SANDIA REPORT

SAND2016-7195
Unlimited Release
Printed July 2016

Building the Scientific Basis for Cyber Resilience of Critical Infrastructure

Margot J. Hutchins, Robert Forrest, Derek H. Hart, and Jason E. Stamp

Prepared by
Sandia National Laboratories
Albuquerque, New Mexico 87185 and Livermore, California 94550

Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

Sandia National Laboratories

Issued by Sandia National Laboratories, operated for the United States Department of Energy by Sandia Corporation.

Available to DOE and DOE contractors from:

U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831

Telephone: (865) 576-8401
Facsimile: (865) 576-5728
E-Mail: reports@adonis.osti.gov
Online ordering: http://www.osti.gov/bridge

Available to the public from:

U.S. Department of Commerce
National Technical Information Service
5285 Port Royal Rd
Springfield, VA 22161

Telephone: (800) 553-6847
Facsimile: (703) 605-6900
E-Mail: orders@ntis.fedworld.gov
Online ordering: http://www.ntis.gov/help/ordermethods.asp?loc=7-4-0#online

SAND2016-7195
Unlimited Release
Printed July 2016

Building the Scientific Basis for Cyber Resilience of Critical Infrastructure

Margot J. Hutchins, Robert Forrest, Derek H. Hart, and Jason E. Stamp

Departments 5623, 5624 and 8116
Sandia National Laboratories
P.O. Box 5800
Albuquerque, New Mexico 87185-MS0671

CONTENTS

1. EXECUTIVE SUMMARY

Motivation

Rigorous methods and models are needed to quantify, measure, and increase the cyber resilience of critical infrastructure. An adversary may exploit vulnerabilities in the vital networks such as industrial control systems (ICS) associated with critical infrastructure (e.g., energy, financial, transportation, security), in order to achieve harmful consequences. In cyber systems, the number of vulnerabilities may be large, the attack surface changes over time, and the problem consists of both technical and non-technical factors (e.g., errors in software and human error). Given this complex and dynamic landscape, strategically mitigating risk is important, where "risk" considers both the probability of an event and the consequences if that event occurs. One way to decrease risk is to address consequences by ensuring that critical infrastructure is resilient. In this context, resilience is characterized by the magnitude and duration of a deviation from targeted performance levels, given a disruption.[1] Increasing resilience decreases the consequences of a successful attack.

Scientifically rigorous approaches to address cyber resilience are in the nascent stages; further research is required to develop methods that accurately represent the full complexity of real-world systems and threats. The goal of this project is to further the science for cyber resilience by understanding the relationship between ICS resilience and the resilience of the critical infrastructure (CI) they support. We will identify the operation and design factors that affect cyber resilience of CI, and create systems models that represent the dynamic interplay between these factors and the cyber threats that CI face.

The project will establish methods and models to design and measure the effectiveness of measures aimed at enhancing the cyber resilience of critical infrastructure. Three primary objectives will support this goal: i) creating a framework for resilience for critical infrastructure cyber-physical systems, ii) developing a modeling capability for the dynamic interplay between industrial control systems (ICS) and critical infrastructure (CI), and iii) evaluating the effectiveness of specific countermeasures. This project will leverage existing work in resilience and Emulytics™ at Sandia to create metrics for cyber resilience of critical infrastructure.

Framework

Threats and vulnerabilities to critical infrastructure have long been understood as vital and underappreciated components of national security. The need for increased resilience of critical infrastructure assets in the face of these threats and vulnerabilities is now being more formally acknowledged.

---

[1] Vugrin et al., 2010, "A Framework for Assessing the Resilience of Infrastructure and Economic Systems," in Sustainable & Resilient Critical Infrastructure Systems, K. Gopalakrishnan & S. Peetra (Eds.), Springer.

CI is composed of many subsystems that may be affected differently under various threat scenarios. In power systems, examples of these subsystems include generation, transmission, and distribution. Further, industrial control systems (ICS) support each of those major subsystems, among others. ICS span cyber and physical domains, making them susceptible to the cybersecurity threats and vulnerabilities discussed above.

In order to identify effective risk mitigation strategies for CI systems, we must be able to characterize the overall resilience of a CI system (e.g., electric power), the resilience of its subsystems (e.g., transmission), the resilience of the underlying control systems, and the relationship between resilience metrics within systems and across all levels. For example, resilience of control systems will have an impact on the resilience of the major elements of a power system and vice-versa. With an understanding of resilience at different levels in a system and the relationships among system components, we can identify the elements or systems that have the most influence on overall system resilience. Implementing mitigation strategies that improve the resilience of those influential elements and systems will provide greater improvements to the resilience of the entire system.

Testbed

SCEPTRE, part of the EMULYTICS™ (emulation + analytics) capability at Sandia, is a tool that enables investigation of and experimentation on control systems. It is comprised of two main components. The first allows arbitrarily large control system networks to be modeled in their native protocol via virtual machines. The second simulates physical processes and supports faithful control system behavior by providing relevant process values to the control system network. In the scenario used to exercise the framework outlined above, an industry standard power system software package, PSS/E, is used to simulate the dynamic behavior of power transmission systems. There were challenges integrating data from the physical processes simulation into SCEPTRE because the latter was designed to run in real time and the former requires relatively more time to run. We established a playback method to integrate control system modeling and dynamic physical system simulation that may be used in the future for other experiments. This is a new capability added to SCEPTRE by this project.

An experiment was conducted to test the functionality of the simulated relays and characterize message communication timing – a measure of control system resilience that relates to a measure of power system resilience, sag. The experiment was designed as follows:

- simulate a fault on a power transmission line
- distance relays measure the change of impedance, detect the short in the line, and trip the transmission line for safety
- two relays are monitoring the transmission line from either end of the line.

The mechanics of the simulation and net configuration are described in Appendix A.

Results and Future Work

The testbed picked up the fault and tripped with the same timing characteristics as PSS/E, validating the test setup. Furthermore, the trip message was received at the appropriate relay in the timeframe it would be expected in a real substation. This experiment validates the testbed as an appropriate tool to quantitatively investigate cyber resilience of critical infrastructure. Further, the scenario illustrates an application of the system-of-systems framework for cyber resilience of critical infrastructure, namely, the relationship between resilience at the control systems level and resilience at the transmission-level of power systems. Given that the testbed capability has been validated against real world performance metrics, an appropriate next step would be: i) expanding into larger control system simulations, ii) designing and testing new substation network layouts and iii) investigating impacts of typical network attacks. Additionally, the impacts of typical switch security features can be assessed, new intrusion detection system (IDS) technologies can be tested to understand impact to operations, and network defense in-depth strategies can be investigated as well as their effect on resilience.

2. APPENDIX A: ANNOTATED PRESENTATION

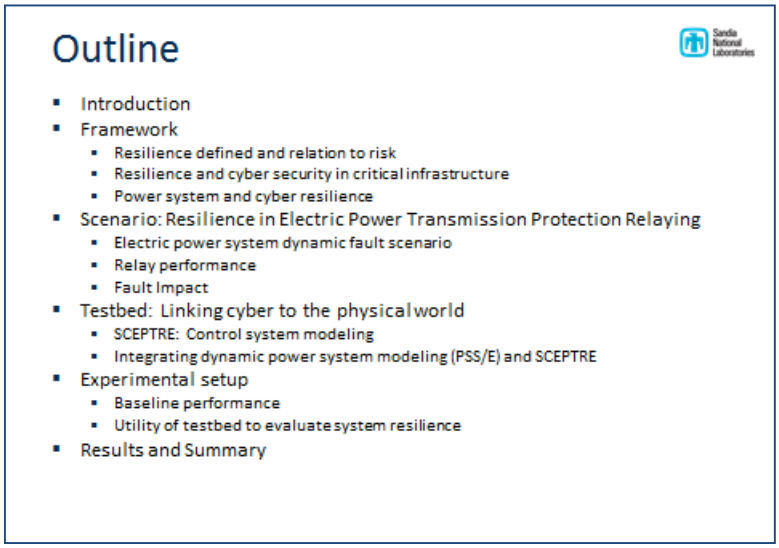

We will begin with an introduction and a framework for the project. We explore the space by defining resilience in the context of critical infrastructure and risk. Then we will discuss the methodology we used to investigate including our specific test bed. We describe SCEPTRE and the new work done to integrate power system simulation with SCEPTERE in a dynamic way. The results and analysis are then presented with baseline performance and a discussion of the utility of using the testbed to evaluate resilience. We will then conclude with a summary and ideas this presents for future work.

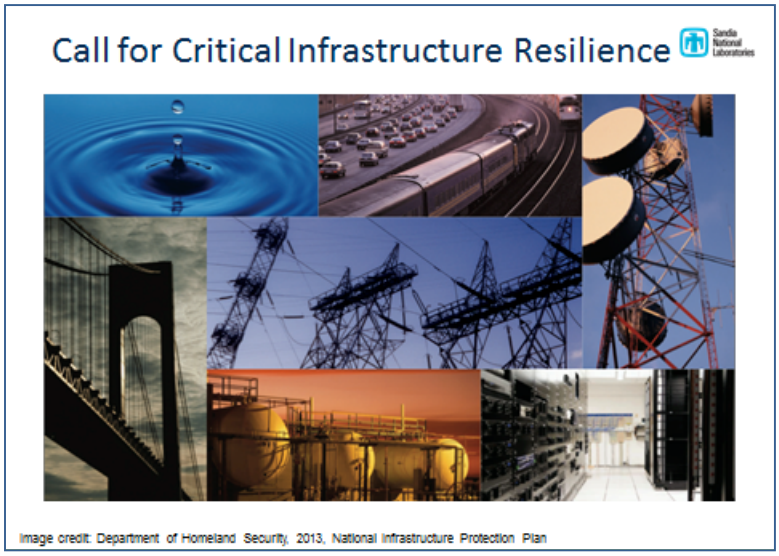

Threats and vulnerabilities to critical infrastructure have long been understood as a vital and underappreciated components of national security.[1] The need for increased resilience of critical infrastructure assets in the face of these threats and vulnerabilities is now being more formally acknowledged. Most notably, the presidential directive (Presidential Policy Directive 21 (PPD-21)) called for the updating of the National Infrastructure Protection Plan (NIPP) to improve approaches to addressing infrastructure resilience. A partial summary is included below.

From the NIPP 2013 Exec Summary:
In February 2013, the President issued Presidential Policy Directive 21 (PPD-21), *Critical Infrastructure Security and Resilience,* which explicitly calls for an update to the National Infrastructure Protection Plan (NIPP). This update is informed by significant evolution in the critical infrastructure risk, policy, and operating environments, as well as experience gained and lessons learned since the NIPP was last issued in 2009. The *National Plan* builds upon previous NIPPs by emphasizing the complementary goals of security and resilience for critical infrastructure. To achieve these goals, cyber and physical security and the resilience of critical infrastructure assets, systems, and networks are integrated into an enterprise approach to risk management.

The integration of physical and cyber security planning is consistent with Executive Order 13636, *Improving Critical Infrastructure Cybersecurity,* which directs the Federal Government to coordinate with critical infrastructure owners and operators to improve information sharing and collaboratively develop and implement risk-based approaches to cybersecurity. In describing activities to manage risks across the five national preparedness mission areas of prevention, protection, mitigation, response, and recovery, the *National Plan* also aligns with the National Preparedness System called for in Presidential Policy Directive 8 (PPD-8), *National Preparedness.*
---

1. R. J. Robles, M. Choi, E. Cho, S. Kim, G. Park, and J. Lee, “Common threats and vulnerabilities of critical infrastructures,” International Journal of Control and Automation, vol. 1, no. 1, pp. 17–22, 2008.

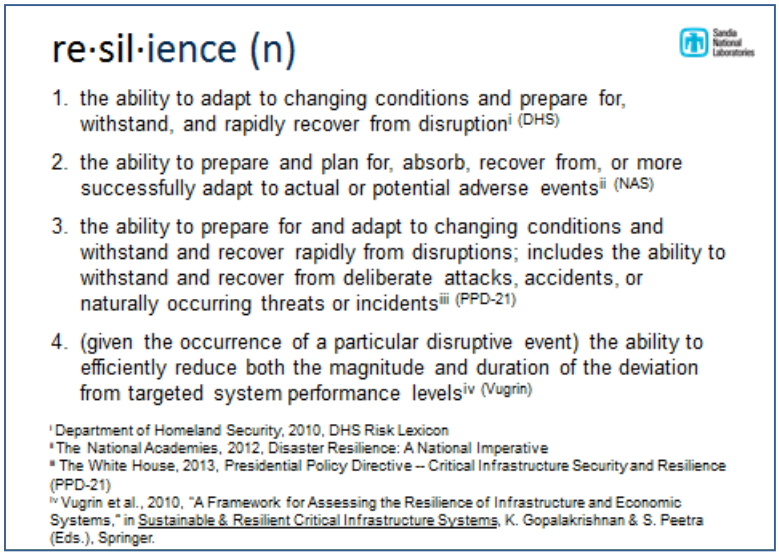

Resilience is a fairly nebulous term that does not yet have a commonly understood formal definition. Above, we see various notable entities defining resilience quite differently. It is instructive to examine a few specific definitions to note differences. The DHS (1) and NAS (2) studies have fairly straightforward definitions but they may not be specific enough to translate into quantifiable, actionable metrics. The PPD-21 definition (3) evolved to some extent, but contains a list calling out specific disruptions in the definition, which may indicate a lack of sufficient clarity. The definition that Vugrin et al, included in their 2010 work lends itself especially well to quantification. They indicated that given the occurrence of a particular disruptive event, resilience is the ability to efficiently reduce both the magnitude and duration of the deviation from targeted system performance levels. This is the definition that most closely tracks what we will use here.

As we will see, one of the challenges with nebulous evolving definitions of resilience is choosing accurate metrics to quantify it.

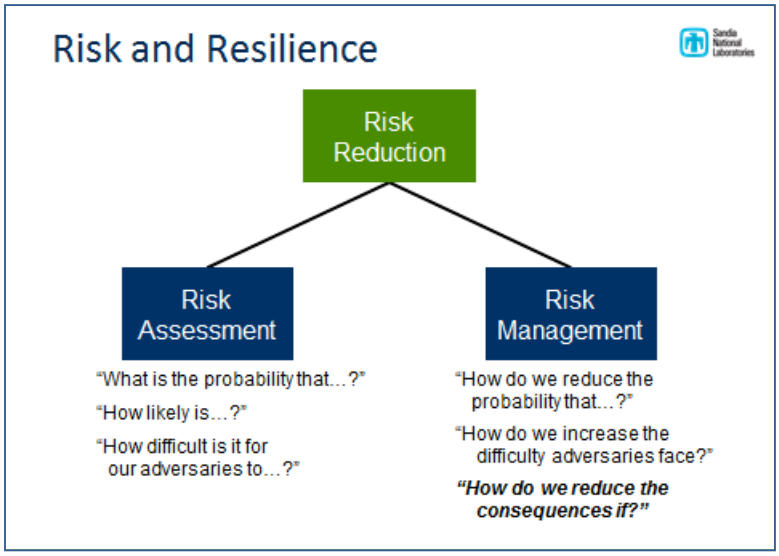

It is important to understand the role of resilience in the greater context of risk reduction. Risk reduction can be broken in to two components: Risk Assessment and Risk Management. Risk assessment is more focused on how large the risk is, i.e., what is the probability of a bad outcome? How likely is that bad outcome? How difficult is it for our adversaries to achieve their goal?

Risk Management is focused on actionable steps to reduce assessed risks. For example, how do we reduce the probability of an undesired consequence? How do we make things more difficult for our adversaries? And, most relevant to resilience, how do we reduce the consequences if a disruptive event occurs?

We, therefore, look at resilience as a very specific way to address a specific component of risk reduction more broadly. If we think of risk as F (threat, vulnerability, consequences), resilience addresses consequence by attempting to reduce the effect after a disruptive event occurs. Simply stated, how do we reduce the fallout if something goes wrong?

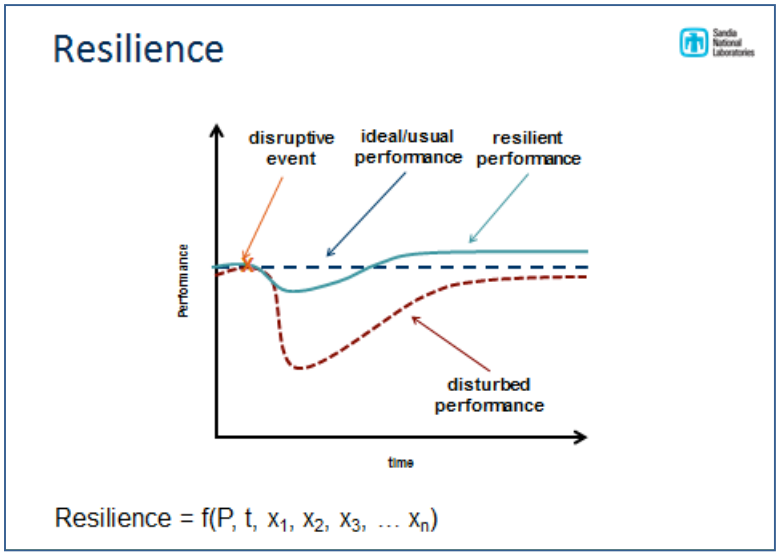

To attempt to get a more intuitive understanding of resilience, we now examine graphically what we previously defined.

The plot shown here shows performance, defined by some metrics, as a function of time. Generally, compared to an unperturbed baseline system (dark blue), a disruptive event occurs at a point in time and alters the system performance (shown in red). Performance degrades for a certain amount of time before recovery starts. Nominal operations are restored and performance returns to baseline. Generically, the loss as a result of the event may be thought of as the difference between baseline performance and the disturbed performance curve. A more resilient system can be thought of as an improvement over the disturbed performance curve, given a disruptive event.

The light blue line indicates how a more resilient system would behave following a disruptive event – the magnitude and duration of the disruption are decreased relative to the original disturbed performance curve. An ideally robust system could come back stronger following an incident, perhaps because new infrastructure has been added or because the system learns something that allows it to operate at a more efficient state.

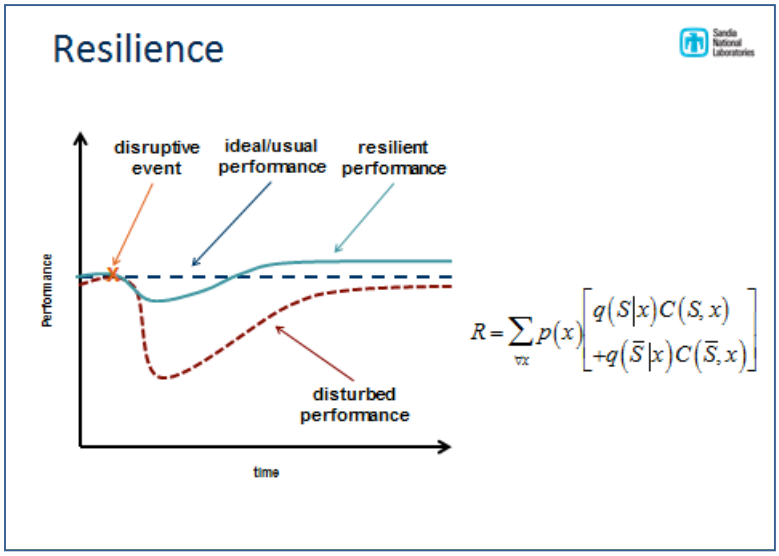

Returning to the context of resilience in terms of a subset of risk, specifically as reducing the consequences of an event, we can now understand resilience in terms of an expression for risk.

The above equation indicates that risk in a system is equal to the sum of the probabilistic risk associated with all scenarios in the system. P(x) is the probability of scenario x occurring. q(S-bar-x) is the probability that scenario x is successful. C represents consequences, assuming the success of a specific scenario. Given success of a certain scenario, performance degrades as a function of time as we have seen previously.

In a system that is relatively more resilient, a particular scenario will have less performance degradation, relative to the baseline.

It is interesting to note that in a real system, according to the formulation above it may not always be beneficial to completely eliminate disruptive events. Truly resilient systems recover gracefully from disruptive events, learn from them and then operate at or above previous performance levels. Totally eliminating these disruptive events would eliminate performance gains realized by addressing disruptive events. In other words, we learn from failures. In an ideal system, we want to decrease real losses from such events but still realize the efficiency gains from learning and engineering more efficient systems. This cycle of learning is common to many systems and industries, most notably the auto and nuclear industry.

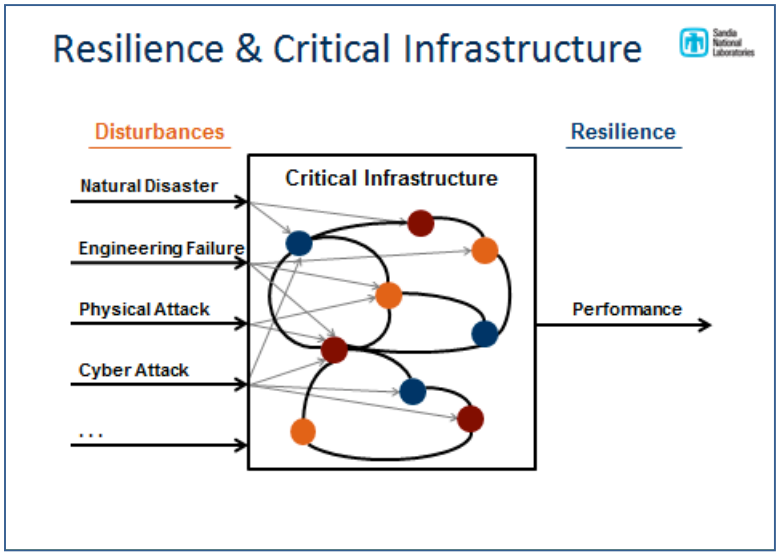

Critical infrastructure is composed of many subsystems that may be affected by various threat scenarios differently. If we understand and can model how a scenario will impact the performance of subsystems and how that affects performance of the entire system, we can identify those areas that have the greatest effect on resilience. At Sandia, NISAC is an example of our capability to model risks, specifically natural disasters.

Cyber systems are increasingly important subsystems to understand in this context. Because of their relative novelty, as well as their complexity, they may well be the most challenging subsystem to understand and quantify in this context.

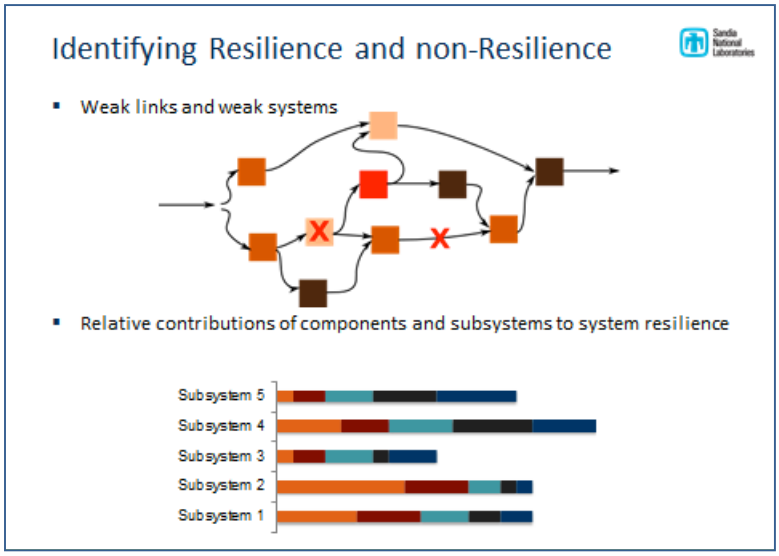

Understanding the components or subsystems that combine to determine resilience is an important first step. Weak links often lead to weak systems, especially if they are single points of failure – that is, they are required for the larger system to perform. As an example, we can map out dependencies of subsystems to understand the relative importance of components. This can allow us to concentrate on vital links, understand relatively important links between systems and find single points of failure.

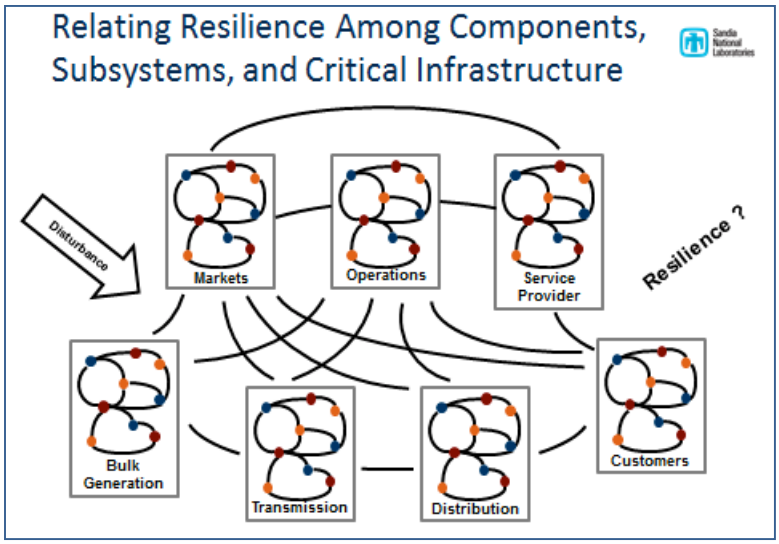

As we mentioned previously, infrastructure is made up of many subsystems, such as generation, transmission, and distribution, as illustrated above. These interconnected systems of systems make up our critical infrastructure. It is vital to identify which components and subsystems can be altered to have the greatest impact on overall system resilience given various disturbances.

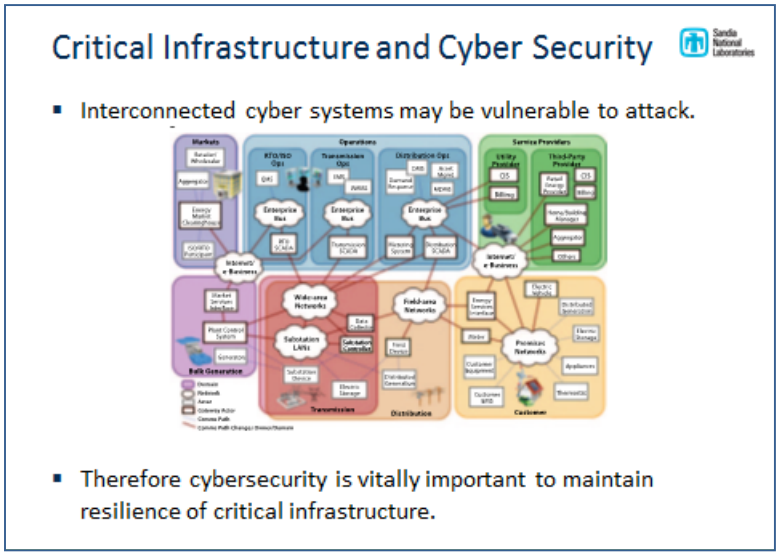

Here is a cyber version of the diagram in the previous slide.

Cyber systems are vital subsystems to understand, and some of the most difficult to quantify in terms of resiliency. All systems have some cyber aspect to them, and, therefore, are susceptible to cyber vulnerabilities.
No matter what system one looks at, control systems in particular play a vital role in operations. Control systems comprise the heart of infrastructure operations. For example a power plant may operate with a large network of SCADA components (RTU, HMI, PLC components) interacting with human operators. Similarly, large scale distribution networks contain control system components that operate on a larger geographic scale.

The diagram above illustrates on a very high level the vital role of control systems in a grid and their interconnectedness.

Source: National Institute for Standards and Technology, *NIST Framework and Roadmap for Smart Grid Interoperability Standards, Release 1.0*, special publication 1108 (Washington, DC: U.S. Department of Commerce, 2010), 35, http://www.nist.gov/public_affairs/releases/upload/smartgrid_interoperability_final.pdf.

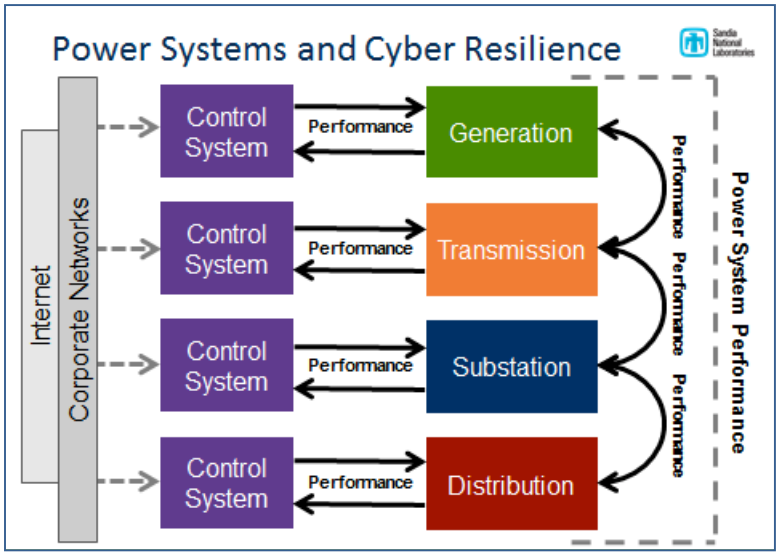

Consider several inter-related power system elements and sub-systems. Control systems support each of the major power system elements shown here (i.e., generation, transmission, substation, distribution). Therefore, the performance of the control systems will have an impact on the performance of the major elements. Further, the performance of a major element will impact the performance of other major elements (not all relationships are shown here for simplicity) and the performance of major elements will impact the performance of the control systems. Vulnerabilities in any of these systems or elements may be exploited to impact performance directly or indirectly. As we know, vulnerabilities in corporate networks may also be exploited to directly impact one device or system and cause cascading degradation in performance that ultimately diminishes overall system performance (or key performance metrics of interest).

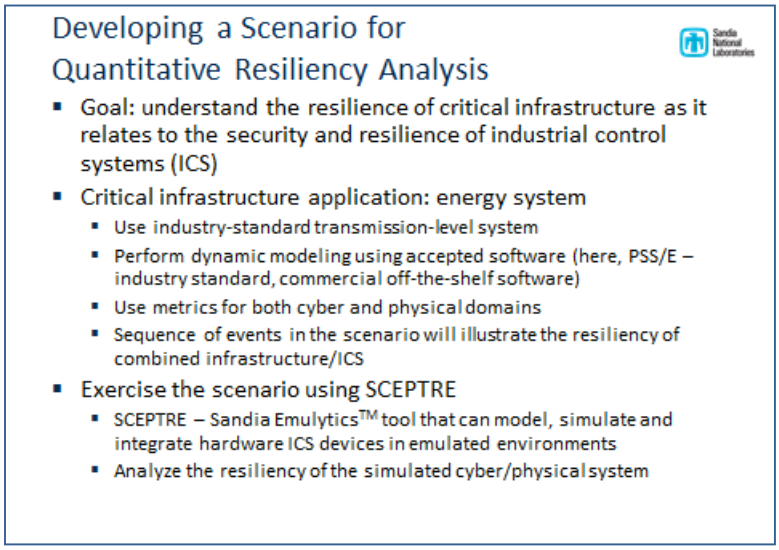

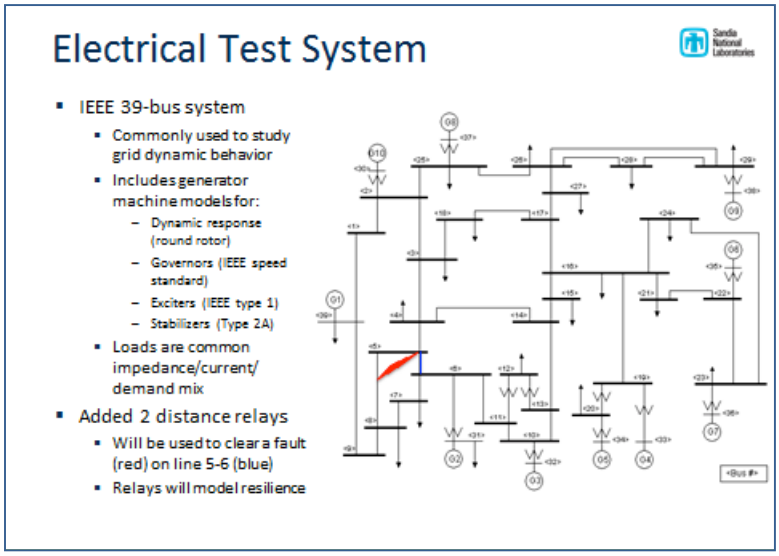

Source: **http://psdyn.ece.wisc.edu/IEEE_benchmarks/.**

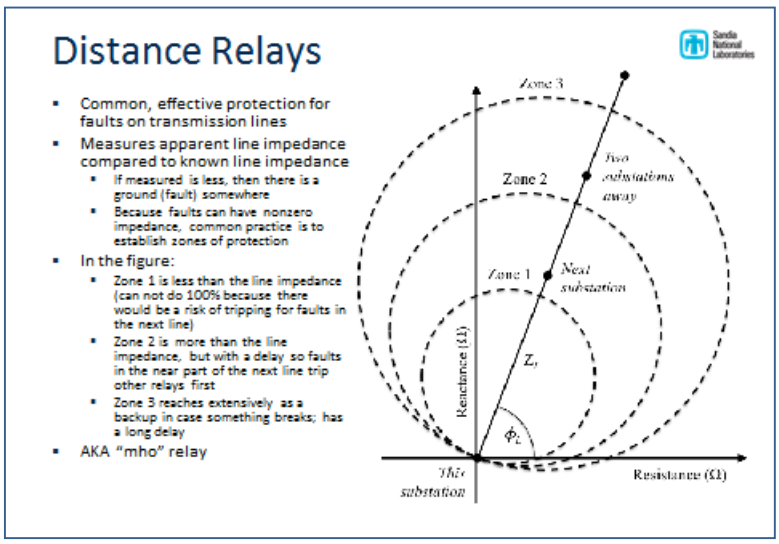

Pilot Relaying for Faster Trip Speed

- Lines trip at both ends to isolate a fault
- For faults in close proximity to one end of a line, one relay will see zone 1 and the other zone 2
- The zone 2 relay will delay tripping
- Solution:
  - If one relay sees a fault in zone 1, then it must be on the line, so it sends a signal to the other end for a no-delay trip
  - Requires communications
  - Certain implementations are called "permissive overreach transfer trip" or POTT
- Modern GOOSE protocol could carry permissive message
  - Generic Object-oriented Substation Event (GOOSE)
  - Layer 2 (likely will be routable in the future)
  - Redundant fast ethernet is common
  - Expecting VERY short communications delay

Impacts of Faults

- Energy imbalance causing machine rotor acceleration
- Causes local frequency and phase oscillations
- May lead to groups of machines "fighting" against each other to maintain frequency, which leads to loss of stability and uncontrolled grid collapse
- Voltage will be locally depressed, called a "sag"
- Generator exciters will change state to maintain voltage and stability
- Relating to resiliency:
  - The length of the sag, and the "nearness" to instability, depend on the clearing time for a fault
  - For the pilot scheme discussed previously, this will change based on the communications delay and the system's cyber security

Initial Data Collection

- Using PSS/E, capture data files for SCEPTRE integration later
  - Voltage magnitude/angle at buses 5 and 6
  - Current magnitude/angle at buses 5 and 6
  - To check model operation, also store the impedances seen by the relays at buses 5 and 6
- Fault duration:
  - 39-bus system is dynamically unstable for relatively short faults
  - For line 5-6, fault must be cleared in 4-5 cycles (a cycle is 16.67ms for a 60Hz system)
- Relay setup (both ends):
  - Zone 1 reach is 80% of the line impedance, no delay
  - Zone 1 reach is 80% of the line impedance, with no additional delay
  - Zone 2 reach is 150% of the line impedance, 18 cycle delay
  - Additional time for the relay to output a trip and the breaker to open: 2 cycles (this is fast but still reasonable)
- Analysis:
  - The sag duration increases linearly with GOOSE delay
  - A delay of 2.6 cycles leads to instability

Scenario Setup (2-cycle GOOSE Delay)

- At time t = 0:
  - Fault on line 5-6
  - At bus 5 end
  - Distance relay at bus 5 enters zone 1 (no trip delay)
  - Starts the breaker trip, total time is 2 cycles
  - Also sends permissive trip to relay at bus 6
  - Distance relay at bus 6 enters zone 2 (18 cycle delay)
- After 2 cycles:
  - Breaker at 5 trips, voltage recovers somewhat
  - Relay at 6 receives permissive trip signal, starts breaker trip (takes 2 cycles)
- After 4 cycles:
  - Breaker at 6 trips
  - Fault is totally removed from the system
  - Voltage recovers eventually

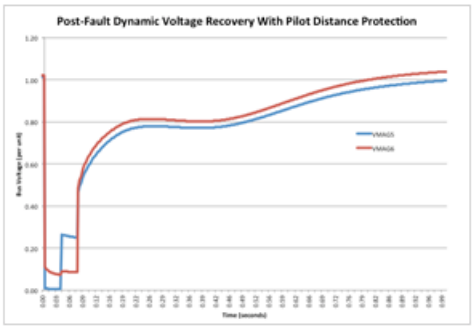

Relating Critical Infrastructure and Control System Resilience

- How do cyber security vulnerabilities, exploitation and defense impact resilience of critical infrastructure?
  - Power systems model provides dynamic characteristics of fault condition (e.g., voltage recovery)
  - A testbed is needed to bridge cyber-physical systems
- SCEPTRE (developed at Sandia) is capable of modeling both the ICS and critical infrastructure processes being controlled
- Control system resilience can be evaluated based upon the ability to perform under duress
- SCEPTRE can be used to assess control system performance and resilience as it relates to critical infrastructure resilience

Why SCEPTRE?

- SCEPTRE models both the control system network and underlying physical process
- integrates standard protocol stacks, devices communicate and interact via actual SCADA protocols
  - Uniquely positions SCEPTRE to answer questions about cyber resiliency
  - Can put hardware in the loop and monitor using standard tools (Splunk, Wireshark, etc)

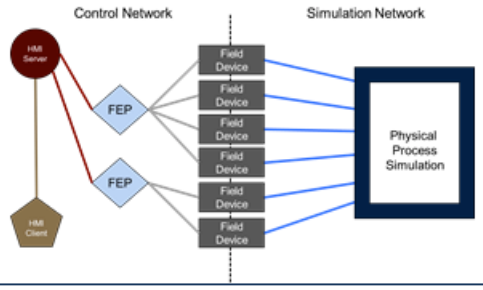

SCEPTRE, part of the EMULYTICS™ (emulation + analytics) capability at Sandia, is a tool that enables investigation of and experimentation on control systems. It is comprised of two main components. The first allows us to model arbitrarily large control system networks in their native protocol via VMs. The second simulates physical processes and translates these processes into values that the control system reads and feeds to the control system network. In our experiment the physical processes will be tied to PSS/E to simulate power transmission systems.

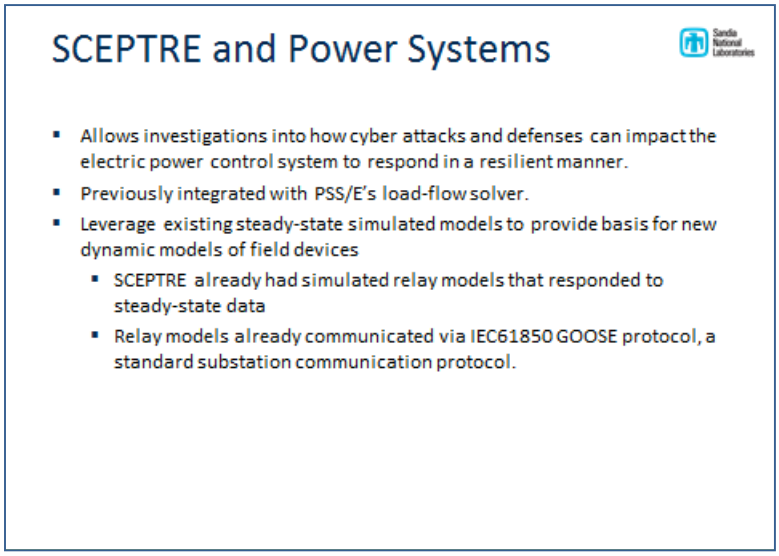

Because SCEPTRE simulates control system devices with VMs, it's possible to instantiate realistic control system networks, communicating via actual SCADA protocols. We can then interact with this virtual network with standard network tools, like vulnerability tools and network scanners.

SCEPTRE electric power control systems were tied to PSS/E as part of the Secure and Sustainable Energy Futures (SSEF) Mission Integration Program Management (MIPM) project in year 1. PSS/E is an industry recognized power transmission system planning software package developed by Siemens. PSS/E has a Python scripting interface that makes integration into the SCEPTRE environment relatively simple. From inside SCEPTRE experiments, control system events can initiate solves in the PSS/E software, which returns a new steady state to the control system devices. Due to the acceptance of PSS/E as a industry standard used around the world, the effects modeled on the power system can be viewed with higher confidence.

Generic Object Oriented Substation Events (GOOSE) is a control mode that is defined as part of the IEC-61850 standard.

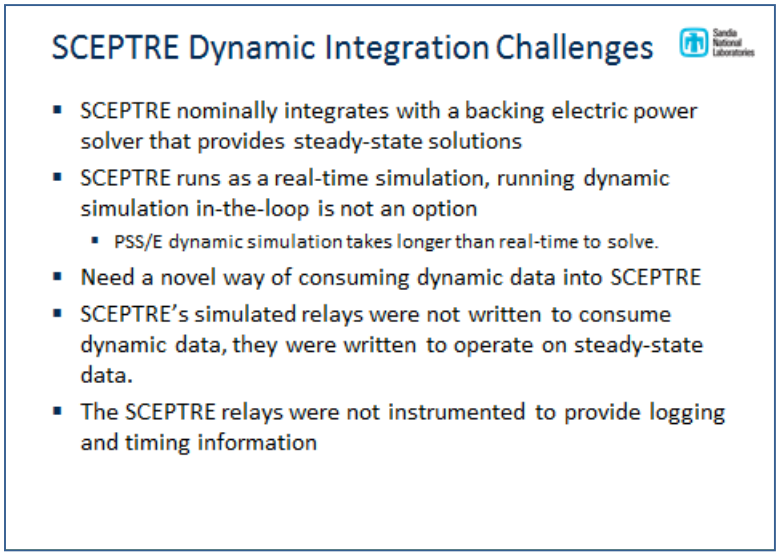

SCEPTRE was designed to run in real time, so there are some challenges we encountered integrating dynamic data (simulating real physical processes). We designed a novel way to do this that may be used in the future for other experiments.

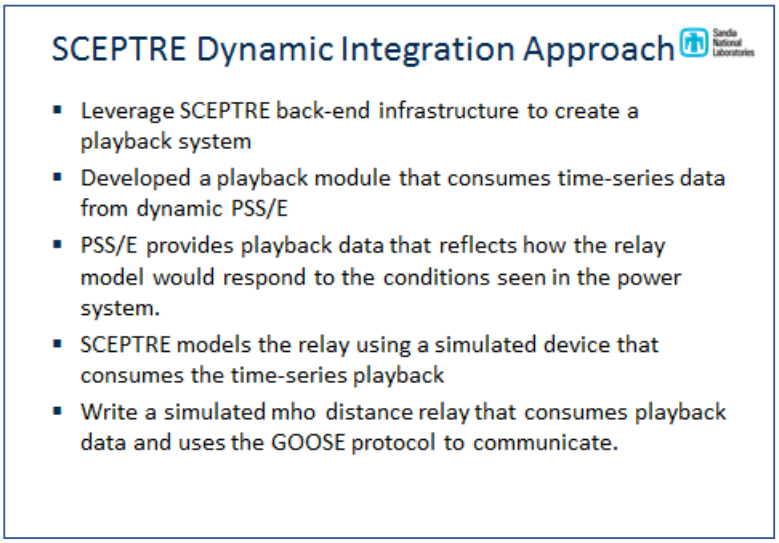

The playback method described here is the approach we used for dynamic integration.

Experimental ICS Testbed Description

- 2 mho distance relays were modeled in a pilot protection scheme
  - Relay 1 (monitoring branch 5 to 6) in the system trips first
  - Relay 2 (monitoring branch 6 to 5) will trip when it receives the GOOSE message from the first relay
- Each simulated relay was deployed as part of an Emulytics™ environment in SCEPTRE.
  - 2 Hosts were used to deploy the relays
  - Connect through the SDN switch on the host and move data across the real infrastructure switch.
- Playback system runs on another virtual machine and pushes data to Relay 1

We describe the experimental testbed here. We are simulating a trip in a power transmission line. Distance relays are supposed to measure the change of impedance, detect the short in the line, and trip the transmission. The two relays are monitoring the transmission line from either end of the line. Relay 1 will detect the trip from one end and send a GOOSE message to relay 2 to trip the line from the other end. This is the process we are simulating. The mechanics of the simulation and net configuration are described here and shown in the next slide.

Experimental Testbed Diagram

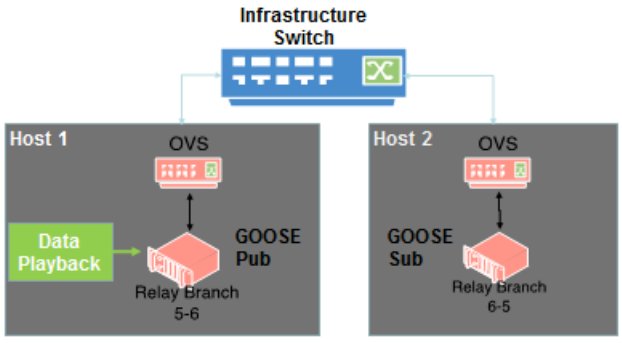

A diagram of the experimental setup.

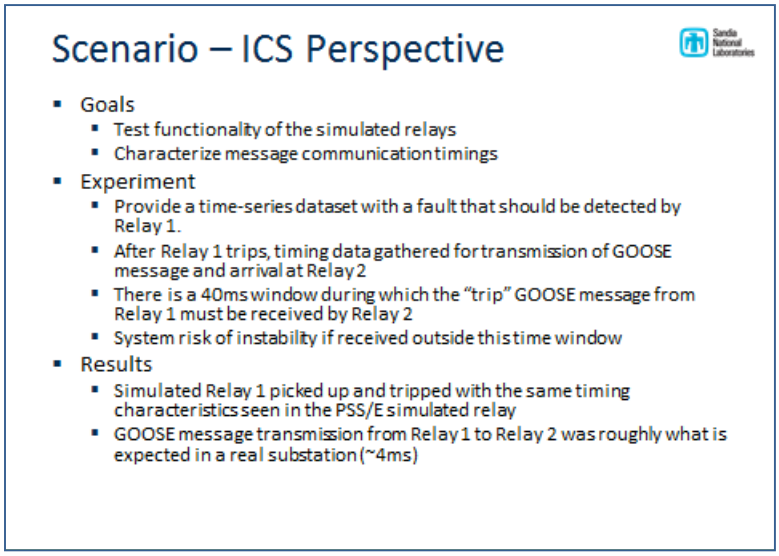

The goals of the experiment are to test the functionality of the simulated relays and characterize message communication timings. The steps of what should unfold after the fault are described above.

The results are promising. We see that our testbed picked up the fault and tripped with the same timing characteristics as PSS/E. Furthermore the GOOSE message was received at Relay 2 in the same timeframe it would be expected in a real substation.

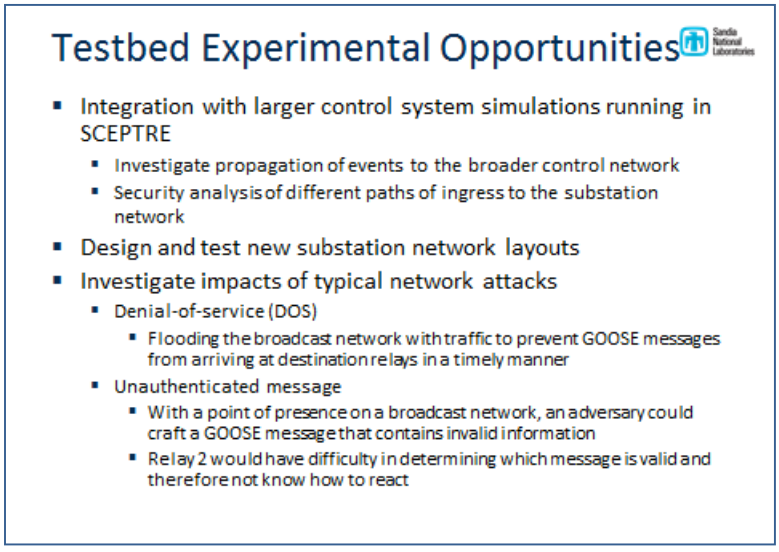

With the results we achieved in the experiment, we can now look to future opportunities to further explore resilience in power systems. Since we validated this capability against real world performance metrics, we have faith that we can begin to expand SCEPTRE into larger control system simulations, design and test new substation network layouts, and investigate impacts of typical network attacks.

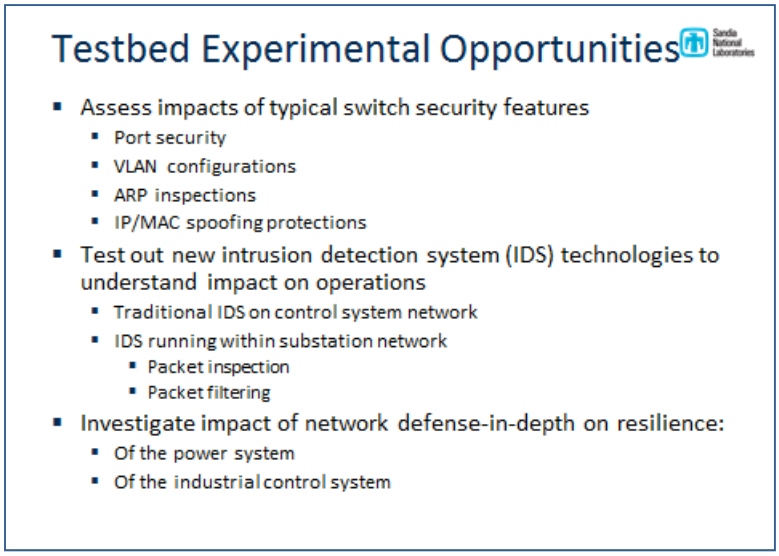

Additionally, we can assess impacts of typical switch security features, test out new intrusion detection system (IDS) technologies to understand impact to operations, and investigate network defense in-depth strategies and their effect on resilience.

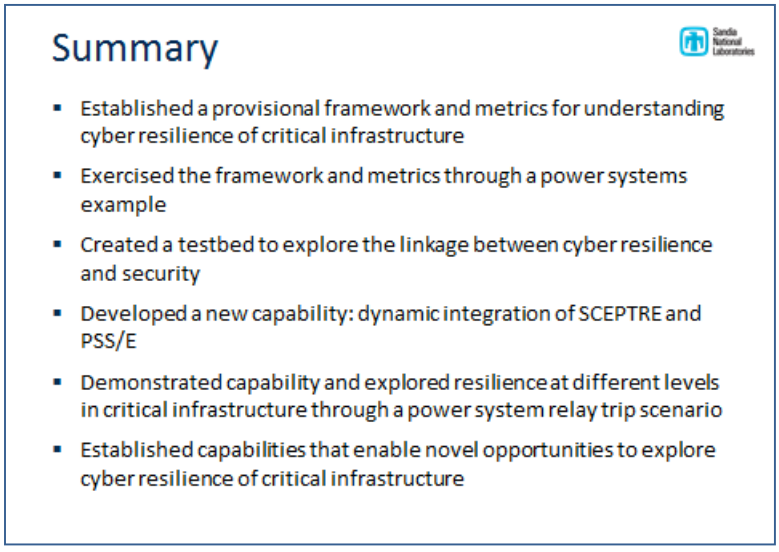

Backup Slides

GOOSE Protocol Basics

- GOOSE is a layer-2 Ethernet protocol that can use either broadcast or multicast addresses (our instances uses broadcast)
- Supports VLAN tagging for logical separation of networks on one physical network
- Open standard that is not brand specific
- No built-in authentication of messages

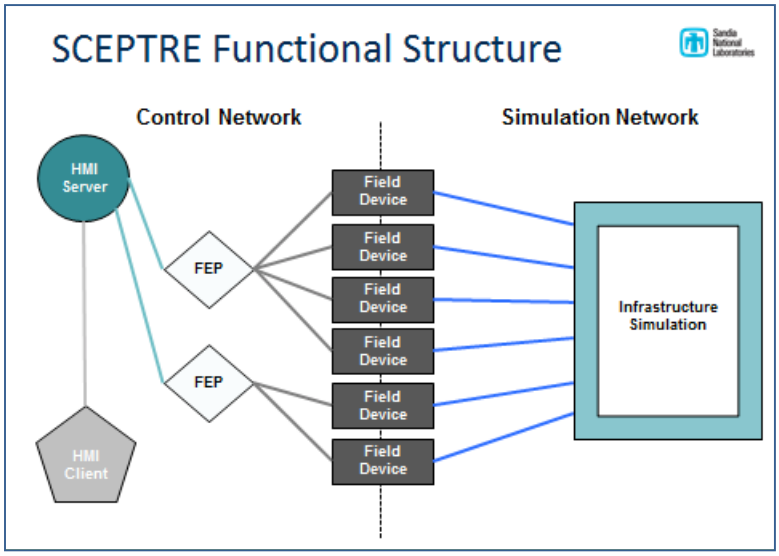
SCEPTRE Functional Structure
Sandia National Laboratories
Control Network
Simulation Network
HMI Server
HMI Client
FEP
FEP
Field Device
Field Device
Field Device
Field Device
Field Device
Field Device
Infrastructure Simulation

3. APPENDIX B

SCEPTRE DEVICE SPECIFICATION: DISTANCE RELAY USING MHO CHARACTERISTIC

(ANSI Type 21)

Jason Stamp
Sandia National Laboratories
Document started 2013 June 24 0809MT
Revised 2013 July 31 1323PT
Spec Version 0.9

1 Introduction

In this doc, we will introduce the concepts necessary to develop a simulated distance relay for power system protection. There are many ways to build a distance relay, but we will be specifically employing the mho characteristic (which will defined later). The initial version will have adequate detail (so as to not require significant reprogramming of the logical steps later) but will make some useful but modest assumptions where ap-propriate. This type of relay is used extensively for protection of high-voltage (115kV and above) transmission lines in power systems. We are modeling an electronic relay that can send network messages.

2 Device Inputs

The actual device has two inputs: one for current, and a second for voltage. The relay would measure and filter each to produce three measurements: the voltage magnitude, the current magnitude, and the phase angle between them. Our simulated relay will have one input, which will either be a sequence with the format:

(*time, current magnitude, voltage magnitude, relative angle*)

or alternatively

(*time, complex current, complex voltage*).

Initially, these will be positive-sequence only (which is the default output of our power analysis software). This places a limitation on the current model: only full balanced faults can be analyzed.

3 Device Settings

The 21 relay will have eight settings emulating configuration selections made during deployment:

- The characteristic impedance for the line (as a two element array of real and reactive components)
- The zone 1 reach setting (in percent)
- The zone 1 time delay setting (in seconds)
- The zone 2 reach setting (in percent)
- The zone 2 time delay setting (in seconds)
- The zone 3 reach setting (in percent)
- The zone 3 time delay setting (in seconds)
- Network notification information for trip alarm messaging

Two additional settings are the bare minimum time for the relay to assert its trip output after a trip is decided, and also the fixed reset time for the device (both are typically in the range of tens of milliseconds). This is typically given in a relay data sheet, but we must select the values here as a part of the modeling process. The concept of trip zones will be covered later.

4 Device Outputs and Status

The device has one analog output, TRIP, that will be associated with a breaker trip input (a 1 will be a trip, and a 0 will be the default non-trip condition). If the relay trips, it must internally store the type of trip (zones 1, 2, or 3). Another output (a networking one) is a network message using IEC61850 that occurs whenever a trip happens and also indicates the trip type (zones 1, 2, or 3). Finally, the model must indicate its status to the SCEPTRE simulation engine (SE) at the conclusion of its simulation interval:

- Status of the output TRIP
- If a trip occurred, the relay also reports the trip time (when the TRIP signal should be associated with breakers, and when the trip alarm network message is sent)
- The current state of the relay (TRIP, RESET, or PENDING RESET; these will be defined later)

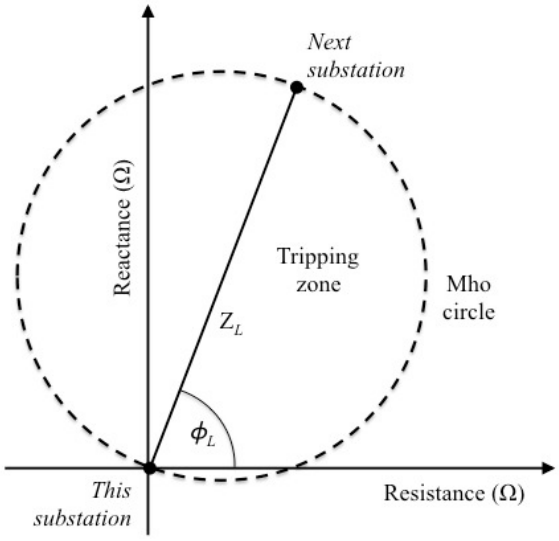

Figure 1: Mho characteristic for a distance relay.

5 Device Operation Logic

A distance relay is used to trip a transmission line if the measured impedance to a fault becomes too low. The theory is that if the measured impedance ever drops below the known impedance of the line, then a fault must exist somewhere along it. This is much more selective than using current settings as fault currents depend on system topology, while impedance to a fault within a single line remains fairly stable.

There are a few caveats to this approach. Impedance is a complex quantity, and the very first distance relays used a complicated arrangement of springs, levers, and coils to cause distance trips based on the ratio of magnitudes of impedance. This led to an undesirable characteristic of tripping for identical impedances measured for both fault currents into the line as well as out of the line (the latter of which clearly indicating that the fault is not within the protected line). Later mechanical improvements resulted in what is called the mho tripping characteristic, which allows for trips in the forward direction only. The result is a circle whose center lies along the plotted complex impedance of the line (as shown in Figure 1). In the figure, $Z_L \angle \Phi_L$ is the impedance of the line (in per unit).

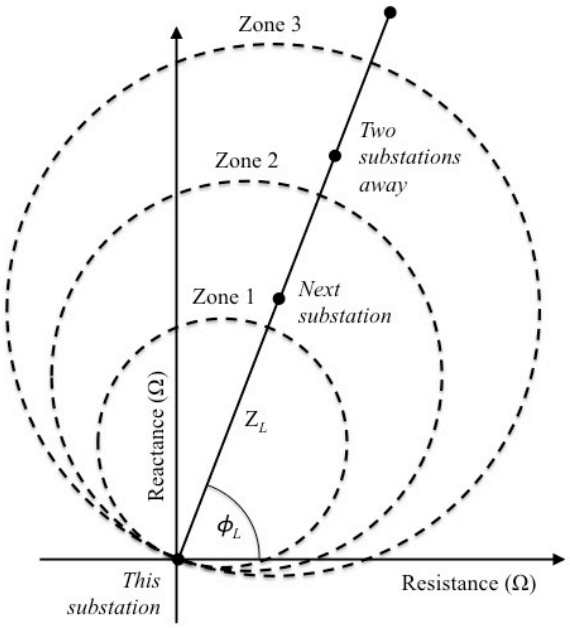

Figure 2: Distance relay zones.

As for settings, we can't set the relay impedance trip threshold to be 100% of the line impedance, since minor fluctuations in the line impedance and unknown fault impedance could make the trip threshold extend beyond the line and into adjacent ones. This would have the undesirable effect of causing a line to possibly trip for faults very near the starting end of adjacent lines. Therefore, a tripping characteristic of 80% of the line impedance with minimal delay is typically used – here, the 80% number is termed the Zone 1 reach setting. A distance relay covers the remaining 20% of the line by setting a second (Zone 2) trip characteristic to extend beyond the end of the line, but with some intentional delay (this way, even though zone 2 reaches into adjacent lines, faults within their areas covered by the first line's zone 2 overreach will trip more quickly on their own zone 1 timing). Finally, each relay also includes a zone 3 setting, which is intended to cover the impedances of both the first line and also the adjacent ones, which acts as a backup in case the primary protect fails on the adjacent lines. Typically, the zone 1, 2, and 3 settings are in the neighborhood of 80&, 150%, and 250% respectively, while the delay for the bigger zone is greater than the smaller one. A diagram of the three zones is shown in Figure 2.

The last essential discussion for the relay model is the tripping logic. Define the following quantities:

Z_L = The line impedance (per-unit)
$L\Phi_L$ = Angle of the line impedance (radians or degrees)
Z_M = The measured impedance at the relay (per-unit)
$L\theta_L$ = Angle of the measured impedance at the relay (radians or degrees)
Z_F = Magnitude of the fault impedance (per-unit)
I_M = Measured current at the relay (per-unit)
V_M = Measured voltage at the relay (per-unit)
d = line impedance fraction (percent)
T_1 = Time delay for Zone 1 trip (seconds)
F_1 = Zone 1 reach (percent)
T_2 = Time delay for Zone 2 trip (seconds)
F_2 = Zone 2 reach (percent)
T_3 = Time delay for Zone 3 trip (seconds)
F_3 = Zone 3 reach (percent)
T_T = Time for the relay to decide to trip (seconds)
T_R = Time for the relay to reset (seconds)
T_D = Delay for the relay sending the trip signal (seconds)

Assume we calculate the measured impedence during a fault as:

$$Z_M = \frac{V_M}{I_M} \tag{1}$$

This is assuming that the voltage and current are given as complex quantities. If magnitudes and the angle difference are given, then (2)

$$Z_M = \frac{|V_M|}{|I_M|} \angle \pm \cos^{-1}\delta$$

The angle is positive when current lags the voltage (i.e. the current phase angle is less) or negative otherwise (a leading current). Assume that for some fault Z_M lies just inside Zone 1, as shown in Figure 3. The displacement of Z_M from Z_L is cause by some nonzero fault impedance (here, it is totally resistive, which is a reasonable assumption).

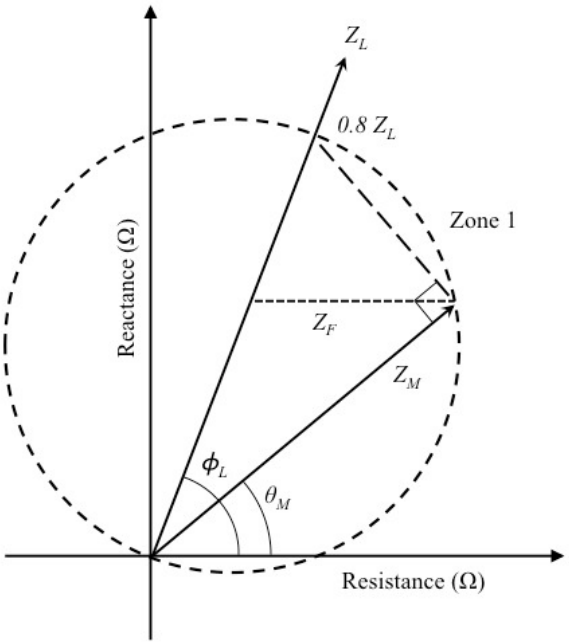

Figure 3: 50/51 relay characteristic curve.

The important issue is that we can algorithmically determine if Z_M is within the zone defined by $0.8Z_L$ using the fact that a right triangle is formed by the diameter and two adjacent chords (one of which is Z_M in the figure), as shown. Using the triangle, if

(3)

$$d = \frac{|Z_M|}{\cos(\phi_L - \theta_M)} < 0.8\,|Z_L|$$

then the impedance lies within zone 1. Similar relationships hold for zones 2 and 3. In Figure 4, we can imagine the right triangles for Z_1 and Z_2, with the former well outside zone 1 (although it might be within zone 3) and the latter easily within zone 1. These sorts of calculations are vastly simplified compared to an actual distance relay, but they will suffice for most introductory research.

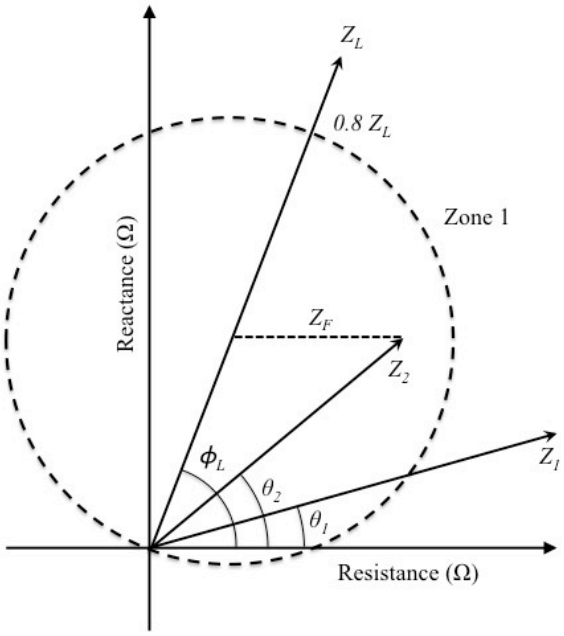

Figure 4: 50/51 relay characteristic curve.

Algorithmically, the trip logic is as follows. The relay will decide to trip at time T_T if any of the the zones measures an impedance within their zone for the requisite time period (the zone delay setting). The relay will actually trip at time $T_T + T_D$.

1. Define D_1, D_2, D_3 as the times in zones 1, 2, and 3 respectively. Set them to zero. Also define a counter $k = 0$.
2. Increment k.
3. Calculate $Z_{M,k}$ and then d_k.
4. For each $d_k < |Z_L| \cdot F_i$ add tl$T = t_{k+1} - t_k$ to the corresponding D_i (where i 2 {1, 2, 3}).
5. Test for trip: if any D_i is greater than the corresponding T_i, then set $T_T = t_{k+1}$ and note which zone tripped. Exit the algorithm.
6. Test for reset: if any D_i has failed to increment at step k, then the zone will reset at time $t_k + T_R$. If this is before t_{k+1} then set D_i to zero; however, if new data points arrive before reset, then they must be evaluated as shown in steps 2 thorough 5. Only if D_i never increments in the interval (t_k, $t_k + T_R$) can the counter for zone i be reset.
7. Loop back to step 2 until the data sequence is exhausted

When the simulation of the relay runs out of samples, we have three possible states:

- TRIP (corresponding to some $D_i = 1$): The device has tripped, in which case the simulation obtains the trip output time and how the device tripped (either on the inverse time element or the instantaneous element).
- RESET (corresponding to all $D_i = 0$): The device is not tripping.
- PENDING_RESET (corresponding to some $0 < D_i < 1$): The device is in between states; this indicates that additional time samples are needed to fully characterize the relay response.

6 Summary

This document should provide adequate information to program a 21 distance relay. The development will aid in the creating of future control devices for SCEPTRE experiments employing dynamic simulation.

7 Model Testing

The sample relay characteristics are show in Table 1, and a sample data sequence is in Table 2. With a given settings, the relay should elect to trip on zone 2 at 0.52 seconds (it actually trips at 0.556 thanks to the additional delay). Zone 2 will start to pick up but reset during the initial few milliseconds after the fault. Leaving the sim to run past the zone 2 trip, we see that zone 1 will pick up instantaneously at 0.56 seconds, and zone 3 would have tripped at 0.7 seconds.

Table 1: Relay settings for the test case.

Setting	Symbol	Value
Line impedance	Z_L	$0.0168 + \boldsymbol{j}0.0899$ per unit
Zone 1 setting	F_1	80%
Zone 1 delay	T_1	0 seconds
Zone 2 setting	F_2	120%
Zone 2 delay	T_2	0.2 seconds
Zone 3 setting	F_3	250%
Zone 3 delay	T_3	0.5 seconds
Intrinsic relay delay	T_D	36 ms
Reset time (all zones)	T_R	15 ms

Table 2: Relay data set: Simple fault occurring at t = 0.20 seconds.

Time (seconds)	Voltage (p.u.)	Current (p.u.)	Angle (degrees)	$\lvert Z_M \rvert$ (per unit)	$d_k / \lvert Z_L \rvert$ (unitless)
0.00	1.021	0.152	-31.300	6.717	127.967
0.05	1.022	0.155	-30.200	6.594	134.117
0.10	1.019	0.161	-28.400	6.329	146.428
0.15	1.018	0.162	-28.200	6.284	147.798
0.20	0.721	3.847	-74.200	0.187	1.385
0.22	0.681	4.116	-79.466	0.165	1.151
0.24	0.702	3.930	-71.844	0.179	1.358
0.26	0.725	3.821	-65.580	0.190	1.563
0.28	0.749	4.064	-65.307	0.184	1.524
0.30	0.758	3.702	-71.199	0.205	1.569
0.32	0.735	3.789	-76.162	0.194	1.400
0.34	0.698	4.134	-80.271	0.169	1.164
0.36	0.650	3.862	-85.176	0.168	1.100
0.38	0.693	3.691	-83.992	0.188	1.243
0.40	0.744	3.778	-82.469	0.197	1.326
0.42	0.683	4.143	-86.938	0.165	1.058
0.44	0.721	4.195	-79.713	0.172	1.193
0.46	0.697	4.408	-78.598	0.158	1.110
0.48	0.634	4.790	-85.075	0.132	0.866
0.50	0.592	4.422	-81.911	0.134	0.907
0.52	0.617	4.295	-82.585	0.144	0.966
0.54	0.602	4.693	-81.077	0.128	0.876
0.56	0.569	4.901	-86.631	0.116	0.748
0.58	0.519	4.649	-80.846	0.112	0.765
0.60	0.468	4.269	-87.455	0.110	0.701
0.62	0.464	4.094	-84.268	0.113	0.748
0.64	0.489	3.747	-87.734	0.131	0.832
0.66	0.492	4.042	-79.799	0.122	0.843
0.68	0.523	3.949	-84.017	0.132	0.877
0.70	0.531	3.736	-76.494	0.142	1.022

In most PDF viewers, a thumbtack icon should appear at left.
Clicking on the thumbtack should extract the original Excel file.

DISTRIBUTION

1 MS 0359 D. Chavez, LDRD Office, 7911

1 MS 0899 Technical Library, 9536 (electronic copy)


Sandia National Laboratories