Chapter 6

DETECTING SENSOR SIGNAL MANIPULATIONS IN NON-LINEAR CHEMICAL PROCESSES

Thomas McEvoy and Stephen Wolthusen

Abstract Modern process control systems are increasingly vulnerable to subversion. Attacks that directly target production processes are difficult to detect because signature-based approaches are not well-suited to the unique requirements of process control systems. Also, anomaly detection mechanisms have difficulty coping with the non-linearity of industrial processes.

This paper focuses on the problem where attackers gain supervisory control of systems and hide their manipulations in signal noise or conceal computational states. To detect these attacks, we identify suitable proxy measurements for the output of a control system. Utilizing control laws, we compare the estimated system output using real-time numerical simulation along with the actual output to detect attacker manipulations. This approach also helps determine the intervention required to return the process to a safe state.

The approach is demonstrated using a heat exchange process as a case study. By employing an explicit control model rather than a learning system or anomaly detection approach, the minimal requirements on proxy sensors and the need for additional sensors can be characterized. This significantly improves resilience while minimizing cost.

Keywords: Process control systems, attack detection, proxy measurements

1. Introduction

Supervisory control and data acquisition (SCADA) systems are vital components in critical infrastructures. Advanced technologies have significantly improved the operation and management of these systems, but they increase the vulnerability to attack [12]. Control systems are often generic computing hosts with complete operating systems and network stacks [7] as opposed to isolated, proprietary systems. This increases the potential for manipulation

of the computational states in a SCADA network [10, 21] and process control signals at the sensors and remote terminal units (RTUs) [4].

In previous work [20], we showed that simple statistical anomaly detection can be bypassed by a knowledgeable attacker, underscoring the need for a multi-sensor approach to detection. We, therefore, proposed a novel approach that analyzes process correlations using additional sensors [15]. This paper presents a formal model of the approach, which utilizes process control laws to directly achieve the goals. A beer pasteurizer is considered as a case study because it is a simple, but realistic, example of heat exchange in a production environment.

Pasteurization involves a series of heat exchanges that are controlled to ensure specific target temperatures required for production. The relations between heat exchange inputs and outputs are captured using material and energy balance equations [1]. An attack may be defined as a "concealed" manipulation of these relations [15], which implies that a process degradation cannot be detected by conventional fault analysis. We assume that the attacker has supervisory access to the plant and can alter setpoint values or sensor values, while hiding these manipulations from plant operators [4, 21].

This paper shows how to identify suitable (composite) proxy measurements for making a determination of the current process behavior. The approach relies on the presence of non-linear relations, so that small alterations in proxy values may be correlated with significant process changes. These proxy measurements support the comparison of the actual behavior of the pasteurizer with its estimated behavior, making direct use of material and energy balances and utilizing real-time numerical simulation techniques. This enables the efficient detection of process signal inconsistencies that indicate the presence of manipulated states. In contrast, non-linearity in industrial processes renders conventional anomaly detection approaches (e.g., statistical analysis or learning systems) computationally infeasible for real-time applications.

The proxy measurements also provide a basis for implementing intervention strategies. In a practical implementation, the sensors used for proxy measurements may reside in an out-of-band network with data being pushed to the network via data diodes for detection purposes.

2. Related Work

SCADA systems are increasingly vulnerable to cyber attack due to their modernization and exposure to untrusted networks [3, 12, 16, 23]. This has led to increased interest in intrusion detection [18]. As in the case of conventional computer systems, intrusion detection research has focused on signature-based approaches (generally) at the perimeter and anomaly-based approaches that address the insider threat and direct attacks on control processes [9, 13, 24].

The predictable nature of SCADA traffic can be leveraged to detect system anomalies [5, 21, 24]. However, a knowledgeable attacker can seize the advantage by manipulating computational states or utilizing signal noise to obfuscate attacks that would otherwise be recognized [4, 20, 21].

We argue that this adversary capability highlights a requirement for additional sensors [2, 6, 20] to provide different points of view in order to detect anomalies [21]. This requirement is underscored by the introduction of sophisticated control processes that rely on multivariate controls and, hence, require more complex forms of supervision [19]. Note that this approach would also apply to traditional control systems. The approach has strong parallels with fault detection strategies in SCADA environments [22], but unlike fault detection approaches, the sensors cannot always be assumed to be reliable.

We use functional models to map systems [17] and identify suitable redundant characteristics for evaluating process behavior. We combine these readings with a process simulation to detect signal manipulation [1], extending the invariant model proposed in [15]. This approach obviates the need for linear approximations as used in explicit control models (see, e.g., Lin, *et al.* [11]).

3. Problem Definition

We assume that an attacker is capable of remote penetration attacks on a process control system and understands the industrial process under control. Furthermore, the attacker may gain unauthorized supervisory access to the system and be able to alter setpoints and sensor readings while disguising this from plant operators [4, 21]. For a complex process, which requires multivariate analysis to ensure that production values are achieved, the attacker may not be able to disguise an attack simply by manipulating signals to hide the attack in signal noise [20]. However, as discussed in the next section, it is possible to conceal the attack by falsifying a subset of control signals and relying on the behavior of other parts of the system to normalize the anomalous signals.

Our focus is to identify a conservative number of additional sensors that could provide an inexpensive, practical and computationally-feasible means of uncovering attacks in real time. It is also desirable to be able to use the detection approach as a basis for intervention, although meeting this requirement is beyond the scope of this paper.

4. Detection Model

Heat exchange is a common industrial process. In our case study, we consider heat exchange in the operation of a flash pasteurizer and model the identification of potential proxy measurements. Subsequently, we use a simulation of the pasteurizer to develop a profile of proxy behavior under different operating conditions to determine information about the state of the pasteurizer by observing the proxy behavior.

4.1 Pasteurizer Simulation

Pasteurization is achieved by a series of heat exchanges between hot and cold fluids, where the hot side setpoint temperature is determined by the flow rate and the cold side temperature by the packaging requirements. In flash pasteur-

ization, the interaction of the pasteurized (hot) and unpasteurized (cold) product is used as part of the temperature cycle with target values being achieved through secondary heat exchanges using steam-heated water on the hot side and a glycol refrigerant on the cold side [15]. The basic equations for heat exchange are:

$$\dot{q} = UA(T_{in,s} - T_{out,c}) \tag{1}$$

$$wCp\frac{dT}{dt} = \dot{w}Cp(T_{in,c} - T_{out,c}) + \dot{q} \tag{2}$$

$$uCp\frac{dT_s}{dt} = \dot{u}CP(T_{in,s} - T_{out,s}) - \dot{q} \tag{3}$$

where $\dot{q}$ is the rate of heat exchange; U is the coefficient of heat exchange for the construction material; A is the heat exchange area; $T_{in,s}$ is the initial hot side temperature; $T_{out,s}$ is the final hot side temperature; $T_{in,c}$ is the initial product temperature; and $T_{out,c}$ is the final product temperature. Equations (2) and (3) represent the respective energy balances of the cold and hot sides where $\dot{w}$ and $\dot{u}$ represent the hot side and cold side flow rates, respectively; w and u are the corresponding liquid volumes; and Cp is the specific heat capacity of the product.

We use a diagrammatic form of these equations in Matlab/Simulink (Figure 1(a)), which we subsequently link together in larger blocks to simulate the main plate heat exchanger called the splitter/regeneration unit. This unit uses the hot and cold beer flows as a primary counterflow heat exchange mechanism (Figure 1(b)).

The secondary heating and cooling actions of steam-heated water and glycol refrigerant at the hot and cold sides may be modeled in separate heat exchange sections using a proportional-integral-derivative (PID) block to mimic valve action. The cooling section is presented in Figure 1(c).

The pasteurization unit (PU) is derived from the flow rate and temperature values using the equation:

$$PU = \frac{w}{\dot{w}}(60)1.393^{(T-60)} \tag{4}$$

where w is the holding volume; $\dot{w}$ is the flow rate; and T is the temperature. This is an oddly dimensioned measure, which was derived by Dayharsh, *et al.* [8]. Note the non-linear relationship between flow rates, temperature and pasteurization values.

The flow rate $\dot{w}$ is given by:

$$\dot{w} = \dot{w}_{max} - \dot{w}_{min}\frac{L_{actual} - L_{min}}{L_{max} - L_{min}} \tag{5}$$

where L is the tank level; and $\dot{w}$ is the flow rate as before. The minimum and maximum tank levels for this specimen are 100 and 220 (cm), respectively; and the minimum and maximum flow rates are 250 and 500 (l/hr), respectively. Flow rates are clamped to their extrema when tank level values exceed their

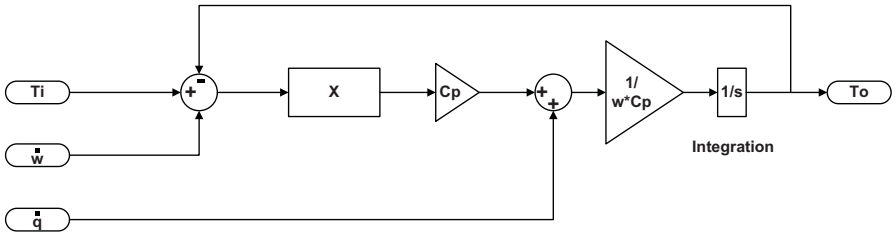

(a) Block diagram (Equation (1)).

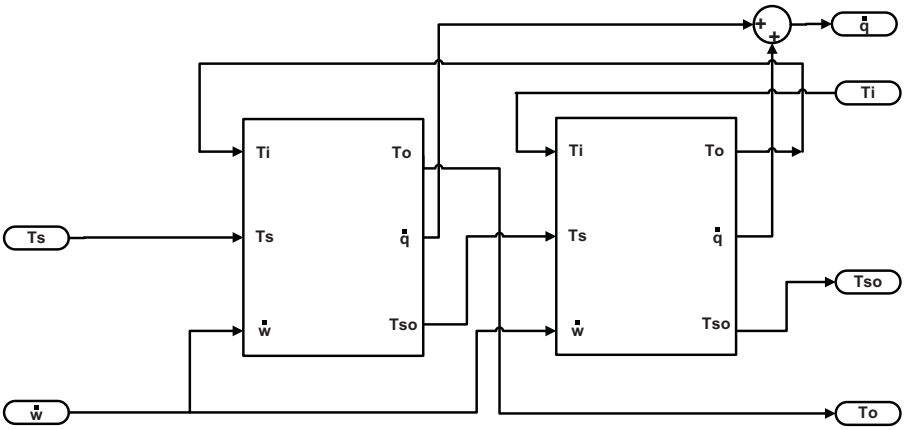

(b) Counterflow heat exchange (splitter/regeneration unit).

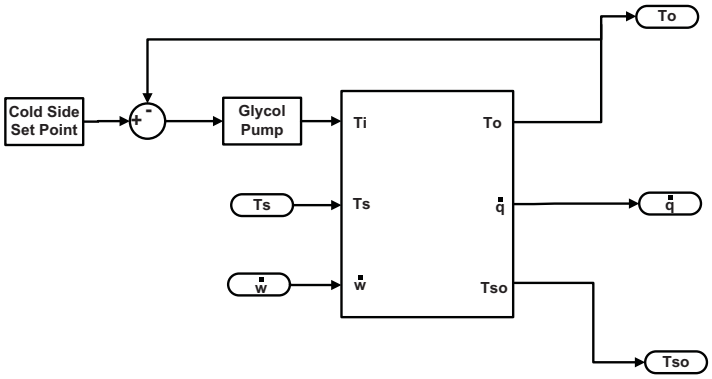

(c) Cooling section.

Figure 1. Heat exchange process system.

minimum or maximum values. The tank levels change almost continuously during pasteurization due to alterations in packaging (called "kegging") rates. We add the appropriate calculations to the model to simulate these setpoints.

4.2 Proxy Discovery

We use an adapted functional causal model of the pasteurizer to visualize significant relations and sensor values under specific attack conditions [17] in order to uncover potential proxy nodes. This is a technique that we have found to be suitable for analyzing small-scale configurations. An algebraic transform of the same technique may be used for large-scale configurations.

Let $\overrightarrow{G} = V(\overrightarrow{E})$ be a graph. Let each $v \in V$ represent a characteristic aspect of pasteurizer functionality (e.g., water temperature). Let each $\overrightarrow{e} \in \overrightarrow{E}$ be a causal relationship between distinct pasteurizer characteristics. We assume that conditional questions may be asked about the state of a node $v \in V$ where it is directly or transitively linked to a node $u \in V$, except where v is also linked to a dominant node w. Where a dominant node exists, its state determines the value of all slave (or invariant) nodes that are tail adjacent. All other relations to invariant nodes are represented as dotted lines. We assume, but do not explicitly show in graph form, that the state of each node is subject to unmeasured disturbances that create a probability distribution over node values. These values are perturbed under an attack, but the perturbation may be concealed by the attacker who manipulates computational states or uses process noise. Clearly, the attacker would attempt to conceal the values of all nodes that are directly implicated in determining the success of pasteurization. Supervisory access also allows the manipulation of certain nodes.

Figure 2 shows the pasteurizer under attack. Potentially manipulated nodes have a shaded ring, while probable concealed nodes have an unshaded ring. All the ringed nodes belong to the set of "covered" nodes whose values may not be known during an attack. Note that PU_{SP} is the pasteurization unit setpoint; TL is the tank level; FR is the flow rate; T_{SP} is the hot side temperature setpoint; S is the steam temperature; W is the water temperature; T_{in} is the initial product temperature; HX_S is the product temperature after heating in the splitter; HX_W is the product temperature after being heated by the water; HX_H is the product temperature at the end of the holding pipe and pre-regeneration; PU_{est} is the estimated PU value; HX_R is the product temperature post regeneration; T_{out} is the product temperature at the end of the process after glycol cooling; C_{SP} is the cold side setpoint; and G is the glycol temperature.

According to Equation (4), the temperature and flow rate determine the pasteurization value; and these values are central to the material and energy balance equations in the heat exchange process (Equations (1-3)). These observations suggest that it may be possible to estimate the pasteurization values from the heat exchange performance and vice versa. Finally, note that the heat

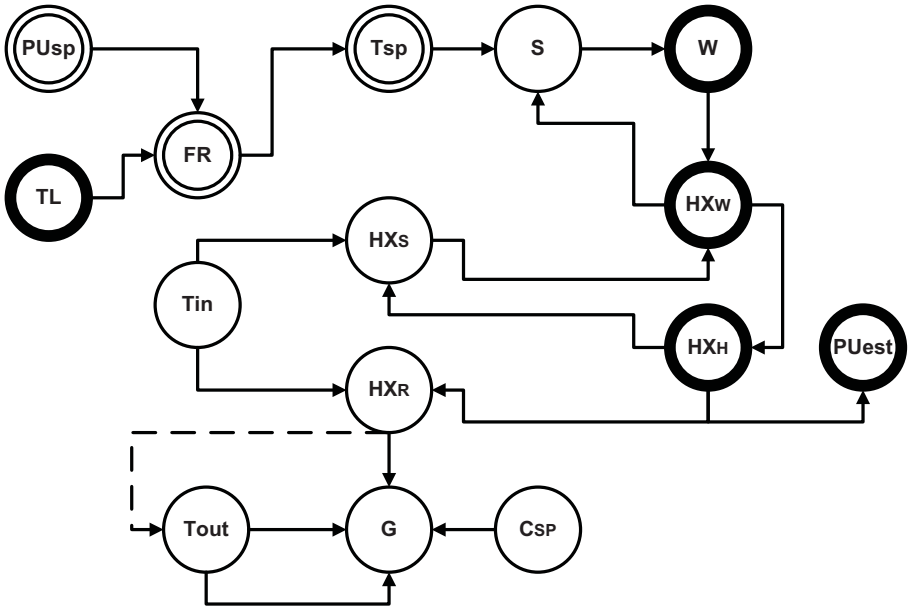

Figure 2. Adapted functional causal model.

exchange output values are uncovered at HX_S and HX_R, potentially enabling their use as proxy measurements for determining the success of pasteurization.

5. Analysis of Heat Exchange Profiles

Based on the assumptions about attacker capabilities and assuming no insider collusion, three attack strategies are possible:

- Lower the *PU* by spoofing a lower tank level, thus increasing the flow rates relative to temperature using negative error values.
- Lower the *PU* by lowering the water temperature and, hence, product temperature relative to the flow rates using positive error values, or equivalently by resetting the *PU* setpoint.
- Combine the above two strategies in a single attack.

In the first two attack strategies, the adversary has to cover all the relevant sensors so that they appear to show consistent values. In the third strategy, the attacker may omit to cover the hot side temperature as a stepped approach to jointly lower the water temperatures and raise the relative flow rates. This can be concealed in process noise, where the product temperature remains constant, but the pasteurization process is still degraded. Therefore, it is necessary to show that the proxy measurements identified in Section 4.2 can be used to estimate the flow rate and the temperature and, hence, the pasteurization unit value.

5.1 Pasteurization Profiles

As a preliminary, we show that there are distinct temperature to flow rate ratios for each PU. Upon solving Equation (4) for temperature [15], we plot the temperature values against flow rates for PU values of 40 (nominal value) and 20 (fail or "divert" value) and determine the PU value for a 1°C alteration downwards for the nominal value of $PU = 40$, which represents a significant loss of quality (i.e., we also solve for $T - 1$ with $PU = 28.7$). The results are shown in Figure 3(a).

5.2 Establishing Flow Rate

The temperature of the cooled (pasteurized) product leaving the regenerator tank can be used to estimate the flow rates. To show this, we plot this temperature for four distinct tank levels $T = 120$, 140, 160 and 180 (representing different flow rates) against the nominal pasteurization rate of 40, the divert value of 20 and the quality loss value of 28.7. The distinct banding of temperatures for the cooled product (color coded by flow rate in Figure 3(b)) shows that the flow rate and, hence, the tank level can be estimated. Confidence intervals are estimated at ± 10 l/hr. It follows that it should be possible to detect the first attack strategy that alters the flow rate by spoofing false tank level readings.

5.3 Establishing a Temperature Profile

Next, we show how a concealed alteration in temperature may be detected. Setting the tank level at TL =180 to lock in the flow rate, we present the results of three runs with distinct pasteurization profiles as before (Figure 3(c)). Given the non-linear relationship between pasteurization levels, these temperature differences are on the average sufficiently significant so that, in combination with the flow rate, it is possible to estimate the process success with a confidence level greater than 3σ. Similar results are obtained for the other flow rates. Hence, it is possible to determine the current PU level modulo approximately four units.

5.4 Uncovering Concealed Attacks

Finally, we show how transitions concealed by signal noise can be detected. We assume that an attack causes the PU value to drop gradually by lowering the water temperatures and raising the flow rates in an effort to hide the manipulations in the process signal. We set the tank level to 220 and dropped it in stages to 172, altering the flow rate upwards. We dropped the pasteurization target from 40 to 28.7 (equivalent to quality loss) by altering the water temperature error signal or PU setpoint.

Figures 4(a) and 4(b) show that the product temperatures do not vary from their mean values, while the PU rates drop significantly. Obviously, the attacker could continue this process until the divert or abort values are achieved.

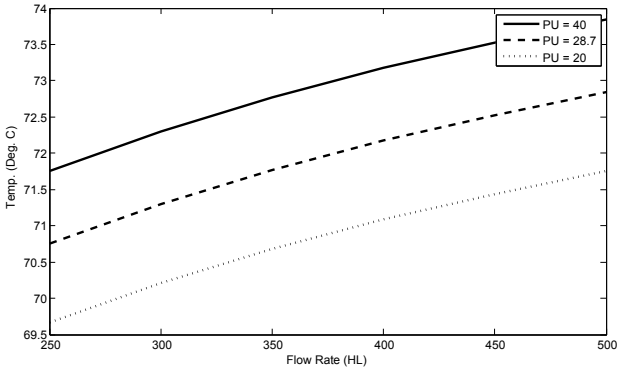

(a) *PU* profiles for flow rate vs. temperature.

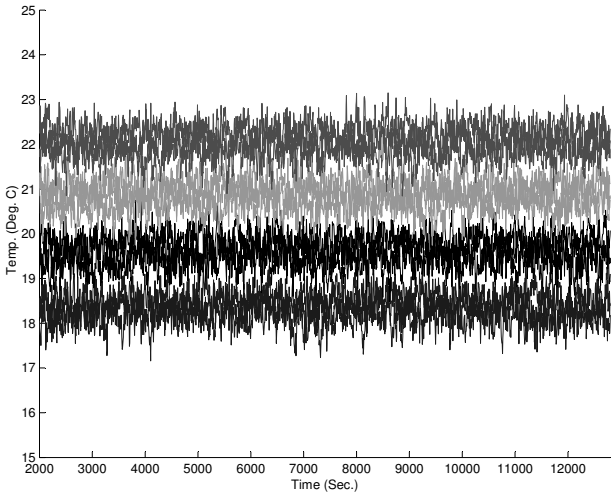

(b) Cooled product temperature as a proxy for flow rates.

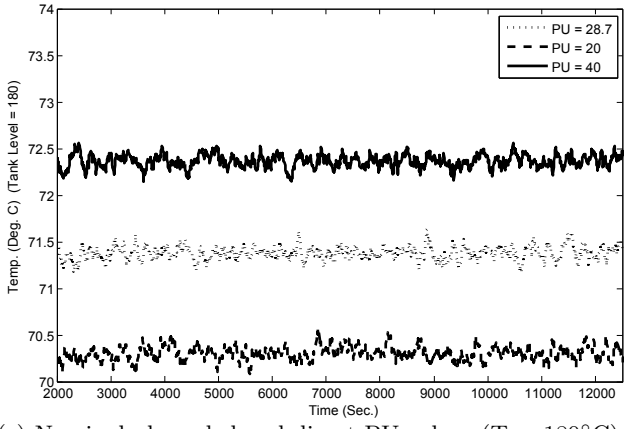

(c) Nominal, degraded and divert PU values (T = 180°C).

Figure 3. Heat exchange profiles.

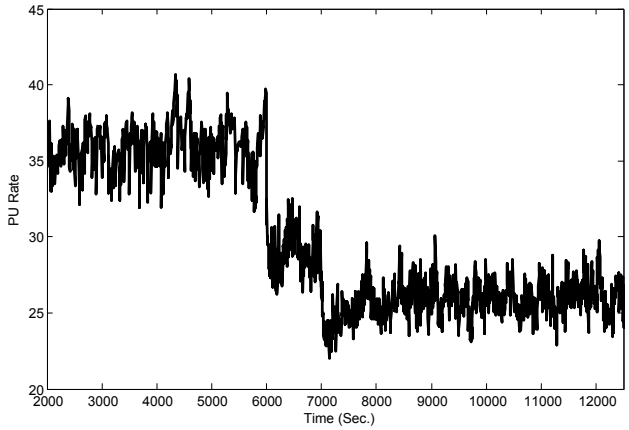

(a) Hiding in signal noise – pasteurization rate alteration.

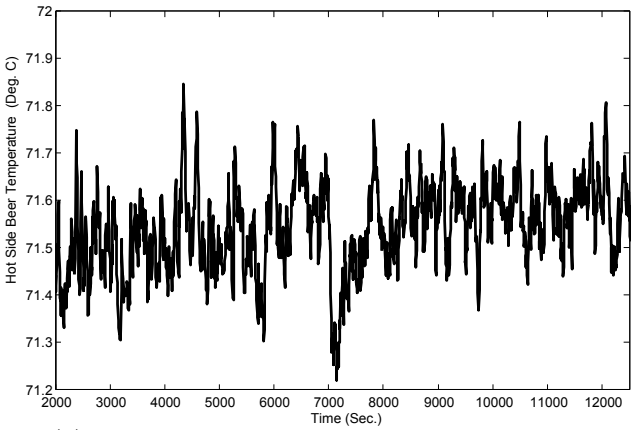

(b) Hiding in signal noise – product temperature.

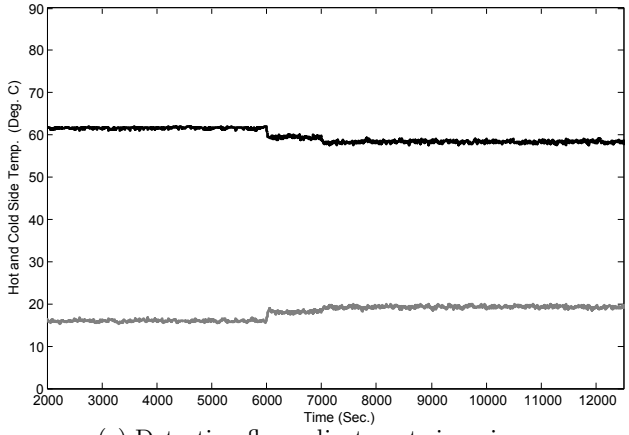

(c) Detecting flow adjustments in noise.

Figure 4. Disguising and detecting alterations in signal noise.

Figure 4(c) shows the difference in the heat exchange profile as a result of the concealed adjustments of the hot side and cold side temperatures. Note that the attack produces a greater contrast in the heat exchange profile compared with the other attacks because both the lower and upper product temperatures are altered simultaneously to create anomalous steps in the heat exchange profile. It follows that it is possible to estimate (and even control) the actual pasteurization rates using this profile. We estimate that confidence levels of $\pm 4PU$ can be achieved assuming a precision of ± 0.5°C in the hot side temperature values and ± 10 l/hr in the flow rate estimations.

6. Discussion

In Section 5, we discussed three possible attacks on the pasteurization process. The first attack lowers the water temperature setpoint so that the product is not heated to the requisite temperature to achieve the target pasteurization value. The second attack raises the flow rate by spoofing lower tank levels. The third attack combines the previous two attacks while keeping the beer temperature invariant to hide the attack manifestations in process noise while degrading the pasteurizer value.

Using well-chosen proxy measurements, it is possible (in combination with numerical simulation) to capture inconsistent and anomalous behavior in a manner that minimizes the computational cost of detection. This approach also supports attack intervention strategies. While this is not strictly relevant for a batch process such as pasteurization (which can restart if there is a fault), there are other processes involving heat exchange where the ability to adjust to attacker actions "on the fly" is necessary because the processes are not easily restarted. Intrusion prevention thus becomes a dynamic process of defending process integrity.

Clearly, the ideal situation is to have a fully redundant set of sensors, but physical and cost constraints along with certification and accreditation requirements make this approach infeasible. Our approach, therefore, seeks to minimize the effort involved in implementing a signal-based anomaly detection mechanism, which is important when dealing with large-scale industrial processes. Although this paper focuses on heat exchange in the context of pasteurization, the approach is applicable to a variety of non-linear engineering processes.

In many cases, there exists the potential to identify output values that are redundant for process control, but that can be used as proxy measurements in combination with real-time control system simulation for intrusion detection. We have previously shown that anomaly detection techniques based on univariate statistical techniques may be unable to distinguish signal noise from an attack, but in the case of a non-linear process, even a difference of σ in the average performance can radically alter the result [20]. Classical intrusion techniques, in general, do not consider process signals. Even if these approaches were to be applied, their reliance on signature-based detection has no validity in the application. Moreover, anomaly detection techniques that rely on

complex statistical analyses (e.g., Markov models) have limited applicability to non-linear systems because they do not accommodate the sharp disparities in process state that result from small alterations in such systems. Learning systems face similar problems because considerable data is required to create a training set that accommodates non-linearity. In case of the pasteurizer, a learning system would have to learn the pasteurization profile of each product that is processed by a pasteurizer, track the performance degradation and estimate state changes. In contrast, a control model of the system encapsulates these aspects in a straightforward manner and provides a computationally inexpensive numerical simulation of process behavior.

Nevertheless, our approach has certain limitations. Some limitations may be introduced by physical constraints such as sensor placement [1] that can reduce the confidence in the results. Distinct processes are associated with different perturbation levels; this can reduce (or increase) confidence levels. However, in most cases, even the process models used to set up the control systems are limited in precision and tuned based on experience rather than physical or chemical models. Thus, the approach is ultimately limited by the precision of these models.

7. Conclusions

Attacks on industrial control systems that involve signal manipulations are often invisible to traditional intrusion detection systems. A promising solution is to use proxy measurements to determine anomalous readings in key process characteristics in a computationally efficient manner while minimizing the need for additional sensors, thereby reducing the accompanying costs. This approach permits the continued safe operation of a process when shutdown is not feasible.

The primary limitations of the approach are process specific and plant specific in nature. Different processes are associated with distinct perturbations. In addition, variations in plant design may not permit the satisfactory placement of supplementary sensors. These factors result in lower confidence levels.

Our future work will attempt to characterize processes that are amenable to this approach. We will also develop a more rigorous adversary capability model. Finally, we hope to combine this approach with other anomaly detection mechanisms (e.g., [14]) to eliminate some of the assumptions imposed on sensor and actuator integrity.

Acknowledgements

This research was partially supported by Vistorm, an HP Company. The authors also wish to acknowledge the assistance of Diageo and, in particular, Mr. Brian Furey for helping validate the model and providing data sets for analysis.

References

[1] B. Bequette, *Process Control: Modeling, Design and Simulation*, Prentice-Hall, Upper Saddle River, New Jersey, 2002.

[2] J. Bigham, D. Gamez and N. Lu, Safeguarding SCADA systems with anomaly detection, *Proceedings of the Second International Workshop on Mathematical Methods, Models and Architectures for Computer Network Security*, pp. 171–182, 2003.

[3] E. Byres and D. Hoffman, The Myths and Facts behind Cyber Security Risks for Industrial Control Systems, Technical Report, Department of Computer Science, University of Victoria, Victoria, Canada, 2004.

[4] A. Cardenas, T. Roosta and S. Sastry, Rethinking security properties, threat models and the design space in sensor networks: A case study in SCADA systems, *Ad Hoc Networks*, vol. 7(8), pp. 1434–1447, 2009.

[5] S. Cheung, B. Dutertre, M. Fong, U. Lindqvist, K. Skinner and A. Valdes, Using model-based intrusion detection for SCADA networks, *Proceedings of the SCADA Security Scientific Symposium*, pp. 127–134, 2007.

[6] M. Coutinho, G. Lambert-Torres, L. da Silva, J. da Silva, J. Neto, E. da Costa Bortoni and H. Lazarek, Attack and fault identification in electric power control systems: An approach to improve the security, *Proceedings of the PowerTech Conference*, pp. 103–107, 2007.

[7] A. Creery and E. Byres, Industrial cybersecurity for power systems and SCADA networks, *Proceedings of the Fifty-Second Annual Petroleum and Chemical Industry Conference*, pp. 303–309, 2005.

[8] C. Dayharsh and H. Del Vecchio, Thermal death time studies on beer spoilage organisms, *Proceedings of the American Society of Brewing*, vol. II, pp. 48–52, 1952.

[9] D. Gamez, S. Nadjm-Tehrani, J. Bigham, C. Balducelli, K. Burbeck and T. Chyssler, Safeguarding critical infrastructures, in *Dependable Computing Systems: Paradigms, Performance Issues and Applications*, H. Diab and A. Zomaya (Eds.), Wiley-Interscience, Hoboken, New Jersey, 2005.

[10] G. Hoglund and J. Butler, *Rootkits: Subverting the Windows Kernel*, Addison-Wesley, Reading, Massachusetts, 2005.

[11] Y. Huang, A. Cardenas, S. Amin, Z. Lin, H. Tsai and S. Sastry, Understanding the physical and economic consequences of attacks on control systems, *International Journal of Critical Infrastructure Protection*, vol. 2(3), pp. 73–83, 2009.

[12] R. Krutz, *Securing SCADA Systems*, Wiley, Indianapolis, Indiana, 2006.

[13] O. Linda, T. Vollmer and M. Manic, Neural network based intrusion detection system for critical infrastructures, *Proceedings of the International Joint Conference on Neural Networks*, pp. 1827–1834, 2009.

[14] T. McEvoy and S. Wolthusen, Using observations of invariant behavior to detect malicious agency in distributed environments, *Proceedings of the Fourth International Conference on IT Incident Management and IT Forensics*, pp. 55–72, 2008.

[15] T. McEvoy and S. Wolthusen, Using observations of invariant behavior to detect malicious agency in distributed control systems, presented at the *Fourth International Workshop on Critical Information Infrastructures Security*, 2009.

[16] P. Motta Pires and L. Oliveira, Security aspects of SCADA and corporate network interconnections: An overview, *Proceedings of the International Conference on Dependability of Computer Systems*, pp. 127–134, 2006.

[17] J. Pearl, *Causality: Models, Reasoning and Inference*, Cambridge University Press, Cambridge, United Kingdom, 2009.

[18] J. Rrushi and K. Kang, Detecting anomalies in process control networks, in *Critical Infrastructure Protection III*, C. Palmer and S. Shenoi (Eds.), Springer, Heidelberg, Germany, pp. 151–165, 2009.

[19] J. Schlesser, D. Armstrong, A. Cinar, P. Ramanauskas and A. Negiz, Automated control and monitoring of thermal processing using high temperature, short time pasteurization, *Journal of Dairy Science*, vol. 80(10), pp. 2291–2296, 1997.

[20] N. Svendsen and S. Wolthusen, Modeling and detecting anomalies in SCADA systems, in *Critical Infrastructure Protection II*, M. Papa and S. Shenoi (Eds.), Springer, Boston, Massachusetts, pp. 101–113, 2008.

[21] J. Verba and M. Milvich, Idaho National Laboratory supervisory control and data acquisition intrusion detection system, *Proceedings of the IEEE Conference on Technologies for Homeland Security*, pp. 469–473, 2008.

[22] X. Wang, J. Lizier, O. Obst, M. Prokopenko and P. Wang, Spatiotemporal anomaly detection in gas monitoring sensor networks, *Proceedings of the Fifth European Conference on Wireless Sensor Networks*, pp. 90–105, 2008.

[23] D. Watts, Security and vulnerability in electric power systems, *Proceedings of the Thirty-Fifth North American Power Symposium*, pp. 559–566, 2003.

[24] D. Yang, A. Usynin and J. Hines, Anomaly-based intrusion detection for SCADA systems, *Proceedings of the Fifth International Topical Meeting on Nuclear Plant Instrumentation, Control and Human Machine Interface Technologies*, 2006.