we have:

$$X_i^a(t) = X_i(t_a - 1).$$

Translated into the real world scenario, the controller's input register assigned to storing measurements of a particular sensor will not be overwritten by a fresh value during the next control cycle run as would happen in a normal case.

4 Experimental Results

One of the original applications of the TE test process is *process diagnostic* [7]: testing and evaluation of process performance and reaction to new or unknown conditions. We analyze its resilience to cyber-physical assaults. Following the principle of "weakest link" we evaluate the impact of the attacks in the presence of IDV(8). There are three metrics readily available to evaluate the result of plant operations: *product quality* defined as *G mol%* in the product flow, *operating costs* and plant *shutdown time* (SDT) due to exceeding of safety constraints. Each simulation run generates in total 53 plots: 41 XMEAS, 11 XMV and operating costs. Moreover a real-time production monitor is available. We scrutinized the process for different types, times and durations of the attacks as well as for different magnitudes of sensor signal manipulations. Below we present some characteristic results to demonstrate how analyzing process reaction to intentional manipulations can be used to improve the robustness of PCS.

4.1 Integrity Attacks

Our analysis shows that the sensitivity of control loops to integrity attacks varies greatly. Attacks on certain sensors increase the variability of plant dynamic but do not endanger plant operations safety. Attacks on other control loops lead to shutdown with a SDT range from 20 minutes to more than 8 hours. The results of the simulations are summarized in Table 3. This table gives *only a notion* of control loop behavior under the attacks. A full evaluation would require a thorough individual analysis of each control loop under different types and modes of attack.

Impact on plant safety. Plant safety issues in general refer to two aspects. One is process safety itself, to prevent unwanted or uncontrolled chemical reactions. The other is equipment safety, which aims at preventing equipment failure or breakage. An example would be preventing pressure in the reactor exceeding safety limits to stave off reactor burst. There are 8 safety provisions implemented in the model with predefined thresholds as specified in Table 2.

A chemical reactor is typically the heart of an industrial process and will probably be a priority target for the adversary. A straightforward attack would

Table 3. Simulation results of integrity attacks

XMEAS	Sensor	Attack	Impact	SDT
(1)	A-feed rate	X^{low}	High reactor pressure	3 h
		X^{high}	High variability	–
(2)	D-feed rate	X^{low}	High stripper liquid level	4.5 h
		X^{high}	Low stripper liquid level	3.5 h
(3)	E-feed rate	X^{low}	High stripper liquid level	4.5 h
		X^{high}	Low stripper liquid level	2.5 h
(4)	C-feed rate	X^{low}	High reactor pressure	0.35 h
		X^{high}	High reactor pressure	0.9 h
(5)	Recycle flow	X^{low}	High reactor pressure	3.3 h
		X^{high}	High reactor pressure	6.5 h
(7)	Reactor pressure	X^{low}	High reactor pressure	8 h
		X^{high}	High operating costs	–
(8)	Reactor level	X^{low}	Low separator liquid level	1.5 h
		X^{high}	High stripper liquid level	1.2 h
(9)	Reactor temperature	X^{low}	High reactor pressure	1.8 h
		X^{high}	High reactor pressure	0.3 h
(10)	Purge rate	X^{low}	High operating costs	–
		X^{high}	High variability	–
(11)	Separator temperature	X^{low}	High variability	–
		X^{high}	High variability	–
(12)	Separator level	X^{low}	High separator liquid level	6 h
		X^{high}	Low separator liquid level	3.5 h
(14)	Separator underflow	X^{low}	Low separator liquid level	7 h
		X^{high}	High stripper liquid level	6.5 h
(15)	Stripper level	X^{low}	High stripper liquid level	6 h
		X^{high}	Low stripper liquid level	5 h
(17)	Stripper underflow	X^{low}	Low stripper liquid level	1.1 h
		X^{high}	High stripper liquid level	1.2 h
(31)	$\%C$ in purge	X^{low}	High stripper liquid level	1.5 h
		X^{high}	High reactor pressure	6 h
(41)	$\%G$ in product	X^{low}	D- and E-feed variability	–
		X^{high}	D- and E-feed variability	–

be forging pressure sensor reading as P^{low}_{reac} to provoke pressure rise. However, the response to this attack has slow dynamics and and it takes 8 hours to succeed. The attacker is also free to decide on the duration and frequency of her assault. Let the attacker launch her attack for 2 hours and wait for 8 hours in a cyclic fashion. We notice that the controller can recover the system state to the normal pressure level within 3 hours (Fig. 4(a)). We then investigate the impact of more frequent attacks on the pressure sensor (every hour). As can be seen in Fig. 4(b), such a strategy is not helpful in achieving an unsafe pressure rise. In contrast, the mean pressure level decreases. Although the illustrated timing attack strategies were not optimal from the attacker's point of view, timing parameters of the attack are an important dimension for process resilience analysis.

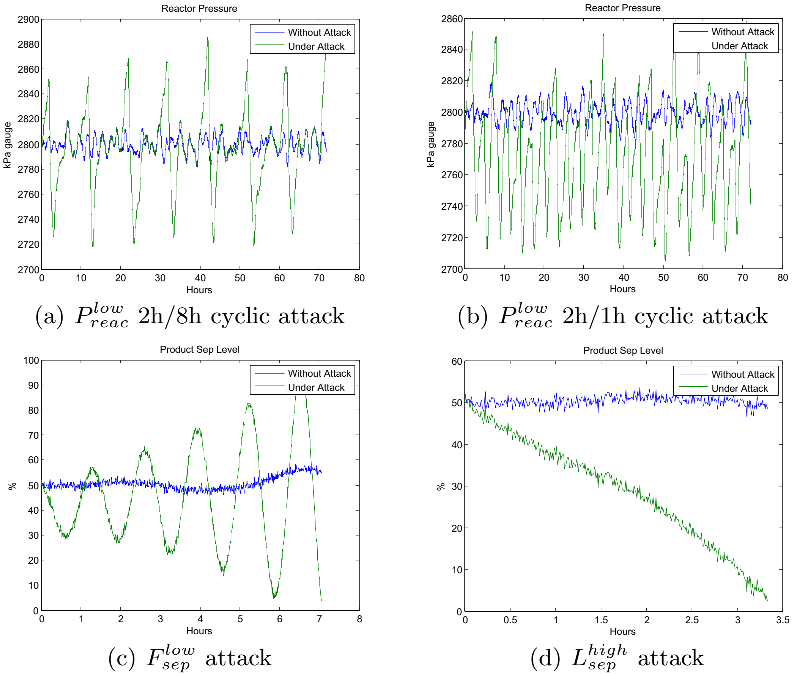

(a) P_{reac}^{low} 2h/8h cyclic attack

(b) P_{reac}^{low} 2h/1h cyclic attack

(c) F_{sep}^{low} attack

(d) L_{sep}^{high} attack

Fig. 4. Impact on integrity attack

To maximize the impact a more knowledgeable attacker might prefer to attack a temperature sensor because the rate r of the reaction depends on temperature T in an exponential fashion [14]:

$$r = A_f e^{-E_a/RT} f(C_i).$$

A small increase in temperature causes an unproportionally big increase in pressure. The TE process does not have an integrated heat exchange and the pressure is controlled by the gas purge valve which is very small and therefore not effective in controlling rapid pressure rises. Moreover reactor temperature usually requires a tight control with a proportional-integral-derivative (PID) controller [14], whereas in the model all controllers are PI controllers.

It is apparent that P_{reac} exhibits slow dynamic under one attack and fast dynamic under another attack. There are many different factors which influence the behavior of a physical phenomena and of a control loop under attack. Among others are the kind of a relationship between the interdependent physical parameters and the way a physical phenomenon is being controlled, in particular, the configuration of the control loop which includes the choice of the MV, type of the control algorithm and tuning parameters of a controller (PI coefficients).

Oscillation is a very undesirable process behavior and is a prominent symptom of deteriorated control. Attack F_{sep}^{low} causes an oscillating response throughout the entire TE plant which eventually shuts down in 7.5 hours on low separator liquid level (Fig. 4(c)). In the normal case the responsible controller should be re-tuned to avoid oscillation. Howbeit it turned out that such process response to the assault can be also beneficial. For example, attack L_{sep}^{high} also leads to a shutdown on L_{sep}^{low}. However as can be seen in Fig. 4(d) in this case the level decreases rectilinearly which significantly reduces SDT.

The results from Table 2 tell us nothing about the sensitivity of control loops to the magnitude of the manipulation. One of the dimensions for analyzing the process is measuring STD under the aggressive integrity attacks. In this case readings of a sensor i are forged as boundary values of the complete sensing range: X_i^{min} and X_i^{max}. The analyzed model exhibits most resilience to the attacks on XMEAS{7;14;15}. However the simulations reveal that aggressive attacks have no impact in case of tampering with XMEAS(7); slight impact in case of XMEAS (15) and significant impact for XMEAS (14) leading to a shutdown in 8 minutes.

Attacks on certain sensors have local effect and on others – plantwide. For example, attack F_C^{high} has as a consequence shutdown on high reactor pressure and attack F_E^{high} on low stripper liquid level. Fault propagation is usually undesired and especially worrisome for the cases with short SDT. In this case the operator would have a limited time window for identifying the root cause of the unwanted behavior and for taking corrective measures.

Attacks on three liquid levels (reactor, separator and liquid) and related sensors exhibit consistent coupled effects. This is because levels are usually assessed jointly during the individual design of their control loops [14]. If the attacker is aware of the volumes of the tanks, she can launch an attack strategically to maximize the impact. Reactor has the biggest volume ($16m^3$) and claiming its levels as being high or low will cause a chain effect of a high amplitude, which separator ($4.9m^3$) and stripper tanks ($4.4m^3$) cannot accommodate.

Economic impact. The economical performance of a plant can be estimated twofold. Firstly, via operating costs, expressed in $\$/h$ and calculated as [7]:

$$(purge\ costs)+(product\ costs)+(compressor\ costs)+(steam\ costs)=total\ costs.$$

Operational costs are primarily determined by the loss of raw materials. The purge rate has the greatest impact on cost due to material losses in purge and costs of running the compressor which cycles non-condensed components back to the reactor feed. Any attacks which have an impact on the purge flow will have a proportional impact on the production costs. Our analysis reveals that those integrity attacks which cause great variation of process dynamic have an impact on the hourly costs rate. However they just cause higher cost fluctuations without influencing the mean value or compound costs.

Another economic indicator is product quality $G\ mol\%$ in the product flow. No conducted attack caused an impact on product composition which could be

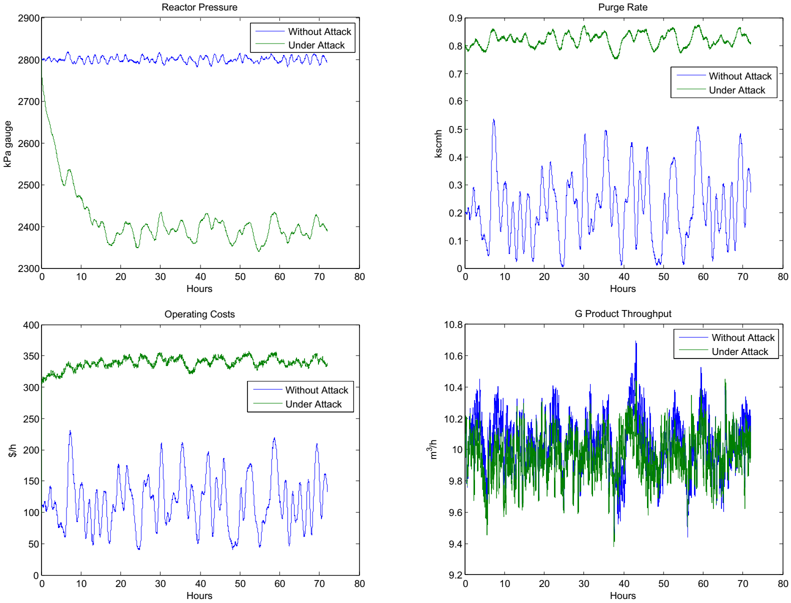

Fig. 5. Economic impact of P_{react}^{high} attack

considered as harmful. This indicates that the implemented control configuration is very robust in keeping the required product composition. Therefore, we introduced an additional metric – *product throughput* – to estimate the influence of an attack on the final product output:

$$F_{strip} m^3/h \times Gmol\% = Gm^3/h.$$

The most effective attack on plant economy is reporting P_{reac}^{high} in order to decrease reactor pressure. As shown in Fig. 5 the controller responds to such an attack by opening the purge valve. This in turn causes a decrease of pressure to a very low level. High losses in the purge will result in corresponding significantly increased operation costs. Also, since the feed of the reactants will be regulated to the lower rate, the output quantity of the final product will also decrease.

4.2 DoS Attacks

As discussed above, during the DoS attack on a sensor a controller stops receiving fresh measurements. As a result, the controller will keep generating control commands based on the last received reading. In a certain sense a DoS attack is similar to an integrity attack with the only difference that the adversary has no direct influence on $X_i^a(t)$. Instead an adversary can take advantage of the timing parameters of an attack, such as its starting time t_a and the duration T_a. Fig. 6 demonstrates the outcome of the DoS attacks on the reactor pressure sensor at

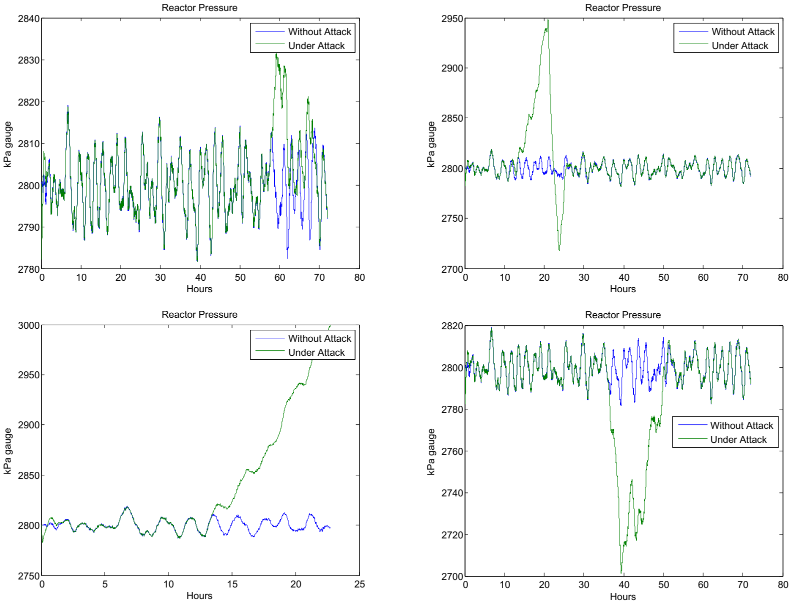

Fig. 6. Random DoS attack on P_{reac}, T_a=10 h

a random time with T_a=10 h. As can be observed, depending on t_a the impact varies from negligible to a shutdown. Apparently, the attacker has to develop a strategy to determine the optimal attack time t_a. Furthermore we performed initial evaluation of how T_a influences the adversary's chances to bring the plant into an unsafe state. The simulations revealed that for T_a <10 h the chances are rather low, whereas for T_a >15 h the chances are rather high.

4.3 Application of the Results

The analysis conducted helps do discover the weaknesses of a process design in the presence of cyber attacks. Our initial exploratory research into a process can be used for designing countermeasures as starting points to improve the security posture of industrial processes.

Security aware control strategy. The design of any control system (as of any engineering system) starts with the requirements. A viable control strategy not only satisfies operational and economic goals but is ideally also able to absorb the greatest anticipated disturbance. Although disturbances are considered as being fortuitous events, long process operation history has accumulated substantial experience about the types of possible operational disruptions. Results of the process-aware security assessment of a plant can equally serve as an input to the design of the control strategy.

In practice, it is hardly possible to design a single control structure capable of accommodating all operational objectives. Therefore often one or more alternative control strategies are developed in parallel to compensate for the weaknesses of the other control configurations. This is called dynamic controllability. One of the most widely applied techniques for alternative control is the usage of an *override controller* [13] which can take command of a MV away from another controller when otherwise the process would exceed some process or equipment limit or constraint. Such selective control keeps the equipment running although perhaps at a suboptimal level. Attack L_A^{low} is similar to the disturbance IDV(6) from the TE challenge problem. No basic *regulatory* control strategy can successfully reject this disturbance. Therefore most of the developed control structures are modified with overrides to handle this situation [21], [14]. The approach of using overrides can be similarly applied to compensate for the other process impairments caused by the cyber-attacks described above.

The value of classifying control loops according to their importance for plant operations is also recognized by control engineers [14]. Based on the assessment of process resilience to the attacks, sensors and control loops can be categorized based on their impact on plant safety. Those that entail safety compromise in minutes (e.g. attacks on T_{reac} or F_C) could be more closely monitored and tightly controlled. Moreover, additional protective measures could be applied to important sensors and controllers, e.g. anomaly detection techniques specific to cyber-physical systems [16], [4].

Another approach to improve the survivability of physical processes under cyber-attacks is resilience-aware network segmentation. As proved in [9] such network design can significantly improve the tolerance period that would give operator more time to intervene. This is a hybrid strategy when control and network configurations can be beneficially considered jointly.

Human Response. Requirement for better human responses to abnormal situations is a recognized industrial problem [1]. Many safety accidents happen because of the non-identification or late identification of process degradation as well as because of wrong corrective actions. Operators could be trained to recognize abnormalities which might be caused by intentional manipulations (in contrast to natural events) and to divert irregularities away from production- or safety-critical to non critical variables. Results on control loops resilience to DoS attacks can be used for intervention action, e.g. for temporary disconnection or switching off of suspected equipment.

5 Attacks on Situational Awareness

Industrial process dynamic is monitored by operators via a Human Machine Interface (HMI) console around the clock. Upon observing an undesired process behavior, an operator takes corrective measures to bring the process back into its steady state. Moreover, if the operator attributes the disturbances as being of unnatural causes, she can initiate an immediate incident investigation. Out of

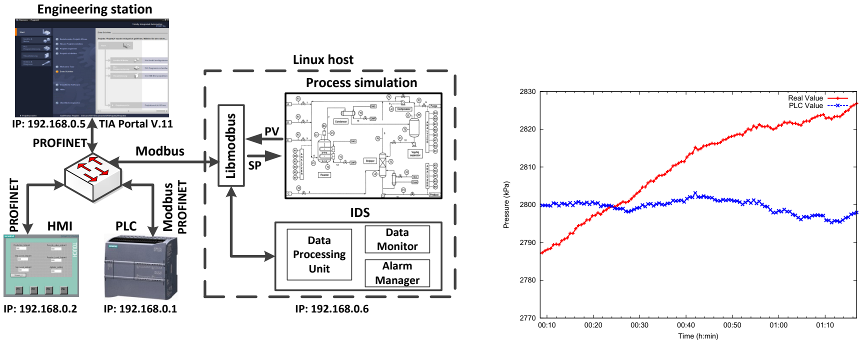

(a) IDS architecture (b) Data Monitor

Fig. 7. Data inconsistency detection

this considerations the attacker might prefer to hide the real field data from the operator. Let the adversary's goal be to raise the pressure in the reactor to an unsafe limit without the operator's awareness. One of the possibilities to achieve this is to record steady state process data and replay them to the operator during the attack. As a result, the operator loses *situational awareness*. This is one of the most dangerous attacks on process control. If the attacker manages also to manipulate the safety limit value or suppress the safety systems communication link, the reactor can actually explode and injure personnel in its vicinity [3].

To model and to detect such type of attacks we have implemented an experimental framework in the form of a hybrid process control environment as depicted in Fig. 7. It is based on the Siemens SIMATIC S7-1200+KTP400 Starter Kit hardware and industrial protocol Modbus/TCP. We use the libmodbus library [12] to enable communication between the simulated process and the HMI. The Programmable Logic Controller (PLC) polls selected PV to display them in the HMI and forwards the setpoints to the process. Modbus protocol utilizes Client-Server communication model. Therefore it is required to install Modbus Server and Client on the PLC.

We implemented attacks on situational awareness through manipulation of the PLC code. During the initial stage of the attack, the PLC records process measurements during normal plant operations. When the attack begins, the PLC sends stored data to the HMI whereas the real field data remains undisclosed. To detect this we implemented an experimental IDS engine. We monitor data flows between the process and the PLC and between the PLC and the HMI. Any discrepancy in the process value between indicated data flows will indicate an *attack on data consistency*. To watch over the specified data flows on one hand we query the output registers of the PLC for the data which should be displayed in the HMI. On the other hand we capture the traffic between the process and the PLC. If an inconsistency in data is detected as shown in Fig. 7(b), an alarm is generated by the Alarm Manager.

6 Final Remarks and Future Work

Establishing control objectives is a first step in a plantwide control design procedure [14]. Therefore the requirements related to the security aspects of plant operations should be determined upfront and included among the set of the control goals.

Conducting process-aware cyber-risk assessment helps in discovering the weaknesses of process design in the presence of cyber attacks. However, examination of the complete set of controllers under multiple types and modes of attacks is an onerous task. Moreover, this activity will inevitably clash both with the usual low availability of time and resources to perform such an analysis and with a lack of expertises on how to recognize, locate and respond to the attacks. This area of research still needs to be advanced from the process engineering standpoint.

Plant stability is another crucial performance characteristic with a direct impact on the global plant bill. Attacks on certain sensors cause higher variability in plant dynamic without challenging safety constraints. However such fluctuations are highly undesirable for two reasons. Firstly, they increase movement of the valves which not only wears out the equipment, but also introduces additional disturbances. Secondly, they cause variations in the input and output streams of the plant which in turn negatively affect interdependent up- and down-stream operational units.

Operational targets and security requirements may conflict and have to be considered in conjunction. For instance, it was shown that the optimal operating steady state condition for P_{reac} is as close as possible to the upper shutdown limit of 3000kPa and for L_{reac} to its lower bound [20]. In this case the attacker will be able to bring the system into an unsafe state quickly. To ensure secure operations it would be desirable to maintain a sufficient safety margin. However, maintaining a safety margin for P_{reac} of at least 100 kPa is equivalent to a 5% increase in cost [21].

The consequences of the attacks were not always predictable. For example, manipulations of feed flow sensors provoke very diverse system reactions. Also the time it takes to achieve the attack goal varies from a few minutes to a few hours. The attacker would need to compromise different sensors if targeting plant safety, operating costs or plant stability. Therefore attacking a sensor at random might not help an attacker to achieve her goal at the first attempt. However, conducting multiple attacks may raise suspicion. We believe that targeted attacks are to proceed with espionage attacks, e.g. [6].

Future research will concentrate on subsequent experimental work on process models to develop a systematic approach to cyber-security assessment of industrial control systems. Further work remains to be done on the TE model: (1) analyzing the impact of DoS attacks on the other sensors; (2) studying the impact of the timing parameters of the attacks, in particular in case of DoS attacks. Finally we would like to explore the opportunities of responding to attacks by the means of process control, namely the dynamic reconfiguration of the process control when confronted with abnormal behavior.

References

1. Abnormal Situation Management (ASM) Consortium: Official website, `https://www.asmconsortium.net/` (retrieved: June 2013)
2. Anderson, R., Fuloria, S.: Security economics and critical national infrastructure. In: Economics of Information Security and Privacy, pp. 55–66 (2010)
3. U.S. Chemical Safety Board: Runaway: Explosion at T2 laboratories (2007), `http://www.youtube.com/watch?v=C561PCq5E1g` (2009) (retrieved: May 2013)
4. Cárdenas, A.A., Amin, S., Lin, Z.S., Huang, Y.L., Huang, C.Y., Sastry, S.: Attacks against process control systems: risk assessment, detection, and response. In: Proceedings of the 6th ACM Symposium on Information, Computer and Communications Security, ASIACCS 2011, pp. 355–366 (2011)
5. Chabukswar, R., Sinopoli, B., Karsai, G., Giani, A., Neema, H., Davis, A.: Simulation of network attacks on SCADA systems. In: First Workshop on Secure Control Systems (2010)
6. Chien, E., O'Gorman, G.: The Nitro attacks: Stealing secrets from the chemical industry. Tech. rep., Symantec (2011)
7. Downs, J.J., Vogel, E.F.: A plant-wide industrial process control problem. Computers & Chemical Engineering 17(3), 245–255 (1993)
8. Genge, B., Siaterlis, C., Hohenadell, M.: Impact of network infrastructure parameters to the effectiveness of cyber attacks against industrial control systems. International Journal of Computers, Communications & Control 7(4), 673–686 (2012)
9. Genge, B., Siaterlis, C.: An experimental study on the impact of network segmentation to the resilience of physical processes. In: Bestak, R., Kencl, L., Li, L.E., Widmer, J., Yin, H. (eds.) NETWORKING 2012, Part I. LNCS, vol. 7289, pp. 121–134. Springer, Heidelberg (2012)
10. Gollmann, D.: Veracity, plausibility, and reputation. In: Askoxylakis, I., Pöhls, H.C., Posegga, J. (eds.) WISTP 2012. LNCS, vol. 7322, pp. 20–28. Springer, Heidelberg (2012)
11. Huang, Y., Cárdenas, A., Amin, S., Lin, S.Z., Tsai, H.Y., Sastry, S.S.: Understanding the physical and economic consequences of attacks against control systems. International Journal of Critical Infrastructure Protection 2(3), 72–83 (2009)
12. libmodbus Project: Official website, `http://libmodbus.org/` (retrieved: June 2013)
13. Liptak, B.G.: Instrument Engineers' Handbook. Process Control and Optimizatiol, vol. 2. CRC Press (2005)
14. Luyben, W.L., Tyreus, B.D., Luyben, M.L.: PlantwideProcess Control. McGraw-Hill (1998)
15. McAvoy, T., Ye, N.: Base control for the Tennessee Eastman problem. Computers & Chemical Engineering 18(5), 383–413 (1994)
16. McEvoy, T., Wolthusen, S.: A plant-wide industrial process control security problem. In: Critical Infrastructure Protection V, vol. 367, pp. 47–56 (2011)
17. McIntyrel, C.: Using Smart Instrumentation. Plant Engineering (2011)
18. Ricker, N.L.: Tennessee Eastman Challenge Archive, `http://depts.washington.edu/control/LARRY/TE/download.html` (retrieved: May 2013)
19. Ricker, N.L.: Model predictive control of a continuous, nonlinear, two-phase reactor. Journal of Process Control 3(2), 109–123 (1993)
20. Ricker, N.: Optimal steady-state operation of the Tennessee Eastman challenge process. Computers & Chemical Engineering 19(9), 949–959 (1995)
21. Ricker, N., Lee, J.: Nonlinear model predictive control of the Tennessee Eastman challenge process. Computers & Chemical Engineering 19(9), 961–981 (1995)