A Plant-Wide Industrial Process Control Security Problem

Thomas Mcevoy, Stephen Wolthusen

Chapter 4

A PLANT-WIDE INDUSTRIAL PROCESS CONTROL SECURITY PROBLEM

Thomas McEvoy and Stephen Wolthusen

Abstract Industrial control systems are a vital part of the critical infrastructure. The potentially large impact of a failure makes them attractive targets for adversaries. Unfortunately, simplistic approaches to intrusion detection using protocol analysis or naïve statistical estimation techniques are inadequate in the face of skilled adversaries who can hide their presence with the appearance of legitimate actions.

This paper describes an approach for identifying malicious activity that involves the use of a path authentication mechanism in combination with state estimation for anomaly detection. The approach provides the ability to reason conjointly over computational structures, and operations and physical states. The well-known Tennessee Eastman reference problem is used to illustrate the efficacy of the approach.

Keywords: Industrial control systems, subversion detection

1. Introduction

In industrial control systems, detection and prevention extend beyond the computational model into the physical realm. While protocol analysis may signal anomalies as proposed by Coutinho, *et al.* [2], a skilled adversary can issue apparently authentic commands [18] using legitimate protocols. Analysis may be extended using state estimation techniques, but should not be applied naïvely [10, 16], especially in non-linear environments such as those encountered in the biochemical industry [6].

This paper describes an approach that utilizes state estimation in intrusion detection in combination with path authentication techniques. The approach assumes the existence of an adversary who can subvert channels and system functions [9]. Hence, it is necessary to verify the reliability and independence of channels and functions for message transmission. This is achieved by combining state estimation techniques using proxy measurements [10] with algebraic proofs over structures and operations. The Tennessee Eastman reference prob-

lem is employed as a case study to demonstrate the application of the approach to non-linear systems.

2. Related Work

Industrial control systems are a vital part of the critical infrastructure and are attractive targets for adversaries. Security in such systems is generally weak [3]. Recent research has focused on anomaly detection at the protocol level, since traffic in control networks is well-characterized and, hence, particularly amenable to such techniques [2]. Approaches using physical state estimation techniques have also been researched [15], but these are largely limited to linear systems. However, many industrial systems, including biological and chemical processes, exhibit non-linear behavior or require non-linear control laws, resulting in less well-defined models and limited accuracy [6]. Real-time detection is also an important requirement for these industrial systems [17].

It has been argued that, in the presence of channel compromise, adversaries may use protocols correctly and present syntactically and semantically correct messages, resulting in the failure of conventional detection techniques to signal anomalies [9, 18]. These attacks may also be concealed in noisy processes that are not amenable to elementary statistical analysis [16]. In particular, this is true for non-linear systems [10]. The Tennessee Eastman reference problem [5] is commonly considered in control systems research and pedagogy (see, e.g., [1, 7, 8, 12]). It provides a well-defined problem space for using different control laws. Furthermore, a number of simulation models are available for this problem. The process calculus used to construct the control overlay model in this paper was defined in [9], where an adversary capability model for industrial control systems was also proposed. This paper uses the process calculus model to analyze computational structures and operations using techniques related to probabilistic packet marking and path authentication [4].

3. Control Problem

An attack on an industrial control system is usually accompanied by the use of concealment techniques. Protocol analysis by itself may not detect an attack that uses legitimate protocols. State estimation techniques rely on the integrity of the signals. They can deal with missing data and noisy signals, but not with deceptive or misleading signals from subverted channels. Hence, conjoint reasoning is required over both channels and signals to uncover malicious activity, helping separate false and true signals.

4. Solution Approach

We define a computational overlay for an industrial control system using an applied π-calculus [13]. In the context of the Tennessee Eastman challenge problem [5], we demonstrate the existence of proxy measurements of plant activity that can be used to detect anomalies. However, this requires the ability

to reason about channel integrity. This is accomplished using path authentication methods that can be proven within the algebraic framework. An explicit model of human intervention is not presented, rather we consider operational capability in terms of detection.

5. Process Calculus

The capabilities of our π-calculus variant are specified by:

$$\pi ::= \bar{x}y_{p,r} \mid x(z_{p,r}) \mid \tau \mid \lambda \mid f(\overrightarrow{z}) \rightarrow \bar{x}w, w \mid [x = y]\pi$$

A simplified version of the process calculus was presented in [9], where it was used to represent adversary capabilities. Here, we expand on its functionality to permit proofs over structures and operations.

The capabilities of the process calculus are: (i) sending a name with priority and routing; (ii) receiving a name with priority and routing; (iii) performing an unobserved action (with the special meaning of decision-making); (iv) performing an observable inaction (process failure); (v) name generating function; (vi) replication capability; and (vii) conditional capability. $\tilde{z}$ is used to denote a vector of names. Names are typed as channels, variables or constants.

The operations of the π-calculus are retained and augmented as follows:

$$P ::= M \mid P|P' \mid \nu z\ P \mid !P$$
$$M ::= \mathbf{0} \mid \pi.P \mid M + M' \mid M \oplus M'$$

where P is a process that may be a summation, concurrent, a new process with (restricted) names, or replication. M is a summation that may be null or termination, a capability guarding a process and – adding a variant summation – a soft choice between retained alternatives and a hard choice between mutually exclusive alternatives (see Sangiorgi and Walker [13] for additional details). Hence, a process may partially order its messaging and the exercising of its capabilities in a manner that is not mutually exclusive. For example, the process may send a set of messages in some order. However, it cannot be subverted as an agent of the adversary and also resist such subversion because these outcomes are mutually exclusive.

The name generating function takes a set of parameters and returns a name. In general, it provides a parametric interface to the physical processes or control functions that may be defined by a state space equation or its transform. The function can also be used for other purposes, for example, to simulate automated decision-making or as a cryptographic primitive.

Routing captures the ability of the system to send a message to a process by means of another process, provided the name of the process exists in the intervening process. Routing information may be explicitly coded in the summation or understood implicitly from the process structure. For example, $\bar{x}m_y.\mathbf{0}|x(u).\bar{r}u_r.\mathbf{0} + x(u).\bar{s}u_{[y]}.\mathbf{0}|s(u).\bar{y}u.\mathbf{0}|r(u).\mathbf{0}|y(u).\mathbf{0}$ sends m to x and for-

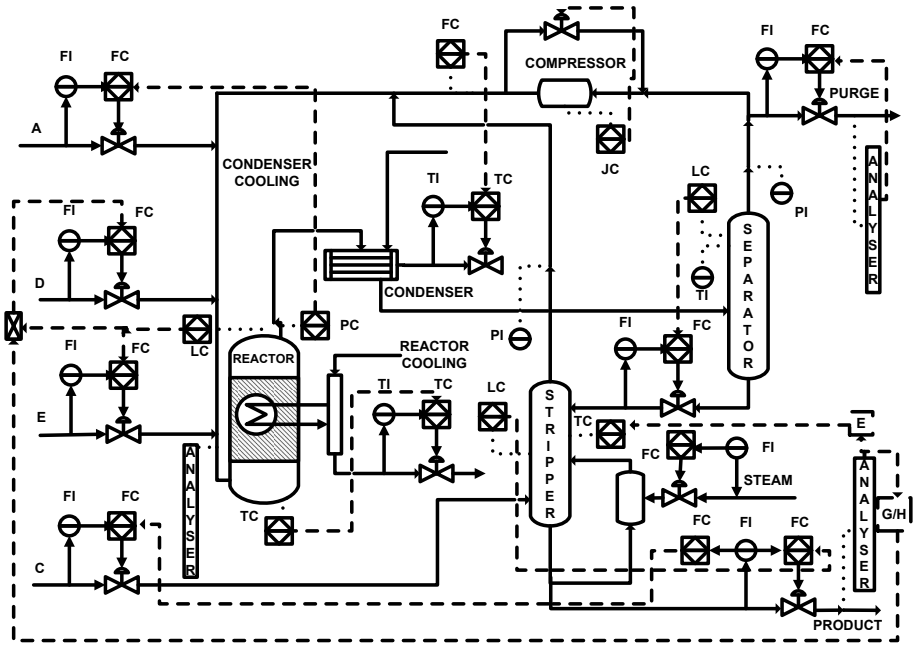

Figure 1. Tennessee Eastman problem under base control [11].

wards it to y, but not to r. Prioritization can be captured by a simple ranking system [9].

Special types of functions are defined using a finite set of labels λ (e.g., delay and message loss). The actions of these properties can be described as necessary. However, they are essentially means for naming invisible process actions that would otherwise be regarded as degenerate terminations. The following equation illustrates one use of labels:

$$((\bar{x}u + x(u)).\mathbf{0} + Loss + Delay) \equiv ((\bar{x}u + x(u).\mathbf{0}) + \mathbf{0} + \mathbf{0})$$

6. Model Creation

This section describes how a suitable state estimation algorithm may be used along with proxy measurements or estimators in combination with path authentication techniques to uncover reliable channels and to maintain system operations in the presence of malicious activity. The Tennessee Eastman challenge problem is used to illustrate the application of the approach to non-linear estimation problems for industrial control systems.

6.1 Tennessee Eastman Problem

The Tennessee Eastman plant is a non-contrived, albeit modified, model of a real chemical process (Figure 1). It consists of a reactor-separator-recycler

arrangement involving two simultaneous irreversible gas-liquid exothermic reactions and two byproduct reactions given by:

$A\ (g) + C\ (g) + D\ (g) \rightarrow G\ (l)$	*Product* 1
$A\ (g) + C\ (g) + E\ (g) \rightarrow H\ (l)$	*Product* 2
$A\ (g) + E\ (g) \rightarrow F\ (l)$	*Byproduct*
$3D\ (g) \rightarrow 2F\ (l)$	*Byproduct*

The plant is open-loop, unstable and highly non-linear. Various approaches to its control have been described [8], which can result in distinct channel architectures, rendering it a suitable candidate for testing a variety of models and techniques.

The gaseous reactants (g) form liquid products (l). Note that the products are not specifically identified and that the process was modified from the real industrial process by the original authors [5]. The gas phase reactions are catalyzed by a non-volatile substance dissolved in the liquid phase in the reactor. The reactor is pressurized and agitated, and uses an internal cooling bundle to remove the heat produced by the reactions.

The products leave the reactor in the vapor phase along with the unreacted feeds, while the catalyst remains in the reactor. The reactor product stream passes through a cooler that condenses the products, and from there to a vapor-liquid separator. Non-condensed components cycle back to the reactor feed via a centrifugal compressor. Condensed components are sent to a product stripping column that removes the remaining reactants.

Products G and H exit the stripper base and are separated in a downstream refining section, which is not included in the problem statement. The byproducts and inerts are purged from the system in the vapor phase using a vapor-liquid separator. The system may be operated in six distinct modes to produce different product mass outputs.

The plant has twelve valves for manipulation, and a total of 41 measurements are involved in monitoring and control. Note that one of the valves is not shown in Figure 1, which only provides closed control loops; the valve is used for higher order control purposes.

Following the base control strategy outlined by McAvoy and Ye [11], most of the variables may be removed from consideration to leave the twelve control variables and twelve manipulated variables shown in Table 1. Hence, for state estimation purposes, depending on the control law used, not all the variables need to be considered. This implies a set of alternative measurements may be available as proxies for the main control variables. This also means that, for state estimation purposes, there are a number of possible measurements in addition to the main ones in the model that can be used for estimation by proxy [10].

Table 1. Manipulated and controlled variables.

Manipulated Variable	Controlled Variable
A-feed set point	Reactor level
D-feed set point	Separator level
E-feed set point	Stripper bottom level
C-feed set point	Reactor pressure
Purge set point	Reactor feed flow
Product set point	Reactor temperature
Stripper stream flow set point	Compressor power
Separator bottom flow set point	Compressor exit flow
Reactor cooling water set point	Separator pressure
Condenser cooling water set point	Separator temperature
Compressor recycle valve	Stripper pressure
Stirrer speed	Stripper temperature

6.2 Tennessee Eastman Overlay

Using our process calculus, we can define a system architecture that satisfies the control purposes. To do so, we define the entities, messengers and agents of the system. By τ, entities make decisions. Messengers pass decisions as names. By $f() \rightarrow$, agents are processes which act on decisions. For example, an operator that is an entity is defined by the equation:

$$Operator := \bar{x}u.\mathbf{0} \oplus x(u).\mathbf{0} \oplus \tau.\mathbf{0} | !Operator$$

where the set $Operator = \{Operator, Adversary\}$ and τ is the decision-making capability. A (simple) controller may be defined by:

$$\begin{aligned} Controller := \nu i((z(e)_1.\bar{z}e.f(p,k,e,i) \rightarrow \bar{z}i_1 \\ + y(k')_2.Controller\langle p,k',e\rangle.\bar{y}k_2 \\ + y(p')_2.Controller\langle p',k,e\rangle.\bar{y}p_2).\mathbf{0} \\ + (y(m).Controller'\langle p,k,e\rangle.\mathbf{0} \oplus Resist \,|\, !Controller) \end{aligned}$$

where the controller may be changed to an agent of the adversary by a malicious message m that represents a successful attack, and R is the ability to resist such an attack with the set $Agent := \{Agent, PlantProcess\}$ representing the agent state. Other examples of control system structures are provided in [9]. They can be used to create the complete system infrastructure.

6.3 State Estimation

State estimation is the problem of accounting for the state of a system in the presence of process disturbances and measurement noise. A general non-linear system can be described as:

$$x_{k+1} = f(x_k, u_k) + W_k \qquad \textit{System Equation}$$
$$y_{k+1} = h(x_k) + v_k \qquad \textit{Output Equation}$$

where x is the state variable vector, u represents the inputs under control, w represents process noise, y is the measured output and v is the measurement noise. Note that x is not known directly, but is estimated based on y; this accounts statistically for both process and measurement noise. We assume that process and measurement noise can be represented as Gaussian white noise with a mean of zero ($\mu = 0$) and a suitable variance (σ^2).

Several state estimation algorithms are available for this purpose. An example is the extended Kalman filter [14]. Note, however, that state estimation techniques in general are defined recursively and hence have a "memory" of the previous states of a system. This distinguishes them from pure correlation techniques where the memory of previous system behavior is lost.

In the case of most industrial systems, it is possible to derive multiple sets of measurements that are functionally independent of each other in control terms. Thus, alternative means exist for testing the reliability of measurements and the ability to substitute one set of measurements for another for control and channel authentication purposes. For example, in the Tennessee Eastman system, influx A in Figure 1 can be measured directly by its flow meter and estimated by the initial flow analyzer, pressure controller and also inversely estimated based on D and E, C, G and H. Both the estimation techniques can be used and their results compared to identify inconsistencies and determine the integrity of channels and functions.

7. Model Application

We assume the existence of an adversary who can subvert channels and functions to act on his behalf. This means that encryption techniques cannot be used to guarantee the freshness or authenticity of messages since the message originators may be compromised by the adversary. In particular, the adversary (or rather his agents) can perfectly forge messages with respect to the protocol formulation and/or directly manipulate physical measurements.

We assume that a set of robust estimators E exist for a system such as the Tennessee Eastman problem, which we can use to detect inconsistent measurements. (The estimators are derived by simulation.) The goal is to clearly mark channels and sensors (controls) as reliable or unreliable to avoid an unnecessary system shutdown. To do this, it is necessary to prove that a set $n = |E(\cdot)|$ of independent channels exists for each estimator. In the case of untainted channels, the associated estimators can be used. However, if all the channels for an estimator are tainted, then a contingent estimator can be used provided that its channels are untainted. Clearly, a complete set of fully separated channels provides a trivial (but not minimal) solution. Non-trivial solutions are required

because channels are generally shared by messages due to the convergence of channels onto an operator and resilience characteristics.

Channel independence may be demonstrated by variations on packet marking for path authentication [4]. Several such techniques may be investigated for applicability, considering parameters such as topological constraints.

We illustrate one technique by constructing a set of channels that use a "salt" to mark the route selected by a message. The salt is a shared secret between the channel and the operator. We assume that a set of known routes exist over which we define "normal routes" and "deviations." For each deviation, a salt is added to provide a trace of the path followed by a signal package.

Let the $\{P1, P2, P3, P4\}$ be the controllers and Op be the operator as previously defined. We assume that each controller hashes the message identifiably. We define a set of channels such that each channel may re-route a message to an adjacent channel on message failure ($Loss$). Before doing so, it rehashes the message hash with its salt and attaches its name in order. The channels are defined by the equations:

$$\begin{aligned}
Cn :=& \nu s(\bar{x}_{Dn} u_{[Op]} + x_{Cn}(u)_{[Op]} + Loss.Hash(u,s) \rightarrow \bar{w}_{C(n+1)} z_{[Op]} \\
&+ w_{C(n-1)}(z)_{[Op]} + \bar{x}_{Dn} z_{[Op]}).\mathbf{0} | !Cn \\
Dn :=& \nu s(\bar{x}_{D} u_{[Op]} + x_{Cn}(u)_{[Op]} + Loss.Hash(u,s) \rightarrow \bar{w}_{D(n+1)} z_{[Op]} \\
&+ w_{D(n-1)}(z)_{[Op]} + \bar{x}_{Dn} z_{[Op]}).\mathbf{0} | !Dn \\
En :=& \nu s(\bar{x}_{Op} u_{[Op]} + x_{En}(u)_{[Op]} + Loss.Hash(u,s) \rightarrow \bar{w}_{E(n+1)} z_{[Op]} \\
&+ w_{E(n-1)}(z)_{[Op]} + \bar{x}_{Op} z_{[Op]}).\mathbf{0} | !En
\end{aligned}$$

The overall structure is given by the equations:

$$\begin{aligned}
&\bar{x}_{P1} z x_{[Op]}, 1.P1 | C1 | D1 | E1 | \\
&\bar{x}_{P2} z x_{[Op]}, 1.P2 | C2 | D2 | E2 | Op \\
&\bar{x}_{P3} z x_{[Op]}, 1.P3 | C3 | D3 | E3 | \\
&\bar{x}_{P4} z x_{[Op]}, 1.P4 | C4 | D4 | E4 |
\end{aligned}$$

Note that the topology is deliberately constrained, a characteristic of industrial control systems.

We claim that each message follows a route that is uniquely identified by its origin and its membership in the set of deviations. Let $K_{m,n}$ be a message with n salts and $m = n + 1$ names. The name order must be consistent with the deviations permitted by the topology and must match the salt order. We subtract a name and a salt from K. Let $K_{m,n} = K_{m-1,n-1}$. We treat this as a move in a game. If the move $K_{m,n} \xrightarrow{\alpha} K_{m-1,n-1}$ is not permitted, where α is the trace that is the set of channels between the two marked channels, say K_P and K_Q, then the routing is invalid. If the routing is valid then, the operation can be repeated until $K_{1,0}$ is reached, which should be the expected origin of the message. Thus, the route taken by each message can be identified. Since each

message follows a uniquely identifiable route, an inconsistent message marks a potentially subverted route.

Using set elimination over routes between an orgin and destination, $\sigma_{i,j}\alpha_i - \alpha_j$, the subverted channels can be identified in a probabilistic manner. Hence, if a message is sent independent of an unreliable channel, it may be regarded as reliable; otherwise, it is not reliable. Observing the independence of channels permits the detection of the adversary's action and operation of the plant, even where manipulated signals share routes with reliable signals.

To complete the approach, the set of estimators should also be independent sources of information about the process. A cyclic dependency between estimators must be avoided. For example, if the estimator $A1$ is used to estimate $B2$, and $B2$ to estimate $C4$, and $C4$ to estimate $A1$, then the results become meaningless. Undermining this approach requires the adversary to capture all the salts that are regard as infeasible. In essence, we assume the adversary can only gain partial control of the system.

8. Conclusions

Research in the area of control systems security has shown that attackers can forge protocols or directly manipulate physical signals to mask their activities. In earlier work [10], we have demonstrated previously that proxy measurements can detect such inconsistencies. However, to minimize the reengineering efforts, it is desirable to use measurements that are already present. Combining path authentication with state estimation techniques is an effective means for identifying subverted channels and processes, and, as such, promises to be a rich area of research in the area of control systems security. Our future research will focus on refining the path authentication technique and selecting robust estimators for state estimation by proxy.

References

[1] L. Bie and X. Wang, Fault detection and diagnosis of a continuous process based on multiblock principal component analysis, *Proceedings of the International Conference on Computer Engineering and Technology*, pp. 200–204, 2009.

[2] 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 security, *Proceedings of the Power Tech Conference*, pp. 103–107, 2007.

[3] 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.

[4] X. Dang, E. Albright and A. Abonamah, Performance analysis of probabilistic packet marking in IPv6, *Computer Communications*, vol. 30(16), pp. 3193–3202, 2007.

[5] J. Downs and E. Vogel, A plant-wide industrial process control problem, *Computers and Chemical Engineering*, vol. 17(3), pp. 245–255, 1993.

[6] 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.), John Wiley, Hoboken, New Jersey, pp. 479–499, 2005.

[7] T. Kraus, P. Kuhl, L. Wirsching, H. Bock and M. Diehl, A moving horizon state estimation algorithm applied to the Tennessee Eastman benchmark process, *Proceedings of the IEEE International Conference on Multisensor Fusion and Integration for Intelligent Systems*, pp. 377–382, 2006.

[8] T. Larsson and S. Skogestad, Plant-wide control – A review and a new design procedure, *Modeling, Identification and Control*, vol. 21(4), pp. 209–240, 2000.

[9] T. McEvoy and S. Wolthusen, A formal adversary capability model for SCADA environments, presented at the *Fifth International Workshop on Critical Information Infrastructure Security*, 2010.

[10] T. McEvoy and S. Wolthusen, Detecting sensor signal manipulations in non-linear chemical processes, in *Critical Infrastructure Protection IV*, T. Moore and S. Shenoi (Eds.), Springer, Heidelberg, Germany, pp. 81–94, 2010.

[11] T. McAvoy and N. Ye, Base control for the Tennessee Eastman problem, *Computers and Chemical Engineering*, vol. 18(5), pp. 383–413, 1994.

[12] N. Ricker, Decentralized control of the Tennessee Eastman challenge process, *Journal of Process Control*, vol. 6(4), pp. 205–221, 1996.

[13] D. Sangiorgi and D. Walker, *π-Calculus: A Theory of Mobile Processes*, Cambridge University Press, Cambridge, United Kingdom, 2001.

[14] D. Simon, *Optimal State Estimation: Kalman, H_∞ and Nonlinear Approaches*, John Wiley, Hoboken, New Jersey, 2006.

[15] S. Su, X. Duan, X. Zeng, W. Chan and K. Li, Context information-based cyber security defense of protection system, *IEEE Transactions on Power Delivery*, vol. 22(3), pp. 1477–1481, 2007.

[16] N. Svendsen and S. Wolthusen, Using physical models for anomaly detection in control systems, in *Critical Infrastructure Protection III*, C. Palmer and S. Shenoi (Eds.), Springer, Heidelberg, Germany, pp. 139–149, 2009.

[17] C. Ten, G. Manimaran and C. Liu, Cybersecurity for critical infrastructures: Attack and defense modeling, *IEEE Transactions on Systems, Man and Cybernetics (Part A: Systems and Humans)*, vol. 40(4), pp. 853–865, 2010.

[18] J. Verba and M. Milvich, Idaho National Laboratory Supervisory Control and Data Acquisition Intrusion Detection System (SCADA IDS), *Proceedings of the IEEE Conference on Technologies for Homeland Security*, pp. 469–473, 2008.