Simulating and Detecting Measurement Attacks in SCADA Testbeds

Brandt Reutimann and Indrakshi Ray

Department of Computer Science
Colorado State University
Fort Collins CO 80523

Abstract. Industrial control systems (ICS) are an increasingly target-rich environment for cyber-criminals, terrorists, and advanced persistent threats. Previously researchers have looked into various types of ICS systems including smart grids, gas pipelines, and manufacturing centers to understand how they can be compromised by cyber threats. However, the manner in which ICS systems are attacked is domain-dependent. Test beds are a necessary tool to model the domain and understand these attacks. This work focuses on building a virtual test bed for gas systems and how it can be used to simulate and detect measurement attacks. Such a test bed provides opportunities for cyber security researchers and domain experts to model, simulate, and understand the behavior of a real-world system and its response to cyber attacks.

1 Introduction

Industrial control systems (ICS) are an increasingly target-rich environment for cyber-criminals, terrorists, and advanced persistent threats. Control technologies have become more connected both to the internet, as well as in their internal networks. As a result, the physical systems underlying these control topologies have also become more accessible to malicious actors. Additionally, the critical nature of systems such as power plants or gas pipelines can make them a high priority target for those wishing to ransom, extort or cause damage to a nation's critical infrastructure. Hence, the security of these systems is tremendously important.

Most work in ICS security focuses mostly on a handful of critical embedded systems. When it comes to security this centers around the embedded control systems that operate power grids, gas pipelines and manufacturing centers. These systems can be incredibly sophisticated from both an engineering and IT perspective. Modern ICSs are typically operated by large networks of interconnected controllers and operators. These networks are colloquially defined as supervisory control and data acquisition systems (SCADA). Hundreds to thousands of process logic controllers are connected to SCADA network and historian servers. Operators use these networks for maintenance, monitoring and system control. Sometimes these SCADA systems may even be connected to corporate networks where typical IT security can become just as much of a concern as SCADA system security.

The size and temperament of large SCADA systems leaves several vulnerable attack surfaces. For instance, controllers in large scale SCADA systems tend to be heterogeneous as it can be logistically infeasible to replace or update all of the controllers at the same time. The switching costs involved in replacing or modifying controllers can also be high as gas and electrical systems typically must maintain a availability throughout their life cycle. As a result, in any SCADA system it is likely that there can be entry points for a clever attacker who can take advantage of the "weakest" links in the system. It follows from this that most defenses in SCADA systems security are to isolate the SCADA network from the outside world. This is an extremely effective strategy but not impervious. Often attackers can find ways to tunnel into the SCADA network either through a corporate network, or through other clever strategies. The most popular example of exploiting a SCADA system and its controllers is the Stuxnet worm used to attack the Iranian nuclear program. As an air gaped system, it is commonly believed that the Stuxnet worm was able to infiltrate its prey system through flash drives inserted into the system by an unsuspecting user. [1]. However Stuxnet is just one of many high profile instances of a SCADA system attack. A 2018 US Department of Energy report details the events of 22 high significance attacks between 2000 and 2018 [2]. For instance in one such attack, a man used a laptop and radio equipment to dump millions of gallons of sewage into Queensland waterways.

As SCADA systems are large, extremely complex, and critically important for our daily lives, studying them for cyber security is of obvious value. However, the process of studying these systems is challenging and sometimes potentially dangerous. A large portion of data on SCADA systems is proprietary, leaving many researchers guessing about how to best model and design these systems. In addition, studying live systems can cause safety hazards such as gas pipelines and electrical grids being put under unnecessary stress. That's why the work of studying these systems requires the effort of not only expert computer scientists but also engineers and safety and security specialists. The goal of our work, then, is to demonstrate how a SCADA system can be studied safely and accurately using simulation techniques. By using a simulated virtual test bed, we can easily iterate and redesign our simulations to be as accurate as possible. This also allows us to model cyber attacks without ever harming a real system.

In this paper we experiment on a simulated SCADA system built using a simulator designed by our research team. We propose a gas system model based on public metrics available from Colorado energy companies. Then we apply a suite of attacks that manipulate the sensor readings and feedback given by the simulated controllers in our system. Through these manipulations we are able to observe failures in the simulated gas model. However, we also show that gas systems can react slowly to attacks. This means that attacks on gas systems must be sustained substantially longer than their counterparts on electrical systems.

1.1 Our Contributions

The main problem, of particular interest to us, is understanding what happens when the control layer starts lying to the supervisory system about the state of the physical world, what has been coined in prior SCADA literature as a measurement attack [3]. Though data manipulation and integrity attacks are not unique in cybersecurity, what makes measurement attacks especially dangerous in ICS is the control-feedback loops common to these systems. These control-feedback loops combined with the typical lack of authentication makes these attacks particularly interesting in SCADA systems.

In this work we make several contributions. The first is introducing a model inspired by simulating the Colorado gas system. At time of writing, this model is one of the first that we know of for use in cybersecurity testing. This model is built using an in-house simulation technology developed by our research team. We then use this model we've developed to perform two measurement attack experiments [3]. The first is a single point of failure experiment that explores the impact of compromising a single controller in a large system, while the second explores the impact of manipulating several controllers using a more advanced knowledge of the system mechanics.

The rest of the paper is as follows. Section 2 discusses prior work in the simulation of SCADA systems as well as experimental attacks on these systems. Section 3 describes the design of our gas system model, the way we set up our cyber-attack examples, and how data collection is done. Following experiment setup there are two experiments in Sections 4 & 5, the first is a single point of failure experiment while the second is a more sophisticated measurement attack. Next, in Section 6 we discuss the results of these experiments and what they mean for cyber-attacks on gas systems. In Section 7 we discuss future directions of our work. Finally, in Section 8 we conclude the paper with final thoughts on this work.

2 Related Work

In order to ensure the security of SCADA systems blue teams, white hat hackers, and those wishing to employ cyber-defenses must be able to anticipate and defend against possible attacks. The best way for them to do this is to probe these systems for flaws on their own. Unfortunately, testing ICS can be a complex, difficult and sometimes dangerous process. Traditional methods of testing these systems usually include some sort of traditional engineering validation methods. However, when testing the cyber-security of these systems previous research has shown that there are some serious risks. In one report, the use of ping sweeps caused a robotic arm to swing on a factory floor and in another case caused a system failure that resulted in over $50,000 worth of damage to equipment [4]. As a result of the risks of cyber-security and non-physical testing, a recent paradigm in research has focused on simulating industrial control environments – described in this paper as SCADA systems. Researchers from Mississippi State University have successfully created several physical SCADA systems for their

test bed [5]. This test bed includes a water tank, water tower, small gas pipeline, factory conveyor belt, and smart grid system. The same authors were also able to create virtual models of their water tank, and gas pipeline systems. These virtual models were validated against the physical implementations of the systems [6]. Although these simulations models have actual physical implementations, the physical test beds are extremely simple. These models have single feedback loops and simple control sequences which do not accurately reflect the real-world. As a result, it can be hard to do any security testing that relies on attacking the control system algorithms. Instead these models focus on attacks on the network or physical components of the test bed.

Prior research has shown successful tests of various network-based attacks on virtual SCADA test beds [3]. Several other attempts have been made to model large scale behavior of SCADA systems from a test bed or network-based perspective [7, 8]. However, the aforementioned research either only models the network behavior of SCADA systems or is limited to a small scale physical test bed. As a result there is a shortfall in this previous work. A SCADA system model is either lacking in granularity by only focusing on networks, or lacks larger scale perspective from only focusing on physical control systems. In addition, systems that do have physical implementations are limited by the fact that most test beds only have the resources to acquire so many real components. Survey research of this problem has shown that hybrid and hardware in the loop simulations tends to lead to the best of both worlds [9]. However, these solutions are not as cost effective as software modeling and virtualized systems. This shows that scalability, accuracy, and fidelity to the real world are critical components of these systems. Not only for the development of test beds that can be used for security testing but, also for the design of system simulations from an engineering standpoint. In some sense, the use of simulating full scale SCADA systems can also be considered as a tool for designing and reasoning about large scale gas pipelines and electrical grids.

In addition to studying SCADA simulation in general, the interactions between gas and electrical systems has also been a popular area of study. Failure of gas systems can cause failure in electrical systems and vice versa. Some literature has tried to tackle this interaction by using mathematical models to show the interdependencies of the systems described as a mathematical optimization problem [10, 11]. Other work has expanded upon this by creating models to simulate this interaction. These simulated systems have shown that failure in a handful of gas lines can lead to cascading failures in a corresponding electrical system [12]. However, these models can only consider a limited set of variables as opposed to a high fidelity simulation like the ones we will use in this paper. Previous work in SCADA simulation has focused mainly on the operation of single facilities or models of electrical systems [13].

Although SCADA simulation is a topic that does have a wealth of literature behind it, we believe our work fits into a unique role. At the time of writing we are not aware of any work in SCADA simulation that studies gas systems or their effects on the electrical systems. Especially when it comes to cyber attacks

and vulnerabilities. Although this work does not cover the interaction between electrical and gas systems in depth, we do explore vulnerabilities in simulated gas systems that have the potential to harm a hypothesized electrical system. In addition to being some of the first simulation work to explore gas transmission systems, our work also explores measurement attacks in a new and unique way. Based on our review we have only seen a limited body of work on measurement attacks and how they can be applied to large scale systems. The following research explores this research question by testing several measurement attacks on gas pipeline models. Finally, we believe that our design for simulating SCADA systems provides a unique and simplified architecture that can be expanded in future work to create modular simulations that can be used to study all kinds of cyber physical systems. To our knowledge, most work before this has mostly focused on creating simulations for specific types of systems (mainly electrical) and has neglected a method for creating a more expansive simulation model.

3 Colorado Gas Model and Experiment Setup

3.1 Model structure

The Colorado scale gas model use for our experiments is a substantially sophisticated gas model that requires special design considerations. The main objective in designing these larger models is to make them realistic enough that the system fails because of a cyber-attack and not because the engineering inherent to the model is poor. In order to do this our team consulted with engineers at the Colorado Powerhouse as well as gas system experts to define what constraints the model must meet. The main concern was getting data to scale our model around. However, the process of acquiring accurate data on control systems like the ones we study is very challenging.

For this system we combine data from several sources, as well as using some principles of gas dynamics in order to determine accurate measurements. We used a Kinder Morgan system map to discover the general area of where gas pipes and power plants exist in the state of Colorado [14]. We then extrapolated this map to the nearest large cities and determined rough distances for placement of the gas pipe. We then placed moderate sized power plants near medium sized cities in Colorado and a large power plant in Denver. We also sprinkled a handful of distribution loads to add complexity to the system. Once we determined the placement of power plants, compressors and distribution loads we referenced Xcel Energy in order to discover the peak load capacity in Megawatts for power plants around the state [15]. We converted these peak capacities from Megawatts to required gas in kg/s using the gas heating values from experiment 1.

Once we determined the peak mass flow gas loads for each power plant (in kg/s) we were able to design the diameter of the gas pipes to meet this demand at a nominal pressure of 800 psi using Bernoulli's Equation (ignoring gravity / height differences). As gas loads in this system and real life are very high the pipes have to be very big in order to distribute as much gas as needed. Typically gas pipes come in nominal sizes that might not exceed 36 inches in diameter,

and when engineers need to distribute more gas along a line they will lay several pipes in parallel. Currently our system does not consider this and uses a single large pipe. This may have adverse effects on some of the properties of the gas. For example, there is more friction when using several small pipes than there is in using a single large pipe. In addition, in smaller pipes gas may be flowing at a higher velocity which means that temperature changes can be more dramatic.

3.2 Simulink and Colorado State Simulator

After we determined the structure of our gas model and worked out some basic numerical constraints we simulated it using our SCADA simulation tool. The SCADA simulation tool developed by our research team is a controller and software simulator that interfaces with MATLAB's Simulink to create a workable interface for developing virtual PLCs and integrating hardware in the loop. These real and virtual PLC's (process logic controllers) can then be communicated with using their native SCADA protocols (typically MODBUS - Modicon Serial Bus Protocol). This allows us a window into Simulink simulations using traditional controllers and control system software. As the design considerations and development of this tool are a topic of their own we do not describe them in depth in this paper.

3.3 Automated Switch-Off

Gas systems can become extremely dangerous when pressure or temperature drop rapidly. Rapid drops in pressure can prevent gas delivery to a system, while temperature drops can lead to dew pointing and actual solidified particles traveling through a pipe. In prior simulation experiments we used an external control system to switch off a power plant when its pressure or temperature reached a certain point. For the Colorado gas model we wanted to keep the switch-off control inside of the process model to represent engineers at a power station tripping the station off when the gas reached a certain set pressure or temperature. In the Colorado gas system we add a MATLAB function block that sets the power stations load to zero and switches it off for a number of seconds specified by a parameter at the beginning of the simulation. For this model we used 3600 seconds (one hour) as the shutoff duration for our power plants. This time represents how long after a shutoff it takes the power plant engineers to run through their safety procedures and ramp the system back up to meet the current load. The goal of the trials in our experiments are to get these power plants to trip their automatic shutoff by dropping the gas pressure in the system.

3.4 Dynamic Load Distribution

The Colorado gas model uses a dynamic load distribution algorithm that allows it to simulate a system wide load that different power plants in the system have to unite to meet. In previous experiments every load was specified individually,

whereas in this experiment a single load can be specified in order to model a realistic scenario. The advantage of this approach is that when one power plant in the system trips off, the load switches onto other plants in the system. This is effective for showing scenarios where the loss of a power plant can possibly cause a cascading failure as demands get increased very rapidly on other parts of the system.

Each power plant in the system is specified with a load capacity in kg/s. When determining the load of a given power plant you use it's ratio to the total system capacity to express its proportion of the current load. Let the plant capacity be p_c, the total system capacity be s_c and the current load at time t be l_t. Then the load on the a given plant, i, will be $\frac{p_{c,i}}{s_c * l_t}$. Also we must consider a scenario where the total load on the system exceeds the capacity of the power plants available. In this case, each plant will output at its maximum load and the system's demand will not be met. This scenario can be considered one where the gas system is not generating enough electrical power to meet demand.

3.5 Control Design

The control system for experiments in this paper is substantially more complex than our prior experiments. The compressor stations in the system are designed so that the power capacity of the station can be modified. When compressing gas, the compressor station inspects a value δ_p which is the difference in pressure between the desired set point (in the case of this model it is 800 psi) and the current pressure reading. As δ_p increases, more power is required to compress gas to meet the pressure differential. In order to make the system realistic there has to be a limit on the amount of power that a compressor can use to achieve this goal. This limit is specified in megawatts, with the base value being 5 megawatts. When the system requires more power than is available, δ_p is modified to the maximum compression available for the current maximum amount of power. This plays into the control system as the operators of the system will want to minimize the amount of power being used by compressor stations whenever possible.

The control algorithm for the Colorado gas model increases the power available to upstream compressors based on the current pressure reading at a downstream station or compressor. The controller iteratively checks the state of the system and updates the power available to each immediate upstream compressor. Power is updated in 5 megawatt increments, if the pressure is less than 750 psi then power is updated +5 megawatts while if power is at 800 psi available power is decreased -5 megawatts to save energy. Upstream compressors are identified by consulting a directed graph that represents that gas flow in the system. Each node in the graph is either a power station or a compressor. Every power station is a leaf node by default, as gas does not flow through power stations to other nodes in the system.

3.6 Gas System Scenario

For the purpose of the next two experiments the Colorado gas model is set up in a specific scenario. The scenario chosen represents a time when the gas system is put under a high level of strain. This is important because instances where the system is under high strain are great opportunities for a cyber attack. The current scenario is a 3 day simulation where for the first day the system is running under a moderate amount of load. Although slightly contrived, we believe this scenario is a good starting point for demonstrating the value of SCADA system simulation. The scenario shows the potential impact of data manipulation in a SCADA system. These situations can happen when renewable energy like wind is carrying a large portion of the electrical demand [16]. Then suddenly, on the second day we have a loss of this wind power. As a result, the natural gas power plants have to ramp up to a high load to generate enough electricity to maintain the current usage. After a 12 hour period the required load drops back down representing either the wind power coming back on or the demand for electricity going down (see Figure 1a for a profile of the system load in this scenario). It is important to note that when these power plants ramp up the demand for natural gas throughout the system also dramatically increases. If the compressors in this system are not adjusted to meet demand we can see power plants start to fail as they are not being delivered gas at a high enough pressure (see Figure 1b for an example of what happens when the control system is unresponsive). Thus the goal of measurement attacks in this scenario is to prevent the proper control actions from being taken during this window of high strain, and therefore causing failures of power plants in the system. In the following experiments we say an attack is successful if it causes a power plant to trip off within the window of high strain on the system.

3.7 Data Collection

Data will be collected in several places of the simulation to demonstrate consistency and to show the effects of the compromises in the controllers. The Simulink model logs all data during the simulation, and we will also track data that is coming from PLCs to the simulated operator. This data will be marked on the operator side. As a result we will have the operator query both data from the simulation PLCs and from Oracle PLC so that there is a copy of authentic data as well as a copy of compromised data. We can compare these two sets of data in order to show how the real state of the system is affected by the control actions taken on the compromised data. As a result, we have the capacity to demonstrate that data received by the operator is the same as the data in the model, as well as showing the compromises are affecting the real state of the physical model.

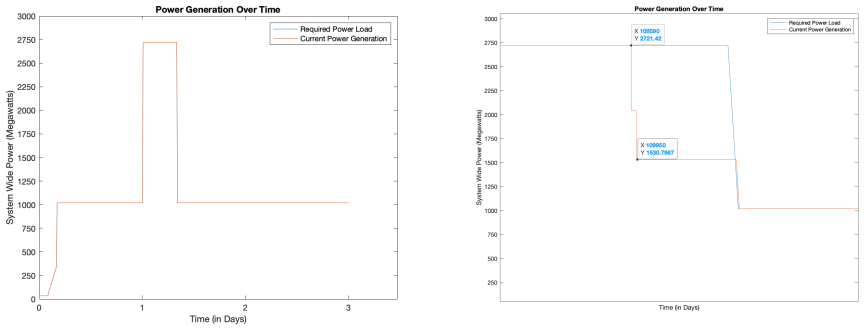

(a) (b)

Fig. 1: (a) Under a period of high stress there is no loss in power if the system corrects properly. (b) Rapid loss of power generation capability can occur without control system intervention.

4 Experiment 1: Single Point of Failure

4.1 Experiment Setup

For this experiment we explored the effect of a single point of failure on the system during the 12 hour window of high strain. When gas loads on the system ramp up pressure starts to dip at all the power plants as they start evacuating their upstream gas lines. When this pressure dips the controller should notice and make more power available to upstream compressors. Therefore, for attacks in this experiment we compromised power plant controllers to falsify their pressure readings. We hypothesized that compromising pressure readings at these power plants would prevent upstream compressors from ramping up and therefore preventing the power plants from getting the gas pressure they need to continue operating.

The compromise used for pressure readings was extremely simple. No matter what the current state of the system is, the compromised power plant will always report 800 psi. For this experiment we ran 5 trials, applying this compromise to each power plant and rerunning the simulation. We can say the attack was successful if a power plant shut offs in the 12 hour window of high strain. The goal of this experiment was to determine if a single compromise could cause a catastrophic effect in the system.

4.2 Results

The single point of failure trials showed interesting results. For 4 out of the 5 trials we saw that even though a measurement attack was taking place the compressor immediately upstream of that attack might not be affected. This is because there can be multiple downstream entities showing low pressure and so

a single compromise will not prevent the controller from updating the pressure at the compressor. For power plants that were isolated the impact of this measurement could have disabled their compressor but since the compressors along the main lines in the simulation were still running there was enough gas being pushed through the system to prevent a failure.

The one trial that caused failure did provide interesting insights. When applying the compromise to the Fort Collins power plant we saw that the Fort Collins compressor was temporarily disabled and this caused the Fort Collins plant to go offline (Figure 2a shows the differential between the falsified pressure reading and the actual reading). In addition, the increased demand brought onto the auxiliary line from Fort Morgan also caused the Fort Morgan plant to trip near the end of the 12 hour window (see Figure 3a). An interesting take away from this trial is that Fort Collins compressor is a critical point in the system as it feeds gas to a lot of major power plants downstream of it. If the Fort Collins compressor does not meet the demand of the downstream plants it can cause increased load on the auxiliary lines.

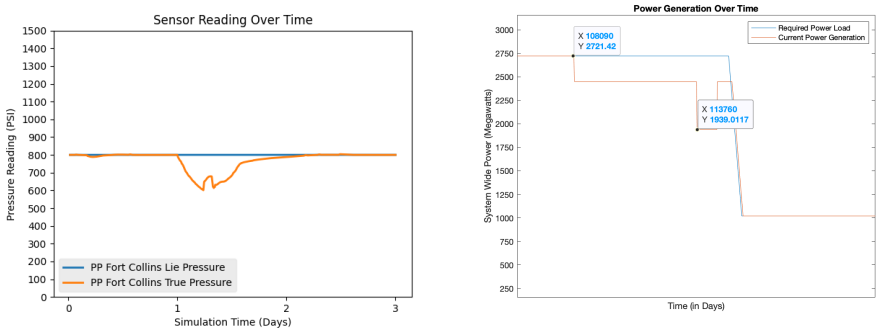

(a) (b)

Fig. 2: (a) Shows the difference between the actual and false readings at the Fort Collins plant. (b) Shows a drop in the ability to meet power demand by the Fort Collins and then Longmont Compressors.

5 Experiment 2: Sophisticated Measurement Attack

5.1 Experiment Setup

For the second experiment we explored the impact of compromising multiple sensors in the system in order to cause a rapid failure. The objective of this experiment is to prevent a compressor from ramping up when it needs to while encouraging other compressors to ramp up when they may not need to. In a way, this attack is similar to the attack described in experiment 2 but slightly more

sophisticated. As we are using falsified readings in an attempt to manipulate the actions of upstream compressors.

For this experiment we are compromising three sensors in the system. Based on experiment 2 we know that the Fort Collins compressor is a particularly critical point in the system. So we compromised the immediate downstream neighbors to this compressor, the Fort Collins power plant and the Longmont compressor. We set both the Fort Collins power plant and the Longmont compressor to always read 800 psi regardless of the current actual pressure. In addition, we marked the Denver power plant as reading low so that the Denver compressor would ramp up, pulling gas down through the path from Fort Collins to Denver. The objective here is to pull gas down the line while preventing the Fort Collins compressor from ramping up to meet the new demand. This should hopefully empty the gas lines and cause failures in the system. The resulting disparity between the true readings and the compromised readings can be seen in Figure 3b.

5.2 Results

The second experiment posed interesting results in the realm of measurement attacks. The compromised high readings at the Fort Collins power plant and Longmont compressor prevented the Fort Collins compressor from ramping up during the period of high stress. In addition, the controller offered an increased power to the Denver compressor to compensate for the falsified low pressure reading. This caused the pressure in the Fort Collins to Longmont line to drop (see Figure 3a) as well as causing an increased demand on the auxiliary line from Fort Morgan. As gas is evacuated out of both the lines coming into Longmont pressure drops rapidly in both the lines. As a result, both the Fort Collins and Fort Morgan plants tripped off in rapid succession of one another. This led to the gas and power generation system not being able to meet its total demand. The total loss of capacity is measured in kg/s of gas load, but translates roughly to a loss of about 800 MW within a period of approximately 5 minutes (see Figure 4b).

6 Discussion

The results of the performed experiments give an insight into the potentially dangerous nature of measurement attacks. In the first experiment we explored how a single point of failure can affect a complex system. The single point of failure work seemed to show that although at critical points in the system there could be a catastrophic effect, more isolated areas tended to show little impact on the overall system. One insight that can be derived from this is that defense in these types of systems may not have to be completely blanketed. It may be sufficient to secure the system from catastrophic failure by identifying critical points and applying efforts to maximize defense at those points.

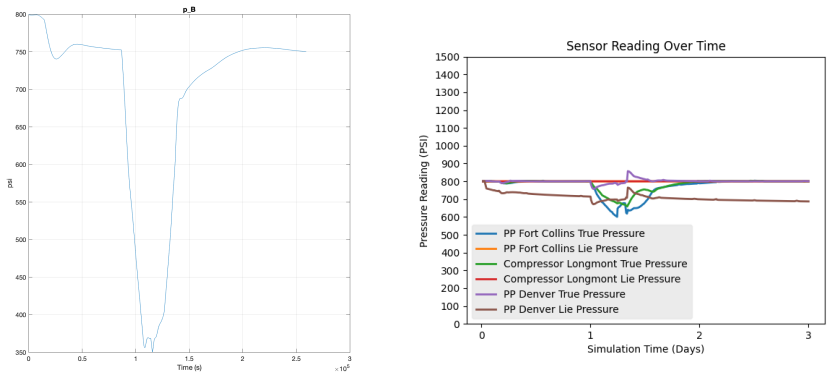

(a) (b)

Fig. 3: (a) Pressure in the Longmont gas line dips under high stress. (b) Real and compromised sensor readings.

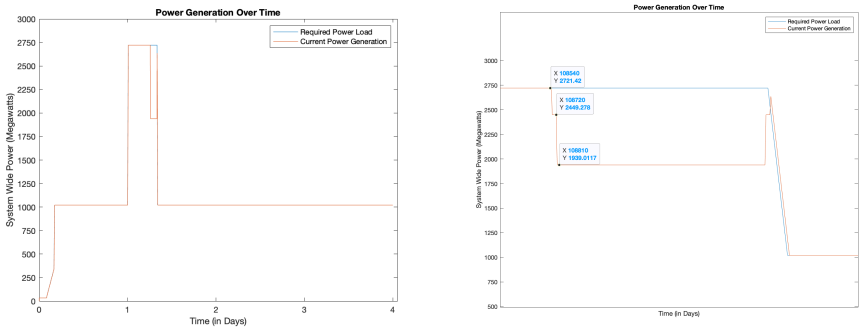

(a) (b)

Fig. 4: Load profiles from the second experiment show that power plants fell offline and failed to meet demand.

The second experiment was extremely interesting and slightly frightening in that it demonstrated that an attacker with intimate knowledge of system dynamics could use measurement attacks to manipulate it's control response. Through measurement attacks the system can be put into a state that causes failures that otherwise would not happen with normal control behavior. One interesting insight from this experiment is that as more sensors in the system become compromised the attacks can become increasingly complex. This likely means that they will also be harder to detect. However, for these types of attacks to happen the attacker must understand how the control system operates, the gas system topology and have knowledge of when the system is in a period of high strain.

Finally, the results of the experiments also expose that the nature of gas systems in response to cyber attack may be very different from that of electrical systems. Failures in electrical systems can occur very rapidly. For example, in the famous northeast blackout of 2003 at around 4:06 pm system operators started noticing abnormalities, by 4:10 pm a power surge severed a large system line that caused a cascade of events leading to a blackout [17]. In another instance, the San Diego blackout of 2011, a system suspended primarily by 500 kV line went down in roughly 10 minutes following a maintenance problem at a remote substation [18].

Unlike the fast most moving makeup of electrical systems, physical reactions in gas systems can be delayed or slow moving. We see this from our results, as even under sustained attack the system takes several hours to reach a drop off in production. Part of this is due to how uniquely our scenario is set up. But these results could also be explained by the fact that it takes a long time (relative to electricity) to physically move gas through pipes. We also believe that this could be why isolated attacks can have less of an effect on the system. In addition, with large gas pipes like the ones used in our system there is a substantial amount of gas storage retained within a pipe. Even when there are jerks in the system, such as rapidly increased load, the storage retained within the pipes can prevent failures as there is not an immediate reaction to the change in system state. For example, in the early scenario where we demonstrated no control response the system did not fail until the end of the 12 hour window. This suggests that there is ample time to react when the system is not operating entirely as it is supposed to. Additionally, placement of gas storage tanks along strategic places within gas lines can provide another source of safety for systems trying to avoid failure during periods of high strain. This does not bode well for attackers trying to damage the system, but it also does not mean the system is safe from strategically devised and coordinated attacks.

We believe the main take away from the earlier experiments then is that hackers would have to sustain an attack for hours or even days to have an impact on the gas system. For example, in our system which is sustained by one primary source of gas - not unlike the 500 kV line in the San Diego blackout example - a measurement attack had to be sustained for nearly 12 hours before there was any effect. However, to advance our hypothesis it would require several

varying gas system scenarios. Our single scenario was intentionally designed to be challenging for a gas system, so that we could demonstrate the value of simulation and measurement attacks. Regardless, the long time for an effect to occur is encouraging in the sense that it would be very difficult to trick system operators for hours on end. Redundant sensors and cross-communication between operators would likely allow them to identify faulty controllers in the system before any failures occur. However, because cyber attacks are fairly unprecedented on gas pipelines it is also unlikely that the average gas operator would make the assumption that their system is under attack, rather than explaining any strange phenomena as faulty sensors or devices. Despite this, if a measurement attack is sustained on a gas system undetected for long enough it could cause a set of rapid cascading failures like we saw in experiment 2.

7 Future Work

7.1 Analyzing and Preventing Measurement Attacks with Machine Learning

In our work we have demonstrated the theory behind measurement attacks, as well as possible dangers they may pose. However, we have not delved into possible methods for detecting or preventing these attacks. As machine learning has become an increasingly popular approach for problems involving large amounts of data, or for deriving insights based on data, it would be interesting to explore methods for detecting faulty or lying sensors in a SCADA system. There are several thoughtful approaches to this problem. One introductory thought may be to use statistical methods for outlier identification. These methods may allow control systems to discover sensor readings as outliers within the time series of data they are contained in. Another approach may be to apply deep learning or classification algorithms towards identifying anomalies within a time series of data. However, even if one could identify outliers using this technique what would be the approach for dealing with these outliers? Would it be wise to block out data from sensors presumed to be lying? Is it still possible to derive useful information from a sensor even when it is lying?

Maybe a more interesting approach to observing the data in these SCADA systems is to identify constraints between and within the data coming from sensors. For example, a constraint within the data may be that a pressure sensor in a normal scenario should never read outside of a certain range. However, a more interesting case may be to identify constraints between sensor readings across a physical system. In the case of one PLC it can be something as simple as temperature and pressure vary together because of underlying properties of gas. In addition, one could explore constraints between PLCs in a model. For instance, several power plants downstream of a large compressor station will likely have pressure readings that vary together in response to changes in the compressors outputs. Being able to identify these constraints could allow an intelligent SCADA system to flag violations in data constraints and report these to system operators.

A primary issue with identifying constraints in these SCADA systems is that as the systems scale, they get increasingly complex and interconnected. The process of manually identifying constraints between different components would require time, clever engineering and would be prone to errors. Researchers at Colorado State University have developed methods for automatically identifying constraints in data [19]. This work has developed a methodology for identifying and verifying constraints in large data sets using a feedback loop with subject matter experts. An expansion of this work could deal with how to apply these automated constraint discovery approaches to time series data. Then applying this approach of constraint discovery to SCADA systems.

7.2 Automatic Discovery of Critical Points for Defense

As uncovered in our experiments, not all points of control are as equally critical to the operation of a SCADA system. Compromise of some controllers can be much more devastating than the compromise of others. In the case of our first experiment the compromise of the Fort Collins power plant controller was able to prevent the Fort Collins compressor from feeding gas to the numerous downstream power plants. This resulted in a partial failure of the system. The research question we can derive from this is: how can we determine which points in the system are critical for operation? Is there an automated way to do this? One approach may be to model real world systems using simulation like we have in the past, and then iteratively apply measurement or command injection compromises to each controller. This would be similar to what we did in 5.5. Then upon inspecting the results you could see which compromises caused failures within the system and continue exploration from that point.

As well as automatically discovering critical points of operation, there is a body of research that could explore ways to incrementally employ defenses to these critical points. If you have limited resources to secure a system the high level insight is that you want to prioritize how you defend points in the system. By using automatic discovery of critical points you can determine which areas have the most potential for harm. Then you can develop a strategy for deploying defenses in these points. In this way you could step toward system security by hardening the most vulnerable areas first and work your way out towards defending less vulnerable spots later.

8 Conclusion

In conclusion, the areas of SCADA simulation, cyber security, and measurement attacks are all interesting paths for research. We believe that improving the realism of SCADA simulations, and investigating possible defense mechanisms are critical research areas for protecting the current infrastructure of power and gas systems. In this paper we had two essential contributions. The first was a question "What are the impacts of measurement attacks in SCADA systems and how can we defend against them?". The second, was how could we build a

model of a gas system to realistically reflect the impact of these attacks. From our work we discovered that an intelligent attacker can cause substantial harm to a system by employing these types of attacks. Additionally, we were able to uncover some interesting observations on the nature of how gas systems respond to cyber attacks. Experiments on the nature of critical points in the system opens a new area of research related to defense of SCADA systems. On top of this, the use of simulated environments is what helped us get to a point where we can start reasoning and asking high level questions about SCADA system security without ever touching a real system. This is an encouraging demonstration then for the use of simulation in SCADA security, as it allows us to experiment with hypotheticals that we would not be able to explore on real systems.

References

1. Ralph Langner. Stuxnet: Dissecting a Cyberwarfare Weapon. *IEEE Security Privacy*, 9(3):49–51, May 2011. Conference Name: IEEE Security Privacy.
2. Kevin E. Hemsley and Ronald E. Fisher. History of industrial control system cyber incidents. Technical report, US Department of Energy: Idaho National Laboratory, December 2018.
3. A. Ashok, Pengyuan Wang, M. Brown, and M. Govindarasu. Experimental evaluation of cyber attacks on Automatic Generation Control using a CPS Security Testbed. In *2015 IEEE Power Energy Society General Meeting*, pages 1–5, July 2015.
4. David P Duggan. Penetration Testing of Industrial Control Systems. *Sandia National Laboratories*, page 7, 2005.
5. Thomas Morris, Anurag Srivastava, Bradley Reaves, Wei Gao, Kalyan Pavurapu, and Ram Reddi. A control system testbed to validate critical infrastructure protection concepts. *International Journal of Critical Infrastructure Protection*, 4(2):88–103, August 2011.
6. Thomas H Morris, Zach Thornton, and Ian Turnipseed. Industrial Control System Simulation and Data Logging for Intrusion Detection System Research. *7th Annual Southeastern Cyber Security Summit*, page 6, 2012.
7. Simon Duque Antón, Michael Gundall, Daniel Fraunholz, and Hans Dieter Schotten. Implementing SCADA Scenarios and Introducing Attacks to Obtain Training Data for Intrusion Detection Methods. *arXiv:1905.12443 [cs]*, May 2019. arXiv: 1905.12443.
8. Bertrand and Olivier Taburiux Masset. Simulating Industrial Control Systems Using Mininet. 2018.
9. Qais Qassim, Mohd. Ezanee Rusli, Salman Yussof, Roslan Ismail, Fairuz Abdullah, Norhamadi Ja'afar, Hafizah Che Hasan, and Maslina Daud. A Survey of SCADA Testbed Implementation Approaches. *Indian Journal of Science and Technology*, 10(26):1–8, June 2017.
10. Carlos M. Correa-Posada, Pedro Sánchez-Martín, and Sara Lumbreras. Security-constrained model for integrated power and natural-gas system. *Journal of Modern Power Systems and Clean Energy*, 5(3):326–336, May 2017. Conference Name: Journal of Modern Power Systems and Clean Energy.
11. Tao Li, Mircea Eremia, and Mohammad Shahidehpour. Interdependency of Natural Gas Network and Power System Security. *IEEE Transactions on Power Sys-*

tems, 23(4):1817–1824, November 2008. Conference Name: IEEE Transactions on Power Systems.
12. Burcin Cakir Erdener, Kwabena Pambour A., Ricardo Bolado Lavin, and Berna Dengiz. An integrated simulation model for analysing electricity and gas systems | Elsevier Enhanced Reader, April 2014. Library Catalog: reader.elsevier.com.
13. Kostas Mathioudakis, Nick Frangiadakis, Andreas Merentitis, and Vangelis Gazis. Towards generic SCADA simulators: A survey of existing multi-purpose co-simulation platforms, best practices and use-cases. page 7, 2013.
14. N. Schubert. KMI System Map, February 2014.
15. Colorado Generating Stations – Public Service Company of Colorado, 2017.
16. Akhmatov Vladislav. Analysis of dynamic behavior of electric power systems with large amount of wind power, 2003.
17. Interim Report on August 14, 2003 Black. Technical report, New York Independent System Operator, January 2004.
18. Jeff McDonald and Morgan Lee. Blackout sparks multiple investigations. *The San Diego-Union Tribune*, September 2011.
19. Hajar Homayouni, Sudipto Ghosh, and Indrakshi Ray. ADQuaTe: An Automated Data Quality Test Approach for Constraint Discovery and Fault Detection. In *2019 IEEE 20th International Conference on Information Reuse and Integration for Data Science (IRI)*, pages 61–68, Los Angeles, CA, USA, July 2019. IEEE.