



**Rocking the pocket book:
Hacking chemical plants
for competition and
extortion**

Marina Krotofil
Jason Larsen

WHITE PAPER, DEFCON 23

FURTHER CO-AUTHORS AND CONTRIBUTORS:

Alexander Isakov • Alexander Winnicki • Dieter Gollmann • Pavel Gurikov

DAMN VULNERABLE CHEMICAL PROCESS

<https://github.com/satejnik/DVCP-VAC> • <https://github.com/satejnik/DVCP-TE>

This research was done in Hamburg University of Technology, Hamburg, Germany.

August 2015



Contents

1	Introduction	5
1.1	Process Control Systems	6
2	Classes of Cyber-Physical Attacks	8
2.1	Equipment damage	9
2.2	Production damage	9
2.3	Compliance violation	9
3	Stages of Cyber-Physical Attacks	11
3.1	Access	11
3.2	Discovery	12
3.3	Control	13
3.4	Damage	13
3.5	Cleanup	14
4	Vinyl Acetate Monomer Process	15
4.1	Process Description	15
4.2	Control Model	17
4.3	Simulation of Attacks	17

5	Attack for Production Damage	19
5.1	Preliminary Analysis	20
6	Attaining Attack Objectives	21
6.1	Access	21
6.2	Discovery	21
6.3	Control	22
6.4	Damage	25
6.5	Cleanup	28
6.6	Discussion	31
7	Damn Vulnerable Chemical Process	33
7.1	Framework description	34
7.2	Comparison of TE and VAM processes	38
8	Conclusion	40
	References	41



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. Complex machines such as aircrafts or robots, building automation systems, smart cities and smart grids, railways and agricultural systems, medical devices and industrial infrastructures in general are examples of cyber-physical systems. Perceived and real security threats affecting cyber-physical systems have been attracting considerable attention in the media, among decision makers, regulators, and in the research community.

Cyber-physical systems span cyberspace and the physical world. The fact that they can cause tangible effects in the physical world and that, thereby, attacks in cyber space can have physical consequences has been a major reason for the current interest in this branch of security. The concern for physical consequences puts cyber-physical systems security apart from *information security*.

On one hand, this is an issue that had to be dealt with already before physical systems were connected to cyberspace. Well-designed systems would have been deployed with appropriate safety measures in place. Conceivably, those measures can restrain cyber-physical attacks once they have transited from cyberspace into the physical domain. On the other hand, those countermeasures were designed under certain assumptions, e.g. physical security protecting access to premises or independence of component failures. Conceivably, those assumptions get invalidated once modern IT systems get integrated with existing physical plants.

Integrating modern IT systems with existing physical systems exposes those installations to new security threats. Some of these threats are well-known in IT security and countermeasures have been studied at length. Those threats are new only because of a new application area. Other threats may indeed be specific to cyber-physical systems. Our work aims at making a distinction between “old” security issues in new settings, and new security aspects intrinsic to cyber-physical systems that would establish cyber-physical security as an object of study in its own rights.

While compromising or disrupting devices or communication channels used to sense or control a physical system is a necessary requirement to attacks aimed at disrupting the physical process, the damage from the attack will be limited if the attacker is unable of manipulating the control

system in a way to achieve her desired outcome in the physical world. After all, breaking into a system is not the same as breaking a system. In order to achieve a desired impact on a control system (like Stuxnet [24]), the attacker needs to assess how her attack will perform at the regulatory control level. Launching such an attack requires a different body of knowledge from the one used in IT security. In particular, attackers need to know how the physical process is controlled, and that includes knowledge of failure conditions of the equipment [25], control principles [36], knowledge of process behavior [30], and signal processing, etc.

Your wish is my command?

It is often claimed that “once communications security is compromised the attacker can do whatever she wants”. These are presumptuous claims. The attacker may well be able to inject any input she wants but this does not necessarily amount to being able to influence processes in the physical world at will. The processes and their actuators have to be properly understood. Process physics and built-in safety measures might get in the way of the attacker.

In our work we consider the physical part of attacks and examine the hurdles an attacker might face when trying to manipulate physical processes, using realistic simulation model of a vinyl acetate plant as a case study. In our work we demonstrate a complete attack, from start to end, directed at persistent economic damage to a production site while avoiding attribution of production loss to a cyber-event. Such an attack scenario could be useful to a manufacturer aiming at putting competitors out of business or as a strong argument in an extortion attack.

Process simulators play an important role in process control studies. Rigorous non-linear process models are useful tools for accurately understanding process dynamics, and thus can be used in both control structure development and validation. At this point in time, very few models of industrial processes are available for analysis by security researchers. To draw general lessons for cyber-physical systems security, be it on the true power of an attacker or on the efficacy of the defenses deployed, many more documented case studies will be necessary.

1.1 Process Control Systems

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, 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. 1.1a. 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. 1.1b). 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.




Figure 1.1: Types of control loops

Industrial Process Control Systems are used to provide autonomous control over a production process through control loops. They monitor the production process via sensors deployed around the product line and interact with the process through actuators. The complexity of modern production processes is usually simplified by dividing the control load into subsystems containing separate control loops. Heavy control loop couplings among subsystems are avoided.

In order to enhance the security of an process control systems from a system perspective, secure control theory which studies how cyber attacks affect the physical dynamics of the system has been explored in recent years.



2. Classes of Cyber-Physical Attacks

Modern industrial plants face multiple challenges – to deliver products at consistent quality and possibly low cost, to manage plant dynamics altered by material recycle and energy integration, to satisfy environmental and safety regulations, and to have a certain degree of flexibility to handle fluctuations such as production rate changes (in response to changing market demand) and feed quality. All of these are the responsibilities of a reliable and efficient control system. Modern plants are now becoming more complex than just the union of a set of unit operations.

In information security, the attacker’s goal may be to exfiltrate information or to disrupt the normal operations of software. In the cyber-physical domain, the attacker’s goal is to disrupt the normal operations of control systems. When weaponizing a buffer overflow, shellcode is constructed that instructs a system to perform specific actions desired by the attacker. Similarly, in cyber-physical exploits the attacker’s payload will contain a set of instructions that manipulate the process, and the choice of instructions depends on the specific impact the attacker wants to have on the process.

So what can actually be done to a process? The effects of cyber-physical attacks can be classified into three groups. Admittedly, the classes outlined are interrelated as damage of one kind may lead to another kind of damage. For example, production can be disrupted through breakage of equipment. Runaway reactions can cause serious safety accidents and equipment breakage. However, a clear understanding of the attack goal is necessary for maximizing attack impact and minimizing the cost of attack implementations avoiding “over-engineering”.

Classes of cyber-physical attacks		
Equipment Damage	Production Damage	Compliance Violation
<ul style="list-style-type: none">• Equipment overstress• Safety limits violation	<ul style="list-style-type: none">• Product quality• Production rate• Operating costs• Maintenance efforts	<ul style="list-style-type: none">• Safety• Pollution• Contractual treaties

2.1 Equipment damage

This class of attacks aims for physical damage of equipment or infrastructure (e.g. pipes, valves). Larsen [25] discusses classes of physical damage. Equipment damage can be achieved in two ways.

Overstress of equipment. Every equipment wears out or breaks at the end of its expected life cycle. Prolonged overstress of equipment can accelerate this process. An example are wear-off attacks on valves due to unstable process control. This type of attack was implemented in the second version of the Stuxnet worm [24].

Violation of safety limits. The second option is to violate safety limits, ideally in some smart way. In this way researchers at Idaho National Labs remotely destroyed a power generator [39]. This type of attack was also realized in the first version of Stuxnet [24]. Those targeting at continuous processes can consult [27] for safety limits of piping infrastructures and related equipment.

2.2 Production damage

Instead of breaking equipment an attacker can go after the production process to spoil the product or make production more expensive. Attacks on production can be divided into three groups.

Product quality and production rate. Attacks may be directed at the product itself – its quality or production rate. Every product has its specification and market prices for a specific quality. The attacker may turn the product unusable or reduce its value. The price of a product may rise exponentially with product purity. Table 2.1 presents relative prices for paracetamol. As can be seen, not achieving the desired product quality can be very expensive.

Purity	Price, Euro/kg
98%	1.0
99%	5.0
100%	8205.0

Table 2.1: Relative paracetamol prices. Source: sigmaaldrich.com

Operating costs. After the process is tuned, the operator's primary task is to keep the process as close as possible to the economically optimal operating conditions. Every plant has an objective cost function consisting of several components which impact the operating costs. It may be loss of raw materials in the purge, premature deactivation of the catalyst, or increased energy usage.

Maintenance efforts. The attacker can impact a production process by increasing the maintenance workload. Maintenance refers to troubleshooting process disturbances and equipment malfunction. For example, rapid operation of a flow valve causes a damaging cavitation process - the formation of vapor cavities in a liquid. Cavitation eventually wears the valve and leads to leaks (requiring valve replacement); also bubbling of a liquid substantially complicates process control.

2.3 Compliance violation

Industrial sectors tend to be strongly regulated to ensure safety and to protect the environment. Non-compliance can attract fines and bad publicity, unlike attacks whose effect can be kept internal to a company.

Safety. Most damaging would be attacks on occupational and environmental safety as they may result in lethal accidents and serious environmental damage. This type of attack in most cases will yield collateral damage.

Environmental pollution Less dramatic would be attacks causing regulatory pollution limits to be exceeded. This can relate to the concentration and volume of gaseous emissions, water or soil contamination and similar. For example, if effluent from an industrial facility fails to meet local regulatory standards, the plant can be fined, and recurrent offenses can lead to plant shutdown. Negative impact on reputation may be a further consequence.

Contractual agreements. Typically this refers to production schedules. Take vaccine production as an illustrative example. Reactions to outbreaks of a disease often lead to political and public pressure. Missing delivery schedules may cause contractual sanctions and bad publicity.



3. Stages of Cyber-Physical Attacks

An attacker targeting a remote process may not immediately be gifted with complete knowledge of the process and the means to manipulate it. An attack may have to go through several stages before the evil goals can be achieved (Figure. 3.1). Perfect knowledge is never achieved and the attacker may need to circle back to previous stages or recursively repeat her exercises at the same stage.




Figure 3.1: Stages of cyber-physical attacks

3.1 Access

Access is the stage most resembling traditional IT hacking. In general, the attacker needs code executing somewhere in the victim's network to manipulate the process and thus has to find some way in.

A process network is usually connected to a business network and a field network along with various regulatory links relevant for any hazardous substances used. Along with all the data streaming off the control network that feeds corporate and third party systems, process control systems have many of the same needs as IT systems. Patches and anti-virus updates must flow into the network. Control commands must also flow out of the network to field equipment. Regulatory data must be sent to various agencies. Sometimes the data must be sent in real time. These data flows can be potential ways into the process network. This stage is largely the same as hacking into any other network.

For inspiration and examples the readers may check the following references. Get your malware in via SCADA application [16] or whitelisting security service [17]. Select attractive vulnerability from [14], find vulnerable Internet facing devices as described in [28], exploit [2]. Rootkitting techniques for field devices (sensors, actuators) can be found in [26].

3.2 Discovery

This refers to discovering information about a plant from documentation. Without detailed knowledge, it is unlikely that an attacker can achieve more than nuisance. Blindly trying to destroy a process by overheating a tank, for example, will probably only result in exercising the emergency shutdown logic and the pressure relief valves.

Espionage. Plants can be highly proprietary. If ten chemical engineers were tasked with building a talcum powder plant, they may build ten very different processes. Even if they were restricted to the exact same chemical process, they could make different choices in the design of the plant. They might choose different vendors for pumps and valves. That in turn would influence the sizing and placement of various pipes and holding tanks. That in turn would change the way the plant is controlled and the design of the control loops.

Espionage and reconnaissance. The attacker must thus reconstruct the layout of the plant and how it carries out its functions. This is the most difficult and time consuming of the stages. There are several data sources that describe the process. The attacker may first study general information on the chemistry, kinetics, and thermodynamics of the physical processes of interest. This can be done by consulting open literature as well as proprietary information of process design companies. Operator screens are meant to be a human readable. Regulatory filings will describe the inner workings of safety or environmental related subsystems. Engineering diagrams may be stored in change management systems electronically so changes in the physical process can be matched to changes in the control logic. This part of the discovery stage may also involve espionage.

Typical documents on plant design are Piping and Instrumentation Diagrams (P&ID), One-Line Diagrams which often contain information on safety interlocks, Cause & Effect Diagrams, Cable Schedule Diagrams, Project Interconnection Diagrams and others. Instrument I/O (input/output) Lists contain a list of instruments which serve as input or output of the control system. This list will contain such information as type of instruments and their location, range of set points, instrument tags and loop numbers, service descriptions, etc.

The indispensability of this attack stage have been realized by attackers LONG time ago. From what is made public, the APTs targeting industrial and research organizations were in action already in 2006 [1] and probably earlier [37]. The massive spying on global oil, energy, and petrochemical companies were carried out in 2009 onwards [31, 7]. The attackers appeared to be especially interested in design documents, formulas, manufacturing processes and research materials. While a large number of espionage campaigns took place in the past years [37], one of them deserves closer

attention as it also included reconnaissance capabilities [13]:

Using OPC, the malware component gathers any details about connected devices and sends them back to the C&C for the attackers to analyze. It appears that this component is used as a tool for intelligence gathering. So far, we have not seen any payloads that attempt to control the connected hardware.

As always: we are living in interesting times. It is just a matter of time until “next generation” malware capable of process control (see next section) will be caught in wild.

3.3 Control

In dynamic systems such as cyber-physical systems, the values of process variables change with time according to the laws of physics. However, transitioning of a process from one state to another is in most cases not instantaneous and adheres to the well-known fact that “things take time”. At this stage the attacker tries to discover the dynamic behavior of the process which can be described in the form of simple differential equations $dy/dt = f(y,u)$, where u is an independent variable and y is the dependent variable which are related by cause-and-effect relationships.

It is easy to discount the difficulties involved in making a process misbehave in a predictable way. The various pieces of a process are connected together in physical relationships. Increasing the temperature in a vessel will usually increase the pressure as well. Adding more chemicals from a holding tank will either decrease the amount of another chemical or increase the flow rate in the pipe. The attacker thus needs to develop a knowledge of all the side effects of manipulating the process.

The process is not designed for the attacker. Every part of the process has a normal range and a possible range. Adjusting one part of the process for malicious purposes may have side effects on other parts of the process. The control stage studies what each actuator does and what side effects are possible. It may be possible to turn off a pump, but the side effect is that pressure builds up quickly in an upstream pipe. Not every action can be taken at every level of the process. Instructing a breaker to close while a line is charged may be prevented by an interlock. The attacker may need to hack an embedded controller to override that interlock. The control stage also involves the study of timings. If the damage occurs in seconds, a safety shutdown minutes later will not stop the attacker.

Some part of control can be studied statically, but other parts must be investigated dynamically on the process (process reconnaissance). No diagram will every be detailed enough to accurately predict the travel time of a disturbance down a pipe to the accuracy needed to set up a resonance between two pumps. Since that data must be extracted from the live process, this is a great chance for the defenders to notice the attackers.

3.4 Damage

Once the attacker understands the process and how to control it, she needs to decide how exactly to achieve her goals. There may be several competing scenarios. The attacker will need to develop some sort of measure (a metric) to choose between them. Bouncing some pumps off the floor until they break may be a good idea, but the economic damage to the target plant may be far less than, e.g., poisoning the catalyst in the reactor.

Physical damage to a process may not come to a process engineer’s mind first. Engineers asked to attack a process tend to come up with what is known as “salty cookie” scenarios. A group of engineers asked to attack a cookie factory all hit on variations of putting too much salt in the cookies

so they would become uneatable. However, in a real world scenario the actual damage to the food factory came from simultaneously creating too much product and disabling the emergency flushing system. The resulting clogged pipes had to be physically replaced after both water pressure and chemical means of clearing them failed.

Accident data can be a good starting point when studying damage scenarios for a process. If a particular type of process has gone wrong in the past by accident, it stands to reason that an attacker may be able to make the process fail in a similar way by design.

While the access phase is the one most familiar to a traditional IT hacker, the damage phase is the one least familiar. It often requires the input of subject matter experts to understand the full range of possibilities.

3.5 Cleanup

In traditional IT hacking, a goal is to go undetected. In most process control scenarios, this is not an option. If a piece of equipment is damaged or if a plant suddenly becomes less profitable, someone will be sent to investigate. An attack will change things in the real world that cannot be removed by simply erasing the log files.

The cleanup phase is about creating a *forensic footprint* for investigators by manipulating the process and the logs in such a way that the analyst draws the wrong conclusions. The goal is to get the attack blamed on operator error or equipment failure instead of a cyber event.

An example of a cleanup phase would be to show the operator a process out-of-control, making her take a particular action. When investigators ask the operator if she was manually manipulating the process when it malfunctioned, she will answer in the affirmative. Another example would be damaging an actuator upstream of the attack to focus investigators towards a previous stage in the process.

Synopsis

Access	Traditional hacking
Discovery	How is this place built and controlled?
Control	What and how much can I change?
Damage	What evil things can I do?
Cleanup	What will they think happened?
<hr/>	
	Mix and repeat until done



4. Vinyl Acetate Monomer Process

Vinyl acetate monomer is a large-scale commodity chemical and is an essential chemical building block used in a wide variety of industrial and consumer products. VAC is a key ingredient in resins, intermediates used in paints, adhesives, coatings, textiles packaging, automotive plastic fuel tanks and many other final products. Detailed information about the product including regulatory, health, environmental, and physical hazard information can be found e.g. on the web page of The Dow Chemical Company [38].

4.1 Process Description

In the VAC process, there are ten basic unit operations, which include a vaporizer, a catalytic plug flow reactor, a feed-effluent heat exchanger (FEHE), a separator, a gas compressor, an absorber, a carbon dioxide CO₂ removal system, a gas removal system, a tank for the liquid recycle stream, and an azeotropic distillation column with a decanter.

The route for vinyl acetate manufacturing used in the process model is the same as employed in today's manufacturing and involves seven chemical components. Ethylene C₂H₄, oxygen O₂, and acetic acid HAc are provided as both fresh and recycled feeds and are converted into the vinyl acetate with water H₂O and carbon dioxide CO₂ as byproducts. The fresh C₂H₄ stream contains an inert component C₂H₆. The following reactions take place in the reactor:

Main reaction: $C_2H_4 + CH_3COOH + \frac{1}{2} O_2 \rightarrow CH_2=CHOCOCH_3 + H_2O$,

Side reaction: $C_2H_4 + 3 O_2 \rightarrow 2 CO_2 + 2 H_2O$.

The reactor contains tubes packed with a catalyst. Both reactions are highly exothermic and require tight control of the reactor cooling. The side reaction of ethylene combustion to CO₂ is highly undesirable as it lowers the conversion and complicates the removal of the reaction heat. Details of the ethylene combustion kinetics in the synthesis of vinyl acetate are presented in [15]

The reactor effluent is sent to the separator, where gas and liquid are separated. The vapor from the separator goes to the compressor and the liquid stream becomes a part of the feed to the

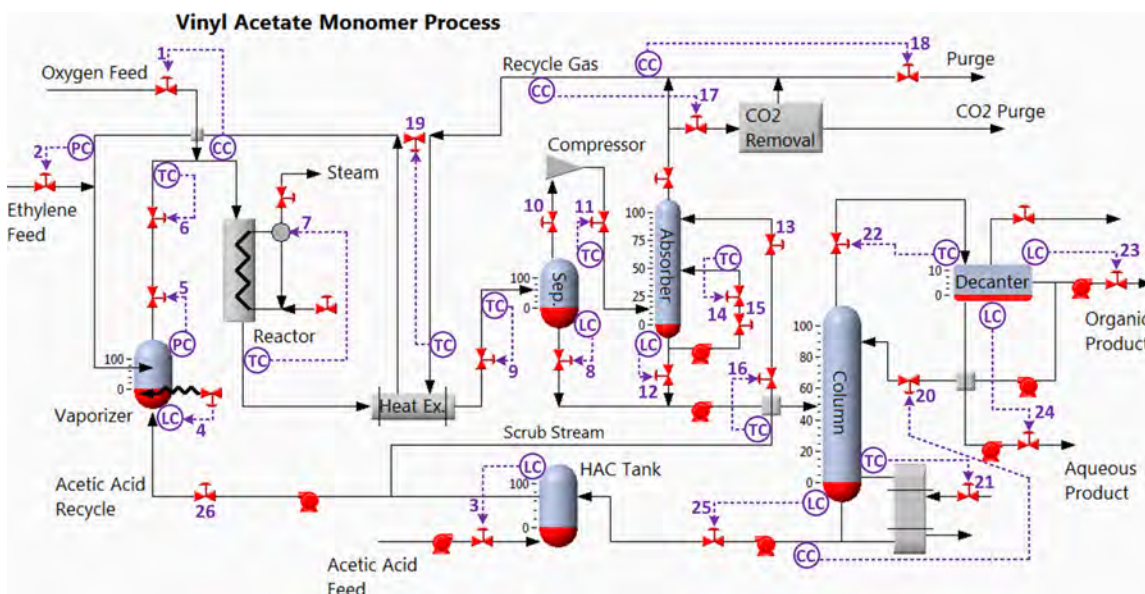


Figure 4.1: Vinyl acetate monomer plantwide process control structure

distillation column. The gas from the compressor is recycled back to the reactor through the absorber and the CO₂ removal system. The liquid products, VAC and water, are withdrawn from the decanter. Fig. 4.1 illustrates the process flowsheet with the locations of control valves.

Readers are referred to [29] and Chapter 10 in [9] for a detailed process description, including the reaction rate expressions, steady state process data and the major aspects of each unit operation. To protect the proprietary information of any specific VAC production facility, the kinetic data, process flowsheet information, equipment data and modeling formulation in the published process came from the sources in the open literature (see [29] and references therein).

Safety constraints

Two key safety constraints exist in the process. Exceeding either of the safety limits will shut down the process via interlocks:

- O₂ concentration must not exceed 8 mol% anywhere in the gas recycle loop to remain outside the explosivity envelope of ethylene. More on the limits of oxygen concentration in gas mixtures can be found in [40].
- The pressure in the gas recycle loop and distillation column cannot exceed 965 kPa (140 psi) because of the mechanical construction limit of the vessels.

Operating constraints

The process constraints must be maintained to ensure efficient production without interruptions for maintenance. They are specified as the following upper and lower bounds for some of the process variables:

- The peak reactor temperature along the length of the tube must remain below 200°C to prevent

mechanical damage to the catalyst requiring shutdown and catalyst exchange.

- Liquid levels in the vaporizer, separator, absorber base, distillation column base, and decanter must operate within the limits of 10-90%.
- Reactor inlet temperature and the hot side exit temperature from the heat exchanger must remain above 130°C to avoid condensation of liquid.
- Organic phase in the decanter must contain less than 600 mol/million of acetic acid to prevent product contamination.
- The VAC composition in the bottoms stream must remain below 100 mol/million to prevent polymerization and fouling in the reboiler and vaporizer.

In contrast to the TE test process, the VAC model is not accompanied with an objective operating costs function for process control optimization. Instead, the economic objective was formulated as balancing trade-offs in maximizing vinyl acetate production and recovery with minimizing carbon dioxide productions and energy consumption.

4.2 Control Model

Similarly to the authors of TE process, Luyben and Tyreus proposed a set of control requirements but did not suggest any process control scheme challenging the research community to come up with their own control approaches, e.g. [30, 33, 9, 35]. The majority of the control design implementations were kept proprietary due to high modeling costs. To make the process model available for a wider range of users McAvoy et al. have developed a simulation model of VAC for Matlab [6] which we use in our experimental work.

The process model includes 246 states, 26 manipulated variables $XMV\{1-26\}$, and 43 measurements $XMEAS\{1-43\}$. Readers are referred to [6] for a complete description of process model formulation, assumptions, and implementation. The process model utilizes a control structure proposed in [30]. Fig. 4.1 depicts the location of control loops. The numbering of control loops follows the numbering in Appendix 2 in [6]. Some manipulated variables are fixed and therefore their control loops are not shown.

4.3 Simulation of Attacks

In cyber-physical hacking the attacker aims to cause tangible impact on the process. In the context of cyber-physical systems the attacker can either modify the control algorithm or tamper with controller inputs and outputs. Input data flow refers to the process data measured by sensors. These are controlled variables (CV) and process measurements. Output data represent control flow of manipulated variables (MV) which update states of the actuators. Communication channels from sensors and to actuators are susceptible to communication jamming (DoS) and data manipulation/injections attacks. We extended the Simulink model to simulate integrity attacks and DoS attacks on the sensors and manipulated variables (actuators). In this work we limited attacker's capabilities to direct integrity attacks on manipulated variables (e.g., change output buffers of the controller or inject packets).¹

Let $Y(t)$ be a manipulated variable at time t , $0 \leq t \leq T$, where T is the duration of the simulation; time is discrete. The attack interval T_a is arbitrary and is limited to the simulation run time. In our setting, we simulate the compromised manipulated variable Y' as follows:

¹The attacker can manipulate state of the actuators also indirectly by forging sensor readings as we presented in [23].

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

where $Y^a(t)$ is the modified manipulated variable.

During a *DoS attack* on a controller signal freshly generated manipulated variables do not reach the actuator. If $Y_j(t)$ is a manipulated variable for actuator j and the attack starts at time t_a , we have:

$$Y_j^a(t) = Y_j(t_a - 1)$$

where Y_t^a is stale manipulated variable (the last MV received from the controller before the DoS attack).

In the context of Process Control Systems DoS attacks are similar to integrity attacks. The only difference is in **how** the attack value is brought about: by choosing DoS approach the attacker has to attack at a *specific time* (e.g. when a valve is all the way open or closed). The advantage of the DoS attacks is that they can be used to manipulate the process even if control traffic is authenticated and integrity protected.



5. Attack for Production Damage

'You can do unfocused and uncontrolled magic without a wand but to do really good spells, yes, you need a wand.'

-Joan Rowling on limits of magic in Harry Potter (2001)

An attacker with an objective beyond simple mayhem will want to reliably manipulate the process and thus implement her attack in such a way that the process is still under control but in a way she needs to accomplish her goals. In the context of cyber-physical systems the “focused magic” is achieved with control theory methods.

Controllability and Observability

Controllability and observability represent two major concepts of modern control system theory [11]. These concepts were introduced by R. Kalman in 1960.

Controllability: In order to be able to do whatever we want with the given dynamic system under control input, the system must be controllable.

Observability: In order to see what is going on inside the system under observation, the system must be observable.

Controllability is about whether one can design control input to steer the state to arbitrarily values. Observability is concerned with whether without knowing the initial state, one can determine the state of a system given the input and the output. Controllability and observability are dual aspects of the same problem: the process must be observable in order to be controllable.

In this work we consider the scenario of economic damage to the plant. One way to influence production costs is to make process inefficient by inducing process disturbances and/or provoking control loops instabilities to increase maintenance efforts. However, creating loop instabilities can be a risky option as the attacker may lose control over the process herself. In the case of unskillful manipulations, the process can even become completely uncontrollable. In contrast, targeting loop effectiveness allows her to remain in control of the process and to adjust disruptive actions over time.

To *persist* with her malign assault, the attacker would want to attract attention from process

operators as little as possible meaning preventing alarms flashing on the operators' screens. Designing an attack scenario is a matter of art as much as economic consideration. The cost of attack can quickly exceed damage worth. Out of this consideration we decided to manipulate the process without triggering alarms to save efforts on alarm suppression or process measurements spoofing. To do so the attacker needs to see the status and alarms for the entire VAC process. One of the good places could be a historian which are often mirrored in near real-time to the business network so that they can be queried by various regulatory processes.

5.1 Preliminary Analysis

Distilled products represent the most valuable commodities leaving a refinery or chemical plant. Therefore maximum economic damage could be achieved by destroying the pipe that carries the final product into the storage vessels. This attack is certainly effective, but it will also be noticed and quickly fixed. For a prolonged damaging effect the attacker needs better scenarios.

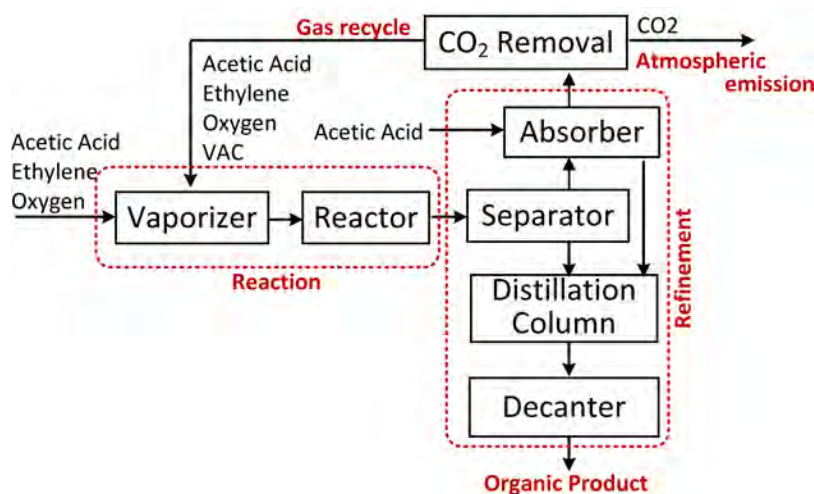


Figure 5.1: Simplified scheme of VAC plant

The VAC plant can be roughly divided into two parts: reaction and refinement (Fig. 5.1). The refinement section is responsible for VAC distillation to ensure a final product meeting rigid specifications. The refinement process consists of multiple units so the attacker has many opportunities to tamper with the process, but the operator also has many opportunities to notice changes and take compensating actions. Moreover, the operator may send impure product for extra refinement. In contrast, upsetting the reaction process inside of the reactor reliably yields reduced production of a useful product. With that, we resolve to attack the reactor unit.

Most factors influencing the reaction process can be described by either deactivation catalyst caused by high temperatures, reduction of reaction rate due to wrong ratio and/or preconditioning of the chemical components or reduction of the primarily reaction due to wrong material and energy balance. These will be the attack scenarios considered in this work.



6. Attaining Attack Objectives

We now have a rough idea about how our specific malicious goal can be possibly achieved but not yet a final attack design. An attack will need to be delivered as a series of commands that manipulate the process to achieve the desired effects. Using the model described above, a specific set of attacks can be selected and developed for final delivery in the malign payload.

6.1 Access

When considering hacking into a VAC process somewhere in the world, one must remember that this is a real facility with real-world needs. Data needs to flow, predictions need to be made, and equipment needs to be maintained. It has been said that all chemical plants start up in a state of imminent failure and then remain that way. When working with a model, it is easy to look at a process environment as static and stable. In reality, there are dedicated teams of workers whose sole purpose is to fix and improve the plant year round. Small armies of workers, consultants, and engineers come and go regularly. An attacker can take advantage of any of these interactions.

6.2 Discovery

The chemical approach to VAC manufacturing is not a trade secret. There is a wealth of information on the process itself and how it is typically implemented in a factory. We assume that the attacker has already obtained knowledge about the process such as presented in Section 4.1. Here, we concentrate on discovering measurements and actuators available to the attacker at the reactor section.

In industrial plants, large numbers of process variables must be maintained within specified limits in order for the plant to operate properly. Therefore process monitoring plays a key role in ensuring that plant performance satisfies the operating objectives. Abnormal process operation can occur for a variety of reasons including equipment problems, instrumentation malfunction, disturbances, etc. Sensor and monitoring equipment placement must satisfy three general objectives: (i) routine monitoring for specified limits; (ii) detection and diagnosis of abnormal operations; (iii) preventive monitoring for early indicators of equipment and process upsets. However, utilization of a large

number of sensors increases the required investment and maintenance costs and optimal sensor placement is an active research area. As a result the attacker may face the challenge of lacking the sensors needed to monitor her attack.

Six sensors, XMEAS{1-6} and one flow composition analyzer are available to attacker at the reactor unit, vaporizer{P;L;T}, heater exit{T}, reactor exit{T;F}, and molar concentrations of the seven chemical elements in the reactor feed stream from analyzer, XMEAS{37-43}. Specifically, O₂ concentration is used to monitor the hazardous conditions related to the explosivity of ethylene in the presence of oxygen. What might catch the attacker's eye is the absence of an analyzer in the reactor exit stream. Thus she will not be able to directly obtain measurements of the molar concentration of the produced vinyl acetate in the reactor outflow.¹

There are seven degrees of freedom XMV{1-7} available for control in the reactor unit, three reactants fresh feeds {O₂;C₂H₄;HAc}, two valves to control vaporize{heater; vapor exit}, reactor preheater valve, and steam drum valve to control the reactor temperature. A quick analysis of the process flowsheet shows that all valves except XMV(3) control effects within the reactor unit itself. Thus, part of the acetic acid inflow controlled by XMV(3) is sent into the vaporizer (reactor unit) and another part into the absorber (refinement section). Also, HAc comes from a supply tank. It means any attack on this feed will be buffered by the acetic acid holdup in the tank. Incidentally, we had discovered a disparity between process documentation and process implementation. Thus, inflow of ethylene XMV(2) is used to control the pressure in the absorber and not the pressure in the gas recycle loop. Thus any manipulation of this control loop will have effects on both the reactor and refinement sections.

A sensor measuring a process variable important for safety or operating constraints will have an alarm or interlock set at certain operating ranges (see Section 4.1). A quick search of chemical engineering journals turns up information on the constraints and specifics of the process. Exact values can be discovered in operator screens, controller logic, one-line diagrams, etc. The basic plumbing of the process can be understood from the flow diagram (Fig. 4.1). In a real world scenario, this information would need to be gathered from configuration files and other sources as described in Section 3.2.

6.3 Control

At this stage the attacker explores dynamic responses of the process to various manipulations. Typically, digital controllers are designed based on process models and with very few exceptions the designs begin with the specification of some desired closed loop properties. The more proprietary information on the process dynamics and controller tuning the attacker can collect, the more accurately she would be able to identify the system. Some of the simulation results of the process response to step changes in a few set points are described in [30]. The authors mention that some dynamic behaviors of this process are not intuitively obvious.

Understanding basic factors that influence dynamic effects of process behavior is fundamental to process control analysis. Thus, every controller is designed and tuned to perform at a certain operational range and may lose its authority in different operating conditions. This is among other reasons related to the non-linearity of most physical processes. It means that process response is not proportional to the applied input. For instance, the dynamic behavior of a process being heated from 140 to 150 degrees will be different from when it is further heated till 160 degrees. If the process

¹Chemical composition analysis systems are expensive; their installation must be justified by important considerations such as safety or significant product quality improvement.

was never expected to be heated to 160 degrees, the control algorithm may never have been tested for its performance in this temperature range.

For two control loops to operate successfully in tandem each loop must “know” what the other is doing. Otherwise trying to achieve their respective objectives the two loops may act against the interest of the other. This phenomenon is known as loop interaction. Some control loop interactions occur naturally as a result of their physical and chemical make-up. Some loop interactions may arise as a consequence of process design. Typical examples are heat integration and recycle streams which create the potential for disturbance propagation and the alteration of dynamic system behavior. Without knowing the exact plant configuration it is not always possible to correctly determine whether a certain process response is related to a fundamental property of the plant or is a result of a specific attack parameter choice.

While the attacker may apply many different attack patterns, in this work we consider the two following attacks:

- **Steady state attacks** – step-like attacks which bring the process into the new state and leave it there (Fig. 6.1a).
- **Periodic attacks** – recurring attacks interleaved with process recovery phases (Fig. 6.1b).




Figure 6.1: Types of attacks on process behavior (red line denotes steady state value)

The control phase is mostly about mapping out the dependencies between each actuator and all of the downstream measurements.

Mapping of the dependencies can be done through detailed modeling or observations on the process itself. An attacker can make small changes in the process and then note those changes as they propagate through the various sensors. Whilst working with the model one can observe alarm generation and propagation. If this was an actual process, the attacker would need to make minor changes and then extrapolate how large a change was necessary to cause an unwanted result.

We applied the following strategy to discovering dynamic process behavior. We first identified the steady state MV values. We then increased or decreased MV by approximately 1% for 30 sec and observed the process response measured by sensors. Depending on the response we increased

the magnitude and duration of the manipulation and monitored the process variables for reaching operational or safety constraints.

Steady state attacks (SSA). Not all actuators are suitable for carrying out steady state attacks. $XMV\{4;5\}$ move the process to its operational or safety constraints within a short time (from minutes to few hours) even if manipulated only slightly.

Periodic attack (PA). This attack scenario can be seen as pulse-width modulation of a steady state attack in which pulse amplitude represents *attack value* (position of the valve), pulse width stands for *attack duration* and inter-pulse distance is *process recovery time*. Examining the sensitivity of control loops to periodic attacks is challenging due to the large number of attack parameters.

Initially we looked at attacks directly setting the position of the valve at the recovery phase (to the steady state value or lower). However, this strategy was unsuccessful for the majority of control loops resulting in critical process variables drifting towards their constraints. This is because MVs were set to fixed values without adjustment to the process dynamics. Therefore we decided to leave administration of the process recovery phase to the controllers.

One of the challenges we faced was control loop ringing while manipulating $XMV(5)$ which controls the vaporizer exit flow. Sometimes, digital controllers produce a control signal that oscillates with decreasing amplitude around the final equilibrium level (Fig. 6.2a). This phenomenon is known as “ringing” and is caused by negative real controller poles. Ringing is an unwanted effect as it increases the wear and tear of controller components and can cause system instability within a multi-loop environment. Manipulation of $XMV(5)$ in the negative direction (decreasing its value) causes a large overshoot which leads to process instability. In our simulations it manifested itself in form of impossible computations.




Figure 6.2: Manipulation of control loop 5

We still pursued this attack as reducing inflow of reactants would result in decreased production of VAC. To overcome the above challenge we decided to take advantage of the negative compensation reaction in the process recovery phase. In our approach we slightly increase the flow for 1-2 minutes and let the process recover for 2 minutes or longer (Fig. 6.2b). In the recovery phase the controller decreased the flow to bring the CV to the set point. In this way we were able to achieve reduced

flow (when averaged over time). Overall, there are three control loops in the reactor unit which can become unstable under certain attack parameters. Those are XMV{2;4;5}.

Outcome of the control stage

We have exhaustively tested all control loops for their sensitivity against a large number of attack settings {attack value; attack duration; recovery time}. Although we could establish a good mental model of process behavior, we needed to find a systematic way to categorize controlled loops. We decided on two parameters: sensitivity to magnitude of manipulation (MM) – how much we can change the process – and to required recovery time (RT). If the process can recover in a time equal to the attack time or shorter we consider such such control loops of low sensitivity. Table 6.3 gives the results of our analysis. Sensitive control loops are risky for reliable control.

Sensitivity	MM	RT
High	XMV{1;5;7}	XMV{4;7}
Medium	XMV{2;4;6}	XMV{5}
Low	XMV{3}	XMV{1;2;3;6}

Table 6.1: Sensitivity of control loops

We concluded with an analysis of alarm activation in response to control loop manipulations (Table 6.2). For each MV we noted upper limits of the attack parameters which would allow to manipulate the process without triggering an alarm (not included into the paper due to space limitation).²

Alarm	SSA	PA
Gas loop O ₂	XMV{1}	XMV{1}
Reactor feed T	XMV{6}	XMV{6}
Reactor T	XMV{7}	XMV{7}
FEHE effluent	XMV{7}	XMV{7}
Gas loop P	XMV{2;3;6}	XMV{2;3;6}
HAc in decanter	XMV{2;3;7}	XMV{3}

Table 6.2: Activation of alarms

6.4 Damage

In the previous stage we evaluated the potential to control the process. In the damage phase the attacker tests those controls to achieve the damage desired. This stage is similar to the “what occurs if” approach deployed in a HAZOP analysis (“what happens if this valve is closed?”).

The attacker must choose one or more attack scenarios to deploy in the final payload. This could be arbitrarily chosen based on gut feeling, but given the amount of effort it takes to mount a real

²At the time of conducting this research the mapping results were stored in a gigantic excel sheet. We are currently working on finding a smarter way to store results of process control stage.

attack, an attacker may well use a metric to measure her success. For an economic attack a plausible metric would be the amount of monetary loss to the victim.

Technician vs. engineering answer

The target plant may not have been designed in a hacker-friendly way and may not measure values needed to monitor attack performance. Also process information may spread across disparate subsystems forcing the attacker to invade a greater number of devices. Generally speaking, there are two types of measurements: “*technician answer*” is a qualitative measure, e.g. whether process measurement will decrease or increase while “*engineering answer*” gives a quantitative answer by how much. When exploiting a process, anything requiring an engineering answer is hard as it often relies on data unique to the plant and its current operating mode.

In order to determine monetary loss one needs to measure how much of vinyl acetate is produced in the reactor. As determined in the discovery stage, there is no analyzer installed in the reactor exit feed. Therefore the attacker does not have an engineering answer. However, the rate of the reaction can be qualitatively determined from the reactor exit temperature. A decrease in temperature signals that less reaction is happening in the reactor, so less product is being produced (Fig. 6.3a). This measure can be sufficient to determine whether a specific attack has an effect on the reaction rate, but does not allow to quantify the effect of an attack and select the most effective one. Looking at the process flowsheet, the only location where the attacker would be able to determine the exact amount of VAC produced is the decanter exit. However, this number would be available to the attacker only after hours, at the end of the refinement phase. This may not be a satisfactory option. Moreover, this would require the attacker to exploit additional devices. In our analysis we have not found a way to meet this challenge.




Figure 6.3: Comparison of indirect and direct measurements

In the real world such a challenge would force the attacker to look into controller code or search for process models in the test plant. In our case, we decided to look into the model code. In particular we were interested in the state variables in the reactor unit, used in the internal computations of the process model. We were able to locate “concentration” variables of seven chemical components

in ten sections of the reactor. Through extensive analysis of the obtained data and familiarizing ourselves with the principle of operation of the plug flow reactor employed in the VAC plant we could determine that the concentration of chemicals in the tenth section would be the same or about the same as in the reactor outflow. We still did not have the unit measure of those numbers, which did not sum up to one/hundred or to the total flow. After further investigations we could conclude that molar concentrations of chemical components can be computed according to the formula:

$$MOL_{comp}(t) = \frac{CONC_{comp}(t)}{\sum_{comp} CONC_{comp}(t)},$$

where $CONC_{comp}(t)$ are concentrations of the individual chemical components.

We have verified the obtained numbers with those provided in Table 5 in [29]. Since the total reactor exit flow is directly measured in the plant, we could compute the amount of vinyl acetate produced:

$$Outflow_{comp}(t) = MOL_{comp}(t)[\%mol] \times F_{react}[Kmol/min],$$

where F_{react} is measured reactor outflow, XMEAS(6).

The production of vinyl acetate is shown in Fig. 6.3b. The rate of VAC production indeed coincides with the reactor outflow temperature profile. Knowing the molar production of VAC we could finally quantify production loss in dollar equivalents as:

$$Cost = VAC_{out}[Kmol] \times 86.09[g/mol] \times 0.971[$/kg],$$

where $86.09[g/mol]$ is VAC molar weight and $0.971[$/kg]$ is VAC price as given in [29]. To verify the numbers obtained we compared the amount of VAC produced in reactor over a period of time (Fig. 6.4a) with the amount of VAC leaving the factory as final organic product (Fig. 6.4b); the numbers matched.

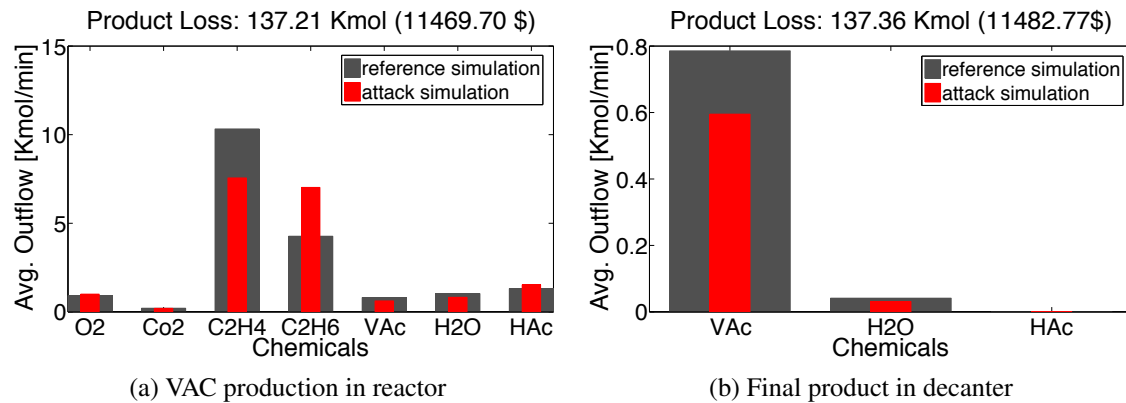


Figure 6.4: Vinyl acetate production, 24 hours

With a suitable metric for evaluating the impact of the attack we could finally start searching for effective attacks. We have established reference values of the steady-state production to determine loss (or gain) as a result of the attack. Certain attacks have caused an increase of vinyl acetate

produced. This does not necessarily mean that the overall financial gain of the plant as such increased as production may have come at the cost of increased operating costs. However, we did not pursue this investigation. Therefore we did not consider attacks causing product gain as “successful”.

Production loss	SSA	PA
High, $\geq 10.000\$$	XMV{2}	XMV{4;6}
Medium, $5.000\$ - 10.000\$$	XMV{6;7}	XMV{5;7}
Low, $2.000\$ - 5.000\$$	–	XMV{3}
Negligible, $\leq 2.000\$$	XMV{1;3}	XMV{1;2}

Table 6.3: Categorization of control loops based on damage potential

Note that attacks on XMV(1), oxygen feed, only have little attack potential, but also that this control loop easily becomes unstable. In addition, this control loop must be manipulated with great care as it quickly reaches its safety limit. Among all XMVs we mark this XMV as of least use.

Outcome of the damage stage

We conclude the damage phase with a portfolio of attacks which can be deployed at any opportune time. By scheduling attack value, attack duration and process recovery time we can control the amount of economic damage we would like to bring about.

Important considerations

Note that our analysis has only determined how much money will be lost. Ideally damage numbers should be multiplied by the chance of success so that a risky high-damage scenario can be compared with a low-risk low-damage scenarios. Precise risk metrics may never be available, but in general attacks that require manipulating more components are considered a higher risk. Attacks that require an engineering answer are riskier than attacks only requiring a technician answer. Finally attacks that must hit a particular measure value or fail are riskier than attacks that simply get more effective the closer they are to the optimal value.

6.5 Cleanup

In our attack scenario we were aiming at attacks that do not move the process towards unsafe conditions. Since we are not causing alarms, the operator may not notice immediately that the process has drifted from the economically optimal operating state. However, process operators may get concerned after noticing a persistent decrease in VAC production and may try to fix the problem. There can be numerous reasons for a process upset and operators are used to them. In this section we discuss how to influence the operator’s believe about what is happening with the process.

Having a human operator in the control loop (Fig. 6.5) turns process control system from pure cyber-physical system into socio-technical system (STS). To take advantage of the operator’s “vulnerabilities” the attacker needs to understand the specifics of the operator’s job and act according to identified weaknesses in operator’s attention, judgment process or standard procedures he has to follow.




Figure 6.5: Socio-technical system

When a plant starts losing money, investigations into that loss of revenue are likely. It may further be impossible to simply erase real-world effects. In this case the attacker needs to convince the investigators that the loss of revenue is due to some cause other than a hacker in the system.

This can be done in a number of ways. For instance, the investigator may be persuaded that the disturbance arose due to operator error. If losses are timed so that they fall in a particular employee's shift every time, one might end up investigating the employee rather than the process.

The investigator may be persuaded that the change is environmental in nature. The attack pattern can be made to increase on rainy days or hot days. Another option is to persuade the investigator that the loss is due to equipment failure. Most chemical plants are harsh environments and components fail regularly. If an attacker either kills a controller through cyber means or waits for a component to fail naturally before starting an attack cycle, the investigator may believe that the problems are the fault of a suboptimal tuning of the new components instead of a cyber attack.

Overall, events that can be used as a decoy for an attack can be grouped as:

- Change in operating conditions: change of set point, change of raw material supplier, new equipment, etc.
- Maintenance work: scheduled or unscheduled.
- Specific events: change in weather conditions, particular operator on duty, etc.

If multiple attacks are chained together they can be rotated when the attacker observes a change in the system. An industrial process, just like software, has to be debugged when it malfunctions. If the attacker changes her attacks based on the debugging efforts of the chemical engineers, future attacks may be attributed to the efforts of the engineer rather than a cyber attacker. Fig. 6.6 illustrates four different attacks which show themselves as symmetric fluctuations of different amplitudes in reactor exit temperature.

It is not possible to see directly into the reactor. Investigators will apply specific metrics allowing them to evaluate the chemical processes in the reactor and hopefully determine potential reasons for the deterioration of reactor efficiency. They will then schedule maintenance work related to the causes identified. For this reason the attacker may keep "playing" different attacks having the same effect on specific chemical processes in the reactor making engineers believe that their maintenance efforts are not bringing the expected results. Typical examples of such metrics would be selectivity and conversion rate. Selectivity is a metric to control catalyst activity. Catalyst selectivity determines the fraction of the ethylene consumed that makes the desired VAC product or in other words how




Figure 6.6: Increasing variation of reactor exit temperature caused by attacks on different control loops

much (in percent) of the primary reaction has been induced by the catalyst:

$$SEL(t) = \frac{VAC_{out}(t)}{VAC_{out}(t) + 0.5 \times CO2_{out}(t)} \times 100,$$

where $VAC_{out}(t)$ and $CO2_{out}(t)$ are molar flows of the respective chemical components in the reactor outflow.

Conversion determines the fraction of the chemicals consumed (converted into product and byproducts) during the reaction. This metric is informative in several ways. For instance there are certain safety limits and best-practice conversion rates, which should not be exceeded. Thus, reduced conversion rate of acetic acid and increased conversion of ethylene suggest an increase of the amount of the undesired secondary reaction. Conversion is computed as

$$CONV_{comp}(t) = \frac{COMP_{in}(t) - COMP_{out}(t)}{COMP_{in}(t)} \times 100,$$

where $COMP_{in}(t)$ and $COMP_{out}(t)$ are molar masses of the chemical components in the reactor in- and outflows.

In addition we introduced a metric to measure reactor efficiency. It computes how much molar mass of acetic acid has reacted, and compares this value to the amount of reacted ethylene. Since the reaction ratio of ethylene and acetic acid in the primary reaction is 1:1, the amount of reacted acetic acid is equal to the amount of correctly reacted ethylene. Relating this value to the amount of total reacted ethylene indicates the percentage of the primary reaction. Efficiency allows similar conclusions as selectivity, however, it is calculated based on the converted reagents rather than on the produced products:

$$EFF(t) = \frac{HAC_{in}(t) - HAC_{out}(t)}{C2H4_{in}(t) - C2H4_{out}(t)} \times 100,$$

Fig. 6.7 illustrates processes in the reactor during the attack on XMV(2). We have decreased ethylene feed at time $t = 120$ minutes. One can see how the attack affects the ratio between primary and secondary (ethylene combustion) reaction: the percentage of the primary reaction drops from 87% to under 82% and the amount of the secondary reaction increases by 4.32% on average

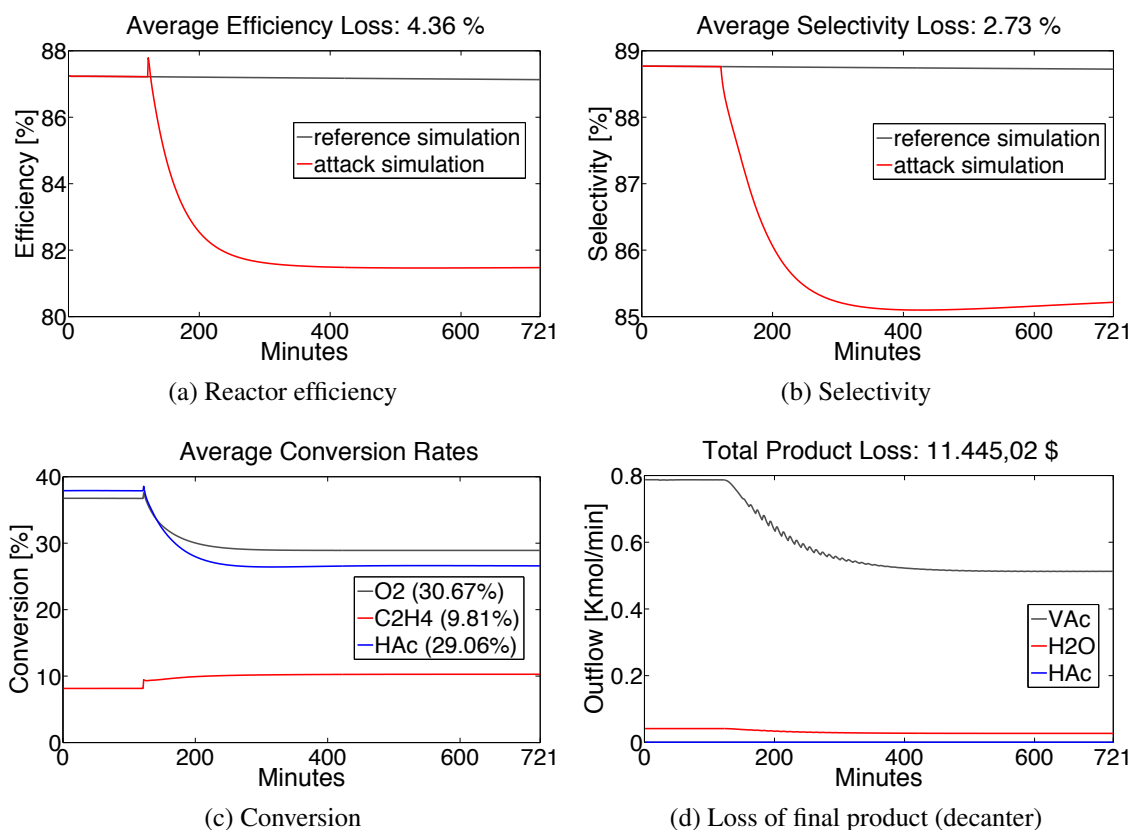


Figure 6.7: Analysis of the physical processes in the reactor

(Fig. 6.7a). Selectivity has also dropped to a lower level (Fig. 6.7b). Since selectivity is calculated based on the ethylene consumed in both primary and secondary reactions, we can conclude that an increase of the secondary reaction has a stronger effect on the reagents consumed than it has on the products produced. In other words, the consumption of reagents is the more expressive metric in this case. Fig. 6.7c plots the conversion rates for the main reagents in the reactor. Ideally, the conversion rate of acetic acid is $\approx 37\%$, around 2% higher than the oxygen conversion. However, due to the attack, the conversion rate of acetic acid drops beneath oxygen conversion, indicating that the reaction kinetics have changed. This is because the newly induced secondary reaction also consumes oxygen (even more than the primary reaction). At the same time the ethylene consumption has increased. Therefore we can conclude that the amount of the primary reaction has decreased (less acetic acid is converted), and the amount of ethylene combustion has increased. The result of the attack on XMV(2) is a significant reduction in production of the final product (Fig. 6.7d)

6.6 Discussion

Our initial attacker model restricts the attacker from crossing any operational constraints. In reality the attacker may suppress alarms while supplying the operator with good process values, e.g. using techniques proposed in our other work [21, 22].

We did check whether it might be beneficial for the attacker to violate operational constraints to cause more damage. Whereas we could almost double the loss in case of steady state attacks, the increase of amount of damage in periodic attacks was modest. Note, that in certain cases violation of operational alarms eventually moves the process into an unsafe state and triggers process shutdown via interlocks. Suppressing safety alarms and interlock triggers requires additional and more advanced hacking.

An attacker is not restricted to a single scenario. Many attacks are not mutually exclusive. It may be possible to attack one part of the process and then later attack another part based on the success of the first. In hacker circles this is referred to as getting “multiple bites at the apple”. An attacker may never have full knowledge of the process and the environment. It may also be impossible for the attacker to fully test her code before deploying it. Chaining together multiple attacks into a single payload maximizes the chance that one of them will have the desired effect.



7. Damn Vulnerable Chemical Process

One of the challenges of cyber-physical security research is the lack of large-scale test beds for studying complex attacks and their effects on physical processes. Building such a testbed requires not only significant financial investment and a specialized personnel but also regulatory permissions to conduct unsafe experiments. Moreover it would involve additional expenses to re-build the testbed if it is broken it as a result of successful experiment. Out of this considerations a more affordable way of conducting cyber-physical experimentation is to use simulation models.

In 1993 Down and Vogel published a model of an industrial chemical process (Tennessee Eastmann process – TE) for the purpose of developing, studying and evaluating process control technologies [10]. For many years their model has proved to be beneficial to the process control community serving as a realistic check on the the industrial and practical relevance of novel process control solutions. Many publications have appeared about the TE process which became one the most popular research process models. Not surprisingly, also the ICS security community has adopted the TE model for studying the impacts of cyber attacks on physical processes (see e.g. [18, 32, 12, 3, 4]).

One disadvantage of the TE model are the missing details about chemical reactions and the equipment involved. The dynamic behavior of the plant was provided in the form of a simulation in Fortran code accompanied by a flowsheet, a steady-state material balance, and a qualitative description of the key process characteristics. As a result this implementation does not allow studying *specific* research questions related to the modeled process. To address the continued interest among researchers to have additional industrial benchmark processes, Luyben and Tyreus published details of an industrial process for the production of Vinyl Acetate Monomer (VAC) in 1997 [29]. This model involves real non-ideal chemical components, a realistically large process flow sheet and consists of several standard unit operations that are typical of many chemical plants with the recycles stream and energy integration. With that the VAC process model goes a step beyond the TE process model.

What makes chemical plants excellent case studies?

Chemical plants are large physical processes with very complex non-linear interdependencies. Their models include disturbances with wide spectrum of dynamic behaviors. Also plant models include simulation of controllers. Chemical plant models have challenging objective functions to maintain: optimal production, safety and minimize production costs. Those functions also allows to numerically evaluate the success of attack from several perspectives.

7.1 Framework description

Damn Vulnerable Chemical Process (DVCP) is an open-source framework developed for cyber-physical security experimentation based on two above mentioned models of chemical plants [19, 20].¹ DVCP allows to study what it takes to convert a cyber attack in to successful cyber-physical attack. The frameworks are useful for working on individual attack instances and complex attacks. Consequently the "defensors" can study process resilience to cyber attacks to develop risk assessment, robust control, attack detecting techniques, etc. The framework can be used standalone or as a physical layer of distributed industrial control systems infrastructure as it is e.g. done at NIST [**nist:testbed**]. Below we summarize our contribution to DVCP framework.

Tennessee Eastman Process

In our framework we use TE Matlab model developed by Ricker [34]. We redesigned the initial model of the TE process to include simulation of data integrity and DoS attacks on sensor signals and the manipulated variables (actuators) as shown in Fig. 7.3. Simulink model is multilayered and care must be taken to ensure correct implementation and propagation of attack parameters through the layers to individual controllers. A functional implementation of an integrity attack on an individual controller at the lowest layer is illustrated in Figure 7.2. Attack data are stored in the work space for further the analysis and visualization. We enhanced the Simulink model with a user interface which allows to set attacks with few mouse clicks. (Fig. 7.3). The user can select attack value, attack time and its duration (predefined or random) as well frequency of attacks (single or periodic). Several attacks can be chained together. Besides, we enabled selection of the sampling frequency of process data.

Original TE model does not allow any randomness in the simulations to guarantee the repeatability of the plant operation disturbances. It means that each simulation run produced identical results. It was a significant disadvantage of TE model as it was not possible to statistically evaluate impact of the attack strategies. We modified the original code by generating a new seed for the random number generator for each simulation run while preserving underlying dynamics of process behavior. It is now possible to switch off randomness or set a specific seed.

Vinyl Acetate Process

The authors of the Vinyl Acetate plant process built a rigorous nonlinear dynamic model of the process to verify the feasibility of the simulation the plant. Details on the assumptions and details of

¹Authors: Marina Krotofil and Alexander Isakov

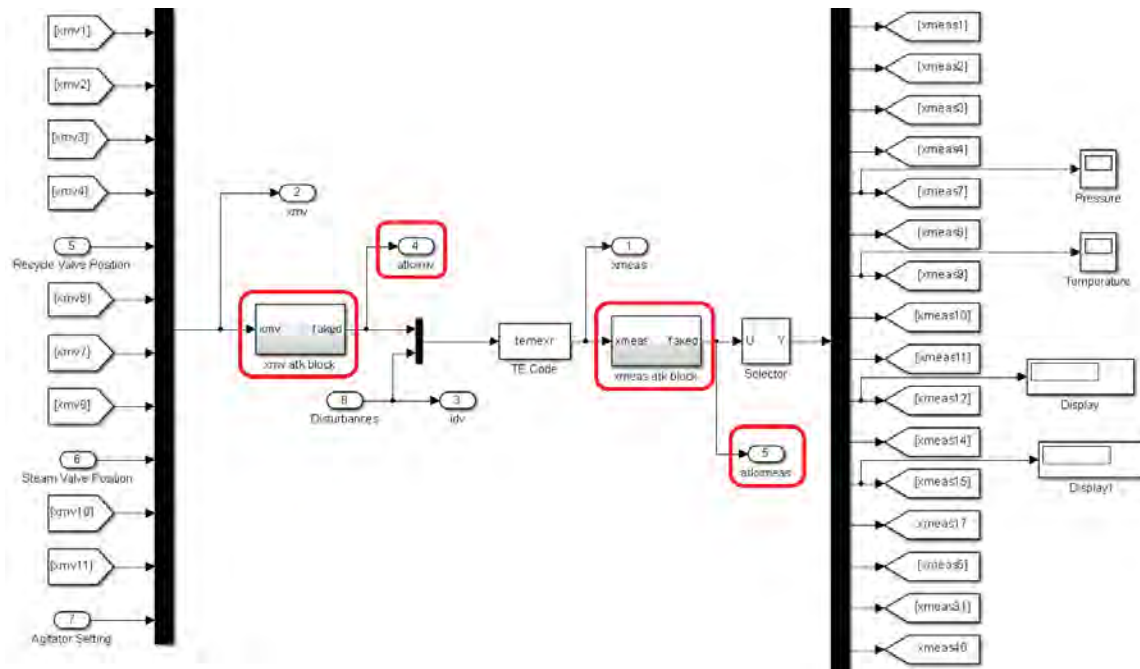


Figure 7.1: Simulink model enhanced with attack blocks

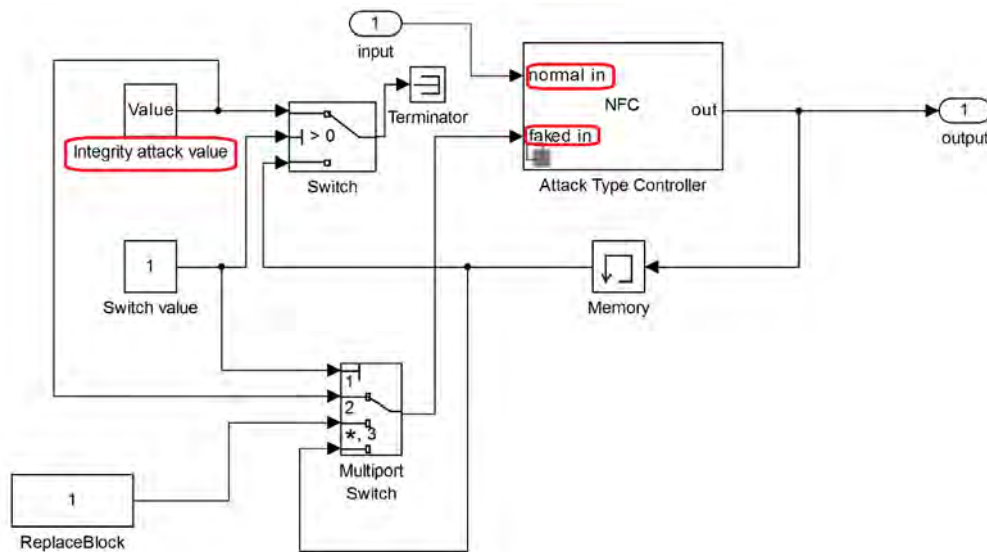


Figure 7.2: Functional implementation of attack modeling in Simulink

the modeling are described in Section 5 in [29]. The simulation model was implemented in TMODES, DuPonts in-house dynamic simulation environment, and thus, is not available for public use. In several academic works on vinyl acetate process models are implemented in specialized commercial simulation tools such as HYSYS [8], Visual Modeler [35], Aspen Plus [35] and others. To make process model available for a wider range of users Luyben et. al. have developed simulation model

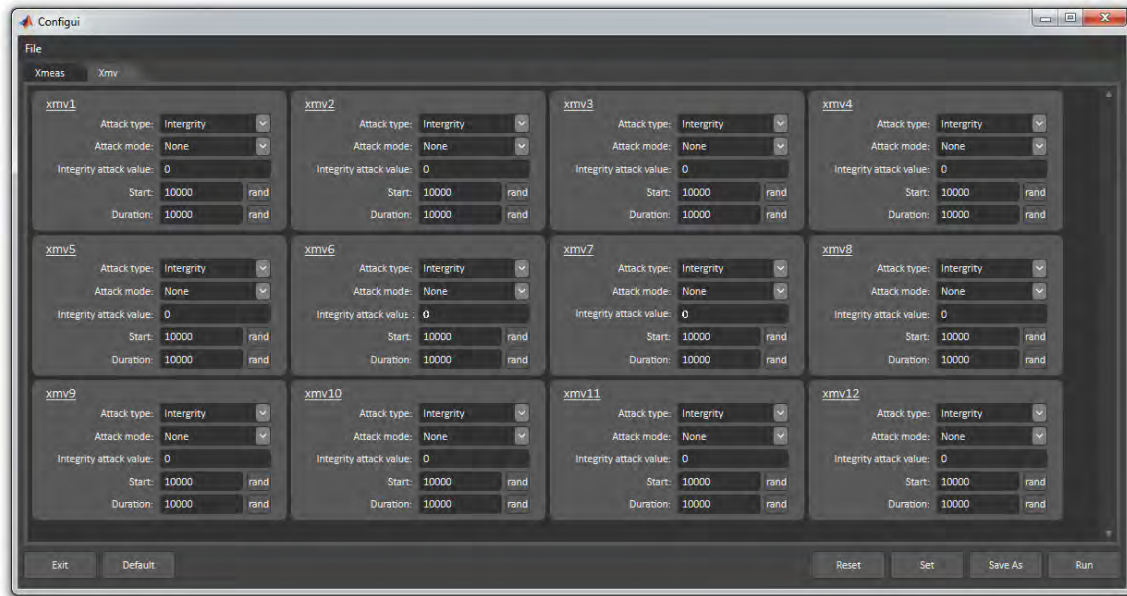


Figure 7.3: GUI for launching attacks

of VAM for Matlab[6, 5]. Both the steady state and dynamic behavior of the Matlab model were designed to be close to the TMODS model.

Originally, process equations had been modeled in Matlab and then translated into C-routines. The C coded is compiled into “MEX functions” and can be called within the Matlab environment. A separate m-file is responsible for the control of the VAC process (scheduling of the C-routines) with a developed multiloop SISO architecture. Additional four Matlab routines were developed for plotting the results of the simulations. No simulation data were output to the workspace for further analysis. The initial model did not provide any interface to process code and any manipulations of the model inputs (e.g. setpoint or controller update) must have been carried out directly in the C code requiring its re-compilation.

Considering the number of variables in the complex VAM process and the inconvenience of manipulating the process within the source code, we have developed a Simulink model of the process.² Simulink provides an interactive, graphical environment for modeling, simulating, and analyzing of dynamic systems at any level of details. Simulink models are compact enough to be understood with moderate effort. The interactive nature of Simulink allows easy experimenting by changing the model and its parameters and immediately observing what happens. Thus, simulation of attacks on a selected model component can be achieved very fast by adding a function block with several lines of code. Such functionality suits well the “what if” nature of cyber-physical exploitation.

To allow users to benefit from the experience curve we intentionally designed a Simulink Model for VAC process similar to one used in the TE model.³ We instrumented the Simulink model with a user

²Initially we have developed user interface and attack codes without building a Simulink model. Several months of experimentation have revealed limitations of such approach.

³We also named manipulated variables as XMV and controlled variables as XMEAS maintain consistent notation between TE and VAM Simulink models.

interface for convenient update of process parameters and setting up attacks on individual components (Figure. 7.4). We also implemented an output of the simulated data to the workspace for further analysis and enabled their automatic visualization (Figure. 7.5). We fixed several implementation mistakes in process code and also made several improvements to its control model to make process more stable.

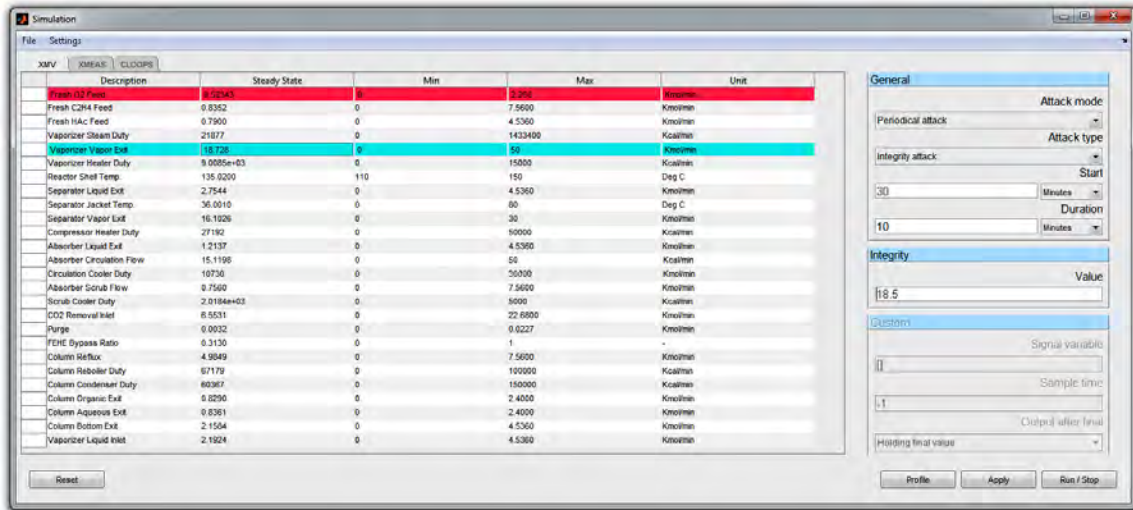


Figure 7.4: User interface for setting up attacks




Figure 7.5: Visualization of simulation results

The biggest advantage of the VAC model is that it is "built" out of realistic "components" and the availability of the source code what allows to modify original model for the use in a wide variety of process control research areas (e.g. change model of the catalyst).

7.2 Comparison of TE and VAM processes

This section gives a short comparison between Tennessee Eastman and Vinyl Acetate process models. Involvement of the real non-ideal chemical components and the size of the process (Tbl. 7.1) make Vinyl Acetate a more challenging process control problem.

	TE	VAM
# sections	5	10
# states	50	246
# measurements	41	43
# MV	12	26
# modes	6	1
# setpoint changes	4	2
# disturbance modes	20	5
Process complexity	Medium	High

Table 7.1: Complexity of the process

Availability of literature on the VAC process specifics (Tbl. 7.2) allows to design targeted attacks. In contrast, absence of the *a priori* information about the TE process makes it an excellent case study for emulation of "grey-box" exploitation use cases. In addition, TE model includes a wide choice of the disturbances with both plantwide and local effects what can make exploitation more challenging.

	TE	VAM
Chemicals and reactions are specified	-	+
Equipment is specified	-	+
Safety constraints	+	+
Operating constraints	+	+
Operating cost function	+	-
Sensor signal noise	+	-
Process randomness	± ⁴	-
Predefined process code	+	-
Literature on the process	± ⁵	


Table 7.2: Specification of the processes

⁴Enabled by us

⁵Since the exact chemistry of TE process is unknown, there is no literature on the process details. However, there is a large number of works on the topic related to TE process control, its safety and security. In contrast, VAM is a real process. There is a large body of literature on the specifics of the individual unit operations and chemical reactions.

In [30] authors specify a number of complexities and interesting dynamic effects of the VAM process. Vinyl Acetate process quickly becomes unstable if pushed outside of the steady state conditions. In contrast, controlling TE process is rather straightforward and the process is more robust to malicious manipulations. Both processes have control loops with short and large time constants and exhibit fast and slow dynamics depending on the input parameters of the process (change in the operating conditions). Due to the differences in process specifications, TE and VAC are suitable for studying *distinct* types of aspects related to cyber-physical exploitation.

Damn Vulnerable Chemical Process: official logo





8. Conclusion

In cyber-physical systems security we are concerned specifically with attacks that cause physical impact. To do so, the attacker has to find ways of manipulating the physical processes in the system. Cyber attacks on process networks may allow the attacker to obtain sensor readings, to manipulate sensor measurements sent to controllers and instructions sent to actuators. To appreciate the effect of such manipulations the attacker has to understand the physical part of her target. The attacker may be impeded by automatic safety measures and may not have access to observations that allow her to monitor the effect of her actions.

The construction of a successful attack has to go through several stages, some can be performed in parallel, some will be performed repeatedly, and some will require expertise on the physical part of the cyber-physical system, an expertise not commonly found in the IT security community.

We have demonstrated this approach in the example of a simulation of a vinyl acetate monomer plant to give some glimpses on the detours an attacker may have to take to reach her goal. Studying the hurdles the attacker has to overcome allows to understand what she needs to do and why. This knowledge useful for eliminating low hanging fruits and making exploitation harder. Analyzing processes when maliciously manipulated enables process operators to discover the weaknesses of a process design in the presence of cyber attacks. The defenders in turn gain insights which additional controls might increase the resilience of physical processes to cyber assaults.



References

- [1] Dmitri Alperovitch. *Revealed: Operation Shady RAT*. Tech. rep. McAfee, 2011.
- [2] Dillon Beresford. “Exploiting Siemens Simatic S7 PLCs”. In: *Black Hat USA (2011)*. https://media.blackhat.com/bh-us-11/Beresford/BH_US11_Beresford_S7_PLCs_Slides.pdf.
- [3] Alvaro A. Cárdenas et al. “Attacks against process control systems: risk assessment, detection, and response”. In: *AsiaCCS '11*. 2011, pp. 355–366.
- [4] Rohan Chabukswar et al. “Simulation of Network Attacks on SCADA Systems”. In: *First Workshop on Secure Control Systems*. 2010.
- [5] Rong Chen et al. *A Nonlinear Dynamic Model of a Vinyl Acetate Process*. <http://www.isr.umd.edu/~mcavoy/VAC%20Material/>. retrieved: June, 2013.
- [6] Rong Chen et al. “A Nonlinear Dynamic Model of a Vinyl Acetate Process”. In: *Industrial & Engineering Chemistry Research* 42.20 (2003), pp. 4478–4487.
- [7] Eric Chien and Gavin O’Gorman. *The Nitro Attacks. Stealing Secrets from the Chemical Industry*. Tech. rep. Symantec, 2011.
- [8] J.P. Contreras et al. “Vinyl Acetate from Ethylene, Acetic Acid and Oxygen Industrial Plant Simulation”. In: *Proceedings of the Computing and Systems Technology Division, American Institute of Chemical Engineers (AIChE) Annual Meeting*. AIChE. 2009, pp. 249–259.
- [9] Alexandre C. Dimian and Costin Sorin Bildean. *Chemical Process Design: Computer-Aided Case Studies*. WILEY-VCH Verlag GmbH & Co., 2008.
- [10] J. J. Downs and E. F. Vogel. “A plant-wide industrial process control problem”. In: *Computers & Chemical Engineering* 17.3 (1993), pp. 245–255.
- [11] Zoran Gajic and M. Lelic. *Modern Control System Engineering*. Prentice Hall, 1996.
- [12] Béla Genge and Christos Siaterlis. “An Experimental Study on the Impact of Network Segmentation to the Resilience of Physical Processes”. In: vol. 7289. LNCS. 2012, pp. 121–134.

- [13] *Havex Hunts For ICS/SCADA Systems*. <https://www.f-secure.com/weblog/archives/00002718.html>.
- [14] *ICS-CERT Advisories*. <https://ics-cert.us-cert.gov/advisories>.
- [15] “Kinetics of ethylene combustion in the synthesis of vinyl acetate over a Pd/SiO₂ catalyst”. In: *Journal of Catalysis* 224.1 (2004), pp. 60–68.
- [16] Brian Krebs. “Chinese Hackers Blamed for Intrusion at Energy Industry Giant Telvent”. In: *KrebsonSecurity* (2012). <http://krebsonsecurity.com/tag/telvent-hack/>.
- [17] Brian Krebs. “Security Firm Bit9 Hacked, Used to Spread Malware”. In: *KrebsonSecurity* (2013). <http://krebsonsecurity.com/2013/02/security-firm-bit9-hacked-used-to-spread-malware/>.
- [18] Marina Krotofil and Alvaro A. Cárdenas. “Resilience of Process Control Systems to Cyber-Physical Attacks”. In: vol. 8208. LNCS. 2013, pp. 166–182.
- [19] Marina Krotofil and Alexander Isakov. *Damn Vulnerable Chemical Process - Tennessee Eastman Process*. <http://github.com/satejnik/DVCP-TE>.
- [20] Marina Krotofil and Alexander Isakov. *Damn Vulnerable Chemical Process - Vinyl Acetate Monomer*. <http://github.com/satejnik/DVCP-VAM>.
- [21] Marina Krotofil and Jason Larsen. “What You Always Wanted and Now Can: Hacking Chemical Processes”. In: *Hack in the Box Amsterdam* (2015). <http://conference.hitb.org/hitbsecconf2015ams/wp-content/uploads/2015/02/D2T1-Marina-Krotofil-and-Jason-Larsen-Hacking-Chemical-Processes.pdf>.
- [22] Marina Krotofil, Jason Larsen, and Dieter Gollmann. “The Process Matters: Ensuring Data Veracity in Cyber-physical Systems”. In: *Proceedings AsiaCCS’15*. 2015.
- [23] Marina Krotofil et al. “Vulnerabilities of cyber-physical systems to stale data—Determining the optimal time to launch attacks”. In: *International Journal of Critical Infrastructure Protection* 7.4 (2014), pp. 213–232.
- [24] Ralf Langner. *To kill a centrifuge*. Tech. rep. Langner Communications, 2013.
- [25] Jason Larsen. “Breakage”. In: *Black Hat Federal* (2007). <https://www.blackhat.com/presentations/bh-dc-08/Larsen/Presentation/bh-dc-08-larsen.pdf>.
- [26] Jason Larsen. “Miniaturization”. In: *Black Hat USA* (2014). <http://http://blackhat.com/docs/us-14/materials/us-14-Larsen-Miniturization.pdf>.
- [27] Robert A. Leishear. *Fluid Mechanics, Water Hammer, Dynamic Stresses, and Piping Design*. ASME, 2013.
- [28] Éireann Leverett and Reid Wightman. “Vulnerability Inheritance Programmable Logic Controllers”. In: *The 2nd International Symposium on Research in Grey-Hat Hacking (GreHack)* (2013).
- [29] Michael L. Luyben and Björn D. Tyréus. “An industrial design/control study for the vinyl acetate monomer process”. In: *Computers & Chemical Engineering* 22.7–8 (1998), pp. 867–877.
- [30] William L. Luyben, Bjorn D. Tyreus, and Michael L. Luyben. *Plantwide Process Control*. McGraw-Hill, 1998.

-
- [31] McAfee et al. *Global Energy Cyberattacks: "Night Dragon"*. Tech. rep. McAfee, 2011.
- [32] Thomas McEvoy and Stephen Wolthusen. "A Plant-Wide Industrial Process Control Security Problem". In: *CIP V*. Vol. 367. 2011, pp. 47–56.
- [33] D. G. Olsen, W. Y. Svrcek, and B.R. Young. "Plantwide control study of a vinyl acetate monomer process design". In: *Chemical Engineering Communication* 192.10 (2005), pp. 1243–1257.
- [34] N. Lawrence Ricker. *Tennessee Eastman Challenge Archive*. <http://depts.washington.edu/control/LARRY/TE/download.html>. retrieved: May, 2013.
- [35] Hiroya Seki et al. "Plantwide control system design of the benchmark vinyl acetate monomer production plant". In: *Computers & Chemical Engineering* 34.8 (2010), pp. 1282–1295.
- [36] Jacques F. Smuts. *Process Control for Practitioners*. OptiControls Inc, 2011.
- [37] *Targeted Cyberattacks Logbook*. <https://apt.securelist.com>.
- [38] The Dow Chemical Company. *Product Safety Assessment: Vinyl Acetate*. <http://www.dow.com/productsafety/finder/vinyl.htm>. retrieved: June, 2014.
- [39] M. Zeller. "Myth or reality - Does the Aurora vulnerability pose a risk to my generator?" In: *Protective Relay Engineers, 2011 64th Annual Conference for*. 2011, pp. 130–136.
- [40] Isaac A. Zlochower and Gregory M. Green. "The limiting oxygen concentration and flammability limits of gases and gas mixtures". In: *Journal of Loss Prevention in the Process Industries* 22.4 (2009), pp. 499–505.