Chapter 8

MODELING AND DETECTING ANOMALIES IN SCADA SYSTEMS

Nils Svendsen and Stephen Wolthusen

Abstract The detection of attacks and intrusions based on anomalies is hampered by the limits of specificity underlying the detection techniques. However, in the case of many critical infrastructure systems, domain-specific knowledge and models can impose constraints that potentially reduce error rates. At the same time, attackers can use their knowledge of system behavior to mask their manipulations, causing adverse effects to observed only after a significant period of time. This paper describes elementary statistical techniques that can be applied to detect anomalies in critical infrastructure networks. A SCADA system employed in liquefied natural gas (LNG) production is used as a case study.

Keywords: SCADA systems, anomaly detection, multivariate analysis

1. Introduction

Supervisory control and data acquisition (SCADA) networks are a key component of the critical infrastructure. These systems are used by operators in modern industrial facilities to continuously monitor and control plant operations. SCADA systems have evolved in terms of the capabilities of their sensors and actuators as well as in their network topologies. SCADA network topologies have moved from simple point-to-point links to arbitrary mesh-type networks, including fixed and wireless links that support large numbers of nodes and overlapping networks.

Although the importance of SCADA systems has been recognized for some time [21], efforts investigating network security issues in SCADA environments have been relatively limited [3, 13]. Ijure, *et al.* [13] identify several security challenges that have to be addressed for SCADA networks: access control, firewalls and intrusion detection systems, protocol vulnerability assessment, cryptography and key management, device and operating system security, and security management.

Please use the following format when citing this chapter:

Svendsen, N. and Wolthusen, S., 2008, in IFIP International Federation for Information Processing, Volume 290; *Critical Infrastructure Protection II*, eds. Papa, M., Shenoi, S., (Boston: Springer), pp. 101–113.

This paper is concerned with the question of whether the implementation of traditional security solutions in a SCADA network will provide adequate levels of security given the constraints and requirements imposed by the application area. The primary requirement is to maintain physical parameters within a set of quality and safety margins and to guarantee suitable reaction times. This is accomplished by gathering data from multiple (possibly hierarchical) sensors and subsystems, and verifying that the readings fall into acceptable ranges based on historical data. However, an attacker with the appropriate knowledge and access can alter correlated process variables to bring a system to a critical state, potentially causing degradation of service or even an outright failure. This paper employs applied statistical methods to detect anomalous behavior in SCADA networks. A case study involving a liquefied natural gas (LNG) production facility is used to demonstrate the utility of the statistical approach.

2. Anomaly Detection in Control Systems

This section provides a brief overview of anomaly detection in control systems followed by an overview of applied statistical methods.

2.1 Anomaly Detection

A control system is a device or set of devices used to manage, command, direct and regulate the behavior of other devices or systems. It typically has four main components: sensors, analyzers, actuators and a communications infrastructure. Sensors determine the state of the controlled system, analyzers determine whether the system is stable or out of control, and actuators are used to maintain the system at (or restore it to) a stable state. Control systems incorporate feedback loops, which may be positive or negative, depending on the application.

Anomaly detection [2] involves establishing profiles of normal process behavior, comparing actual behavior with the established profiles, and identifying deviations from the normal. A profile or set of metrics is determined for each process. The metrics are measures of specific aspects of process behavior (e.g., pressure, temperature or composition).

Anomaly detection methods may be categorized as: statistical methods, rule-based methods, distance-based methods, profiling methods and model-based approaches. This paper focuses on statistical methods for anomaly detection. Denning [9] proposed four statistical models for determining whether an observation is abnormal with respect to previous observations. They are: (i) operational model (where abnormality is determined by comparing a new observation against fixed limits); mean and standard deviation model (where an observation is compared to a confidence interval based on historical observations); multivariate model (where correlations between two or more metrics are taken into account); and Markov process model (used in discrete systems where

transaction frequencies between states and the probability of going from one state to another can be determined).

2.2 Univariate Quality Control Charts

Univariate quality control charts (see, e.g., [27]) can be used to determine if the performance of a process is at an acceptable level. A quality control chart consists of data plotted in time order and horizontal lines, called control limits, that indicate the amount of variation due to common causes. Control must be exerted on both the central tendency and variability, which are accomplished using an $\bar{X}$ -chart and an S -chart, respectively.

Assume that the data consists of m samples of size n for which $S_1, S_2, \dots, S_m$ are the sample standard deviations. The average values of the sample standard deviation $\bar{S}$ are computed along with the overall average $\bar{\bar{X}}$. The corresponding upper and lower control limits for the $\bar{X}$ -chart to control the central tendency are:

$$UCL = \bar{\bar{X}} + A_3\bar{S} \quad LCL = \bar{\bar{X}} - A_3\bar{S}$$

where $A_3 = 3/(c_4\sqrt{n})$ and

$$c_4 = \left(\frac{2}{n-1} \right)^{1/2} \frac{\Gamma(n/2)}{\Gamma[(n-1)/2]}$$

where $\Gamma(\cdot)$ is the gamma function. For the S -chart, we have:

$$UCL = B_6\sigma \quad LCL = B_5\sigma$$

with $B_5 = c_4 - 3\sqrt{1 - c_4^2}$ and $B_6 = c_4 + 3\sqrt{1 - c_4^2}$. Given the control limits, the quality control charts are created by plotting the sample means (standard deviations) in time order in the same plot.

2.3 Multivariate Quality Control Charts

A multivariate approach is used when the data to be analyzed has multiple important characteristics. Such an approach may also be used when processes are assumed to be independent. The T^2 -chart is commonly applied in these situations as it can be applied to a large number of variables. Given the mutually independent vectors $X_1, X_2, \dots, X_n$ of length p where each X_j is distributed as $N_p(\mu_j, \Sigma)$, the control limits of the T^2 -chart are set by assuming that $(X_j - \bar{X})'S^{-1}(X_j - \bar{X})$ has a chi-square distribution [14]. Note that S is the covariance matrix and $(X_j - \bar{X})'$ is the transpose of $(X_j - \bar{X})$. For the j th point, the T^2 -statistic is computed as:

$$T_j^2 = (x_j - \bar{x})'S^{-1}(x_j - \bar{x})$$

and plotted on the time axis. The lower control limit (LCL) is zero while the upper control limit (UCL) is commonly set to $\chi_p^2(0.05)$.

3. Liquefied Natural Gas Production

This section briefly describes the process for producing liquefied natural gas (LNG) [23].

Natural gas (NG) is retrieved from wells, each of which is controlled by a set of valves (“Xmas trees”) that adapts the NG pressure from the well to the pressure in the pipeline system. At this point, monoethyleneglycol (MEG) is injected into the well-stream to inhibit the formation of hydrate, which could block the pipeline. MEG is distributed to the Xmas trees through a control distribution unit (CDU). The CDU also distributes electricity, control signals, hydraulic pressure and chemicals to the Xmas trees. The various well-streams are assembled at the pipeline end manifold (PLEM), where they gather into a single well-stream for transport through the main pipeline. The flow in the main pipeline has multiple phases: natural gas liquids (NGL), condensate (light oil) and a mix of water and MEG.

The well-stream in the main pipeline often arrives in spurts, i.e., the gas and liquids separate and the gas arrives between slugs of liquid. A slug catcher is typically used to separate NG condensate and MEG. Carbon dioxide (CO_2) is then removed from the NG as it would freeze to dry ice during gas liquefaction, which could cause damage later in the process. The NG is already moist and the removal of CO_2 further augments its water content, which would form ice during the cooling process and cause damage. Therefore, the gas is dried before refrigeration. Another important pre-treatment process is the removal of very small quantities of mercury present in the heavier components. This is because mercury could cause corrosive damage to metal components and catalysts that come into contact with the gas stream.

At this point, the NG is ready for fractionation. This involves the separation of LNG from the heavier gas components, known as natural gas liquids (NGL), and the adjustment of the amounts of various hydrocarbons present in the gas. The gases that remain after NGL removal are passed to a “cold box” for cooling to LNG. This is a three-stage process that primarily employs heat exchangers. A byproduct of this phase is nitrogen, which is purified and released to the atmosphere. After NGL is separated from LNG, the NGL undergoes further fractionation to separate ethane and propane from the remaining condensate. Ethane and propane form liquefied petroleum gases (LPG).

LNG/LPG production is energy intensive. The energy requirement to bring the gas from high pressure and relatively high temperature to low pressure and very low temperature is tremendous as the pressure varies from 220 bar to 1 bar and the temperature from 90°C to -163°C . LNG/LPG plants tend to be self-sufficient with regard to energy since they operate gas-driven power plants; this largely eliminates the dependency on external power suppliers and the power grid.

4. LNG Process Attack Points

This section identifies possible LNG process attack points. The focus is on attacks that could halt or degrade LNG production. Outright terrorist acts and sabotage, such as blowing up a storage facility, are not included in the list of attack scenarios. Instead, the scenarios mainly involve subtle manipulations of process control systems and sensors. The scenarios assume that an attacker is knowledgeable about the system.

- **MEG Dosage:** MEG must be present in the well-stream to prevent the water component from freezing. An ice plug could cause a pipeline blockage, resulting in a lengthy shutdown of the plant. Also, the upstream pressure in the pipeline could rise to critical levels.
- **CO₂ Removal:** CO₂ can freeze into dry ice and cause a pipeline blockage, resulting in a lengthy shutdown. Pipeline pressure could also rise to critical levels.
- **Mercury Removal:** In this subtle scenario, the presence of mercury causes pipeline corrosion over the long term. In the best case, this increases maintenance costs; in the worst case, the pipeline could rupture.

The remainder of this paper focuses on how an attacker, by altering the moisture content readings for well-streams, could bring the MEG concentration to a critically low level without it being detected by sensors in the well-heads. The attack is carried out so that the moisture content at each well-head is within the control limit of the stream, meaning that it cannot be detected by univariate analysis. However, it can be detected by observing the correlation between the well-heads.

5. Model Description

This section presents two models, one for monitoring well-streams for unusual fluctuations in the volume flow of water and the other for relating the volume flow of water and the amount of MEG introduced.

5.1 Moisture Content in Well-Streams

Although it is a continuous phenomenon, the moisture content of a well-stream can be represented as a time series. We employ an elementary time series model that includes trend, seasonality and random noise [5]. Each observation X_t of a time series is of the form:

$$X_t = m_t + s_t + Y_t$$

where m_t is a slowly changing function (trend component), s_t is a periodic function of t with period d (seasonal effect), and Y_t is a zero-mean process (random noise and fluctuations). To capture the continuous properties of the

moisture content, we use a random walk to represent fluctuations. For each well i , a well-stream X_{it} is defined. The volume flow from well i is given by Q_i . Thus, the volume flow of water at each time interval is given by the product $X_{it} \cdot Q_i$. Generally, Q_i can be made time-dependent, but we choose to keep it constant for the purpose of our analysis. The attack on the LNG production system is accomplished by introducing an extra constant term in the expression for X_{it} at time t_a during Δ_a iterations. Given the amplitude of the attack $A \in [0, 1]$, the time series has the form:

$$X_t = \begin{cases} m_t + s_t + Y_t, & \text{if } x \notin [t_a, t_a + \Delta_a] \\ m_t + s_t + Y_t + Am_t, & \text{if } x \in [t_a, t_a + \Delta_a] \end{cases}$$

5.2 MEG Dosage

In order to create an elementary model that relates the volume flow of water in a well-stream and the quantity of MEG added, we assume that the well-streams are merged to one stream at the PLEM and that the sensors for measuring the water content in the individual well-streams and the joint well-stream are located at the PLEM. The main consequence is that a latency emerges between the time an attack is initiated (i.e., a change occurs in a well-stream) and the time when the attack is detected. The relationship between the MEG dosage and the water volume flow is given by $Q_{MEG}(t) = f(Q_{water}(t + \Delta_t))$ where $f(\cdot)$ is some function. Due to natural fluctuations in a well-stream, the MEG dosage is not adjusted based on an individual reading, but on a statistical test of whether the value of the current MEG dosage corresponds to the mean of the last k well-stream readings. The process is initiated with an expected water volume flow μ_0 . For every water volume flow measurement, a test is performed to determine whether or not μ_0 is the mean of the last k readings. Assuming that the mean of the readings is μ , a one-sided test on a single sample can be performed using the hypothesis:

$$H_0 : \mu = \mu_0 \quad H_1 : \mu \neq \mu_0.$$

If the H_0 hypothesis is rejected, $\mu_0 = \mu$ holds and the MEG flow is altered according to the function f .

6. Simulation Results

Our simulation experiments consider a system with three well-streams. This section presents a reference simulation to demonstrate how a well-calibrated model is located within the control limits. Next, an attack is launched against all three well-streams and an attempt is made to detect the attack using quality control methods.

6.1 Three Wells with Seasonal Component

The following model is used to express the water content in the three wells:

Figure 1. Characteristic plots of a reference well-stream.

$$\begin{aligned}
 X_{1t} &= 0.3 + 0.03 \cos\left(\frac{i\pi}{100}\right) + 0.0005Y_{1t} \\
 X_{2t} &= 0.4 + 0.04 \cos\left(\frac{i\pi}{80} + \frac{2\pi}{3}\right) + 0.0005Y_{2t} \\
 X_{3t} &= 0.5 + 0.03 \cos\left(\frac{i\pi}{60} + \frac{4\pi}{3}\right) + 0.0005Y_{3t}
 \end{aligned}$$

where $Y_{it} = Z_{i1} + Z_{i2} + \dots + Z_{it}$, for $t = 1, 2, \dots, 3750$ and $\{Z_{it}\}$ is independent and identically distributed random noise. The generated time series points are sampled at a rate of $1/25$. This is done to show that not every point for a continuous process can be sampled. The points are grouped in samples of fifteen elements before further analysis is performed. Figure 1 shows the characteristic plots of the well-streams and the control charts for the joint well-stream.

An attack is launched simultaneously against all three wells; the attack increases the water content of each well-stream by 15%. Note that in order for the attack to be successful, the attacker must have knowledge of the sampling strategy and the grouping of samples. The confidence limits for the $\bar{X}$ -chart and S -chart are set based on historical observations of the process. A total of

Figure 2. Characteristics plots of a well-stream being attacked.

25 runs of the reference process described above were carried out to determine the 95% upper and lower confidence intervals. Figure 2 shows the characteristic plots for the attack, which occurs between iterations 1250 and 1600. As seen in Figure 2(c), the attack can be detected by the change in variance.

Having assumed that the attacker has knowledge about the sampling strategy, we now examine the situation where the attack covers full samples. This means that the samples either contain points that are attack points or points that are not attack points. Thus, internal fluctuations in the samples are avoided. Figure 3 shows one such run where the attack produces no more variation than noise.

6.2 MEG Dosage

The water content is modeled using the series:

$$X_t = 0.3 + 0.03 \cos\left(\frac{i\pi}{100}\right) + 0.0005Y_{1t}.$$

In the attack, a certain percentage of the expected flow is added to the well-stream. The duration of the attack corresponds to fourteen analyzed samples.

Figure 3. Well-streams and S -chart of a well-stream being attacked.

Table 1. Average number of samples before a change is detected.

Amplitude (%)	0	5	10	15	20	25	30
Samples to detect	14.6	6.8	4.1	3.6	2.1	2.3	1.6

Using the statistical analysis described above, we determine the average number of samples that must be considered before the change in the mean is detected (and the MEG quantity is adjusted). The results are presented in Table 1, which lists the average numbers of samples for expected flow percentages ranging from 0% to 30%.

7. Analysis of Scenarios

This section analyzes the simulated scenarios and discusses how the statistical approach works in the case of time series with tendencies.

7.1 Three Wells with Seasonal Component

An examination of Figure 2 indicates that the attack is not detected by the $\bar{X}$ -chart. Specifically, a peak in the sample mean is present, but it does not go over the confidence limits. We ran the attack 25 times and examined the fluctuations in the mean value for the joined well-stream. The attack was detected in 50% of the cases in the well-stream with the smallest volume flow (i.e., with the greatest sensitivity). This detection rate is only three times the false alarm rate in the stream due to random fluctuations. Note, however, that the T^2 -chart detects the attack in all the cases.

7.2 Three Wells with Tendency

It is reasonable to assume that there is a tendency different from zero in the water content of a well-stream. Either there is a known model for the

tendency or the tendency can be predicted either by using a time series model for forecasting or by smoothing and interpolation.

Using the difference between the observed value and the predicted value facilitates an analysis similar to the previous case (three wells with a seasonal component). As a matter of fact, the simulations indicate (but do not confirm) that the smoothing of the signal prior to prediction can help hide attacks.

7.3 MEG Dosage

Several parameters may be adjusted in this scenario. These range from the definition of the time series and its fluctuations to details such as sample size, sampling strategy and sensitivity of hypothesis testing. However, the simulation results show that delays accumulate in large-scale systems where sensors and actuators are located in different physical locations and where the nature of the observed system is such that control actions cannot be performed based on single observations.

8. Related Work

Early work on SCADA security focused almost exclusively on physical attacks [17]. However, intrusion detection in SCADA systems has become an important research area, especially as general-purpose network substrates are employed in control systems, and control networks are intentionally (and sometimes inadvertently) cross-linked with exposed networks. The risks to control networks posed by remote attacks were emphasized in a 1997 White House document [21]. However, much of the research related to SCADA security (see, e.g., [1, 13, 16, 18, 20]) has been driven by security-related incidents that occurred in 2001–2004 [11].

Considerable attention has focused on attacks against electrical power systems [26], although security issues related to other infrastructures have also been investigated [22]. The survivability of distributed control systems and their resilience to attacks, including subversion, is a major issue [4]. Chong, *et al.* [6] discuss the use of adaptive network security mechanisms for systems where service levels must be maintained during attacks. Significant work related to intrusion tolerance systems has been conducted under the MAFTIA Project [8, 24], which built on the results of the earlier Delta-4 project [10]. Lower-level *ad hoc* strategies have been discussed by Haji and co-workers [12]. Bigham, *et al.* [3] have investigated anomaly detection in SCADA environments based on invariant deduction as well as more commonly used *n*-gram techniques. A related approach is discussed by Coutinho, *et al.* [7].

SCADA systems employ multiple types of sensors that are often widely dispersed (especially in the case of the power grid and oil and gas pipelines). Kosut and Tong [15] discuss the application of data fusion techniques to sensors for which Byzantine behavior cannot be ignored. These security concerns apply to sensor data at rest and in transit as discussed by Subramanian, *et al.* [25]. Nguyen and Nahrstedt [19] have addressed the related issue of attack contain-

ment in large-scale industrial control environments using compartmentalization and trust groups.

9. Conclusions

Anomaly detection in SCADA systems has primarily focused on applying general network and host detection techniques. However, the characteristics of SCADA systems, the constraints imposed by real-time industrial environments, and the sophisticated models underlying industrial processes (e.g., state estimator models used for the electrical power grid) require high-level detection approaches as illustrated in this paper. A parallel threat results because attackers with knowledge about process models and SCADA systems can influence or fabricate sensor readings and actuator behavior so that they appear normal to operators. Such manipulations can degrade or disrupt vital industrial processes or force them to operate closer to the margins where a subsequent attack (e.g., a physical attack) could cause significant damage.

Statistical techniques, as described in this paper, are well suited to detecting anomalous behavior in SCADA systems (and critical infrastructure networks, in general). Simplified models and simulations were used in this work to illustrate the main concepts. Our future research will investigate the application of more elaborate hierarchical and composite models. We will also explore the use of multivariate analysis of variance techniques for detecting anomalies in systems with multiple dependent variables.

References

- [1] M. Amanullah, A. Kalam and A. Zayegh, Network security vulnerabilities in SCADA and EMS, *Proceedings of the IEEE/PES Transmission and Distribution Conference and Exhibition: Asia and Pacific*, pp. 1–6, 2005.
- [2] R. Bace, *Intrusion Detection*, Sams, Indianapolis, Indiana, 2000.
- [3] 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.
- [4] P. Bracken, *The Command and Control of Nuclear Forces*, Yale University Press, New Haven, Connecticut, 1985.
- [5] P. Brockwell and R. Davis, *Introduction to Time Series and Forecasting*, Springer-Verlag, New York, 2002.
- [6] J. Chong, P. Pal, M. Atigetchi, P. Rubel and F. Webber, Survivability architecture of a mission critical system: The DPASA example, *Proceedings of the Twenty-First Annual Computer Security Applications Conference*, pp. 495–504, 2005.

- [7] M. Coutinho, G. Lambert-Torres, L. da Silva, E. Fonseca and H. Lazarek, A methodology to extract rules to identify attacks in power system critical infrastructure, *Proceedings of the IEEE Power Engineering Society General Meeting*, pp. 1–7, 2007.
- [8] M. Dacier (Ed.), Design of an Intrusion-Tolerant Intrusion Detection System, MAFTIA Deliverable D10 (Version 4.3), IBM Zurich Research Laboratory, Zurich, Switzerland, 2002.
- [9] D. Denning, An intrusion-detection model, *IEEE Transactions on Software Engineering*, vol. 13(2), pp. 222–232, 1987.
- [10] Y. Deswarte, L. Blain and J. Fabre, Intrusion tolerance in distributed computing systems, *Proceedings of the IEEE Symposium on Research in Security and Privacy*, pp. 110–121, 1991.
- [11] D. Dzung, M. Naedele, T. von Hoff and M. Crevatin, Security for industrial communication systems, *Proceedings of the IEEE*, vol. 93(6), pp. 1152–1177, 2005.
- [12] F. Haji, L. Lindsay and S. Song, Practical security strategy for SCADA automation systems and networks, *Proceedings of the Canadian Conference on Electrical and Computer Engineering*, pp. 172–178, 2005.
- [13] V. Igiure, S. Laughter and R. Williams, Security issues in SCADA networks, *Computers and Security*, vol. 25(7), pp. 498–506, 2006.
- [14] R. Johnson and D. Wichern, *Applied Multivariate Statistical Analysis*, Prentice Hall, Upper Saddle River, New Jersey, 2007.
- [15] O. Kosut and L. Tong, Capacity of cooperative fusion in the presence of Byzantine sensors, *Proceedings of the Forty-Fourth Annual Allerton Conference on Communication, Control and Computation*, 2006.
- [16] T. Kropp, System threats and vulnerabilities: Power system protection, *IEEE Power and Energy*, vol. 4(2), pp. 46–50, 2006.
- [17] E. Murtoviita, J. Keronen, J. Suni and M. Bjork, Visual aids for substation monitoring and security control, *Proceedings of the Third International Conference on Power System Monitoring and Control*, pp. 225–227, 1991.
- [18] M. Naedele, Addressing IT security for critical control systems, *Proceedings of the Fortieth Annual Hawaii International Conference on System Sciences*, p. 115, 2007.
- [19] H. Nguyen and K. Nahrstedt, Attack containment framework for large-scale critical infrastructures, *Proceedings of the Sixteenth International Conference on Computer Communications and Networks*, pp. 442–449, 2007.
- [20] P. Palensky and T. Sauter, Security considerations for FAN-Internet connections, *Proceedings of the IEEE International Workshop on Factory Communication Systems*, pp. 27–35, 2000.

- [21] President's Commission on Critical Infrastructure Protection, Critical Foundations: Protecting America's Infrastructures, The White House, Washington, DC (chnm.gmu.edu/cipdigitalarchive/files/5_CriticalFoundationsPCCIP.pdf), 1997.
- [22] G. Shafiullah, A. Gyasi-Agyei and P. Wolfs, Survey of wireless communications applications in the railway industry, *Proceedings of the Second International Conference on Wireless Broadband and Ultra Wideband Communications*, p. 65, 2007.
- [23] StatoilHydro, The long road to LNG, Stavanger, Norway (www.statoilhydro.com/en/NewsAndMedia/Multimedia/features/SnohvitLNG/Pages/default.aspx), 2007.
- [24] R. Stroud, I. Welch, J. Warne and P. Ryan, A qualitative analysis of the intrusion-tolerance capabilities of the MAFTIA architecture, *Proceedings of the International Conference on Dependable Systems and Networks*, pp. 453–461, 2004.
- [25] N. Subramanian, C. Yang and W. Zhang, Securing distributed data storage and retrieval in sensor networks, *Proceedings of the Fifth Annual IEEE International Conference on Pervasive Computing and Communications*, pp. 191–200, 2007.
- [26] Substations Committee of the IEEE Power Engineering Society, IEEE Recommended Practice for Network Communication in Electric Power Substations, IEEE Standard 1615-2007, IEEE, Piscataway, New Jersey, 2007.
- [27] R. Walpole, R. Meyers and S. Meyers, *Probability and Statistics for Engineers and Scientists*, Prentice Hall, Upper Saddle River, New Jersey, 1998.