Not all SCADA is Equal: Impact of Control Models on ICS Threat Landscape

Antoine Lemay
École Polytechnique de Montréal

Marina Krotofil
Hamburg University of Technology

José M. Fernandez
École Polytechnique de Montréal

Scott Knight
Royal Military College of Canada

There are almost as many ways to implement Industrial Control Systems as there are ways to control industrial systems. This produces a wide-varying range of possible architectures for the SCADA networks that control them. This paper organizes SCADA networks based on these various control architectures in order to evaluate how different control models and architectures can affect the threat model. We observe that distributed control, with control logic in the endpoints, is more susceptible to attacks on SCADA endpoints and that distributed state architectures, with centralized control, is more susceptible to attacks on the SCADA master.

Keywords: Secure control, ICS security, state estimation, control loop

1. INTRODUCTION

The quest for efficiency always drives industry. By increasing the efficiency of day-to-day operations, corporations are able to reduce costs and increase profits. It is no wonder that businesses, including operators of critical infrastructure, leverage efficiency technology such as Industrial Control Systems (ICS). Increasingly, these systems rely on Supervisory Control and Data Acquisition (SCADA) networks to implement industrial controls.

These gains have come at the expense of an increased risk of cyber attacks. The discovery of Stuxnet (Falliere et al. 2011; Langner 2013), a worm designed to perform sabotage on the Iranian nuclear enrichment program, acted as a watershed moment that spurred research to secure ICS against cyber attacks. Unfortunately, research in the field suffers from a lack of publicly available information on ICS in general because of the sensitivity of this information for the operators.

A number of publications have tried to fill the gap by presenting reference architectures for SCADA networks. However, these papers present different, and sometimes even conflicting, reference designs for SCADA networks. For example, in its “Guide to Industrial Control Systems” (Stouffer et al. 2014), the US National Institute of Standards and Technology (NIST) presents a simple control network with Programmable Logic Controllers (PLC),

Human-Machine Interface (HMI) and a control network on a LAN. On the other end, Hull et al. (2013) present a multi-tiered network based around a central server. Wu et al. (2005) go even further by incorporating the SCADA network to other systems such as energy management and marketplaces.

The abundance of architectures in the literature can make life difficult for researchers entering the field. This paper aims to contribute to the community by explaining the reason for diversity among different architectures. We present the operational differences between SCADA systems following a distributed control architecture vs. systems using a distributed state architecture. The paper will show that these differences have impact on the threat model, and the corresponding defensive strategies.

The paper will start by a brief introduction to the basics of control theory followed by the description of control models. Then, a study of the implication of those differences for the threat model in SCADA networks is presented. Finally, we end the paper with a conclusion.

2. BASIC CONTROL

SCADA systems are networked ICS. As such, they implement traditional control principles such as the basic control loop. In a basic control loop, the operator of the system sets a desired state. For

Figure 1: Basic control loop

example, an aircraft pilot sets a desired altitude. The current state of the systems is then evaluated against the desired state to estimate the gap, typically referred to as the measured error. The altimeter might show the pilot that he is 1000 feet under his desired altitude. The measured error is then sent to a controller, be it a microchip, a PID controller or a human brain, which will interpret this gap and generate a set of alterations, or inputs, to the state of the system to reduce that error. The pilot might pull on his joystick to cause the plane to climb. This generates a new state which is hopefully closer to the desired state. The values of the sensors for this new state can be fed back in the process to generate a new measured error and the loop can be repeated until the measured error is close enough to zero. In our example, the aircraft pilot keeps monitoring the altimeter and applying appropriate controls (elevator, engine power, trim) until he reaches the desired altitude. This basic control loop is illustrated in Figure 1.

SCADA systems also follow this model. The desired state is provided by the operator in the form of set points, i.e. operational parameter values at which the system is forced to operate. This is then interpreted by a controller, for example through a ladder-logic program on a PLC. The PLC converts the set point in inputs to send to actuators and alter the state of the system, for example sending a command to a valve to open. The state of the system is continuously evaluated through measurement points (sensors), for example a water level gauge, and the measurements are fed back to the controller to close the loop.

Looking at the Purdue Enterprise Reference Architecture (Williams 1994), there is also a clear hierarchical processing of data with each level operating within its own loop. The direction of the process data flow is bottom-up, while management and control data flows from top to bottom, with each layer adding latency. “Real-time” strategic process monitoring and management is about long-term changes, in order to monitor and correct long-term drifts. Supervisory control is about medium-term strategic monitoring and tactical control (process supervisors and process managers). The focus of operators, and controllers, is short-term

transactional data for regulatory control of a process and to address time critical operational problems.

However, even if all control systems follow this basic control loop, there are still significant differences between implementations.

3. CONTROL MODELS

Applications of ICS are varied. In a sense, each physical system requires its own control implementation. In this section, we will study two implementations at the opposite end of the spectrum in order to illustrate the effects of the control approach on the attack model. However, it is important to keep in mind that modern industrial processes may follow hybrid versions of these models.

3.1. Distributed Control

A good number of industrial processes are operated as a series of functionally independent discrete steps that can be individually controlled. This can be either because they produce a batch of products and are then retooled for the next product, i.e. operating as a cascade of processes requiring few interactions, or because they are complex processes that are only minimally controlled by a control system.

An example could be an assembly line where each robot performs a limited number of tasks with its local control loop. The arm is moved until it is aligned with the rivet hole and punches the rivet in. The piece is then moved to the next step. In this context, the local controller can collect all relevant information and alter the system locally even if a central system will still collect information from all the processes to give operators a global picture. Thus, the entire system is operated as a collection of small subsystems with independent control loops. We will dub this model the *distributed control model* for SCADA networks.

SCADA networks following this model are oriented toward ladder-logic types of architecture. In this architecture, the endpoints of the SCADA network, the controllers, host the programs that are ultimately implementing the control loops, even if the desired state of the system is fixed by operators through a HMI hosted on the Master Terminal Unit (MTU). Should changing set points be insufficient, engineering stations allow plant engineers to alter the logic on the controller by uploading new ladder logic. Figure 2 illustrates a SCADA network in a distributed control environment.

3.2. Distributed State Estimation

Not all systems can be controlled as a series of independent steps. The modern electric grid

Figure 2: Distributed control model

is an example of a system too complex for this kind of control. Shutting down a line in a distribution centre might create impacts on the power generation hundreds of kilometers away. As such, these systems need to be controlled centrally. Similarly, complex, continuous chemical processes require global knowledge and are therefore usually controlled across the entire plant. This makes estimating the state and generating an input that will push the state toward the desired state a complex task.

To address this complexity, it is necessary for the equipment implementing the control algorithm, in this case the MTU, to have information about the global state of the system. Therefore, the state estimation model does not heavily use the ladder-logic in the endpoints, but instead uses the endpoints to perform a *distributed state estimation*. Measurement points collect information about the state of the physical system and send this information to a central database in the MTU. A state estimator, taking as input a vector containing the states of all pieces of equipment, can generate an estimation of the current state of the controlled system and calculates the measured error. This is then sent to a central system that implements control, for example an Energy Management System (EMS) in the electricity sector. This central controller can then compute the inputs required to move the controlled system toward the desired state in a

vector of commands to be sent to individual actuating equipment. Figure 3 illustrates the state estimation model.

3.3. Hybrid Systems

The complexity of modern industrial systems is such that they often require different control paradigms to be combined. For example, while the grid can be controlled centrally, it may still be pertinent to also implement a local control loop for surge protection. Similarly, modern field instrumentation may communicate between themselves to create a distributed state for a chemical process. These systems use a hybrid control model with varying degrees of distribution of control and distribution of state.

4. IMPLICATIONS FOR THE THREAT MODEL

The choice of a control model (distributed state or distributed control) alters the relative importance of SCADA components. As such, it is not surprising that the impact of attacks varies depending on that choice. This section presents an analysis of the impact of an attack on a controller, the falsification of the process measurements, an attack on the MTU and the impact on the timing of the attacks for both distributed state and distributed control models.

Figure 3: State estimation model

4.1. Impact of Attacks on Endpoints

Compromise of a controller is a major risk factor in many SCADA systems. In the scenario where the attacker hacks a controller or RTU to gain local control and does not try to expand his foothold on the network, the consequences of a successful attack are very different in the distributed control and distributed state model.

In systems implementing distributed control, the scope of the control loop is limited to a single operating unit. By compromising a controller, the attacker has access to all the set points, state information and inputs required to alter the state of that unit. As such, he has complete control over the portion of the process that is under the responsibility of that particular controller. This gives a great amount of power to the attacker, especially in the case where the process is just one step in a chain. The flip side of this property is that multiple compromises are required to affect the entire process, if the processes are sufficiently independent.

In distributed state SCADA systems, while the compromise of an end point might allow an attacker to cause some damage locally, he is unlikely to be able to cause much damage to the overall system without attacking other nodes. The centralized controller is usually able to compensate for a single failure provided the physical system possesses the built-in resiliency. For example, an attacker taking a substation offline is likely to cause some local power outages, but the central control logic will reroute the power through other parts of the grid. This greatly limits the scope of damage an attacker can cause without having to expand his attack to other nodes in the network.

Under these circumstances, because of the higher impact of attacks on network end points, greater consideration should be given to preventing attacks against PLCs in networks implementing distributed control.

4.2. Impact of Measurement Falsification

Injecting packets with false measurements, altering the measurements of legitimate packets or changing the way sensor measure the state of the process

represent a great risk because it alters the evaluation of the physical system's state. Therefore, it impacts the control of the system directly.

In systems implementing distributed control, because of the relatively limited complexity of control islands, it is usually easier to determine what false measurements will produce the desired results for the attacker. For example, if the attacker wants a water tank to overflow, sending a measurement showing the tank is half full will induce the controller into thinking more water is needed. The flip side is that, as shown in the previous section, the scope of the impact is limited to a single operating unit unless physical dependencies exist.

For systems with distributed state estimation, the analysis of the impact of false measurements is much more complex. Because the data is processed as part of a global state estimation, the impact of false measurements depends on the significance of specific measurements in the global state. The falsification of a significant value can create far reaching impacts because of the central nature of the controller. Alternatively, the alteration of a non-significant value might not even cause an impact. For example, the aggregate power output of a plant is probably more relevant to the state estimation than the voltage on a distribution line.

The complex nature of the central control algorithm also renders the job of the attacker more difficult. In order to induce the desired alterations of the target system, attackers are required to fully understand the nature of the target system and its control configuration. They are also likely required to possess their own estimation of the global state of the target system suggesting the need for an extensive compromise of the SCADA system (from a large percentage of controllers or from a high impact target such as the MTU). This is likely to require significant reconnaissance on the part of the attackers about the standard operation of the system.

Consequently, an increased attention to the confidentiality of SCADA operation parameters, the disclosure of which is typically considered to have a low impact, is probably warranted in distributed state systems.

4.3. Impact of Attacks on MTU

Because the MTU is the central server with which all controllers communicate, it represents a major target for attackers. In all cases, gaining administrator access on the MTU typically allows the attacker to be able to alter set points across the entire system. However, for attacks on the MTU that do not result in

complete compromise, significant differences in the impact of attacks can be observed.

In SCADA systems based on distributed control, the main role of the MTU is to provide operators with a global visibility of the industrial system and to provide a convenient interface for set points. So, unless there is a significant modification in the operating characteristics of the industrial process that would require the operator's attention, disruption of the MTU is more of an inconvenience. Operations may continue normally using only local controllers and current set points. This property may even allow defenders to shut down the MTU to counter attackers.

In the case of SCADA networks using distributed state estimators, any disruption of the MTU prevents adequate state estimation, and therefore completely prevents the control of the industrial system. Because of the inherent complexity of industrial processes requiring distributed state estimation, the loss of centralized control might completely disable the industrial system. Much like modern fighter jets that require automated piloting systems to remain airborne, the modern grid requires the centralized control algorithms to operate efficiently.

Under these circumstances, the protection of the MTU of SCADA systems using a distributed model is critical. Coupled with the limited impact of individual endpoint compromise, it is clear that the MTU should be the focus of much attention from the defenders.

4.4. Impact on Timing of Attacks

Attacks on ICS that intend to cause physical impacts need to adjust their timing parameters based on the physical system. For example, an attacker wanting to destroy a pipe with a water hammer has to determine the right time to operate the valve, round-trip time of the shock wave and coordinate these two timing parameters to achieve desired destructing effect.

In their work, Krotofil and Cárdenas (2013) have already shown that the timing to bring the system to the state desired by the attacker depends on a large number of parameters including the relationship between interdependent physical parameters and the way they are controlled in particular, the configuration of the control loop which includes the choice of the manipulated variables, control algorithm type and controller tuning parameter. Obviously, these parameters vary wildly from one process to the next. However, in most cases, systems using distributed state architectures have more complex control loops than systems using distributed controls. As such, it falls within reason that calculating the precise timing for those systems is more difficult for attackers.

5. CONCLUSIONS

In this position paper, we described that there are significant architectural differences in SCADA systems depending on the type of control model employed. In distributed control systems focused on local control, there is an increased reliance on SCADA endpoints, both controllers and field equipment (sensors and actuators). On the other hand, in distributed state control process measurements are collected for centralized state estimation and control. This dichotomy implicates that the impact of attacks on SCADA endpoints is higher in distributed control systems, while attacks on central servers can cause more havoc in distributed state estimation systems.

This has implication on future research in the domain of ICS cyber security. First of all, more in-depth investigation of the threat model impact could provide further guidance, notably, the effects of the control architecture on the timing parameters of attacks. Additionally, a better understanding of how each industry sector implements control is critical for SCADA security researchers to best model threats and design effective countermeasures (attack detection and response).

REFERENCES

- Falliere, N., Murchu, L. O., and Chien, E. (2011, Feb.) *W32.stuxnet dossier*. Mountain View, CA, USA: Symantec Corporation, Tech. Rep.
- Hull, J. et al. (2013) Staying in control: Cybersecurity and the modern electric grid. *IEEE Power Energy Mag.*, 10 (1/Jan.–Feb.). 41–48.
- Krotofil, M. and Cárdenas, A. A. (2013) Resilience of process control systems to cyber-physical attacks. In: Volume 8208 of LNCS. *Secure IT systems*. Berlin, Germany: Springer. 166–182.
- Langner, R. (2013, Nov.) *To kill a centrifuge: A technical analysis of what Stuxnet's creators tried to achieve*. Hamburg, Germany: Langner Communications. Tech. Rep.
- Stouffer, K. et al. (2014) *NIST special publication 800-82, revision 2: Guide to industrial control systems (ICS) security*. Gaithersburg, MD, USA: National Institute of Standards and Technology.
- Williams, T. J. (1994) The Purdue enterprise reference architecture. *Comput. Ind.*, 24 (2–3). 141–158.
- Wu, F., Moslehi, K., and Bose, A. (2005) Power system control centers: past, present, and future. *Proc. IEEE*, 93 (11). 1890–1908.