Sandia National Laboratories

Potential and Limitations for using Data Analysis in the Detection of Cyber Attacks on Cyber-Physical Systems

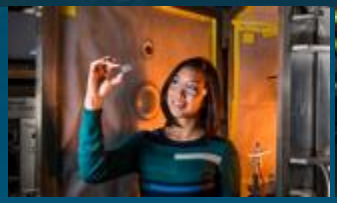

PRESENTED BY

John Mulder

SAND No: SAND2018-13531 C

Cyber-Physical Systems

A cyber-physical system is a system where a physical process is controlled or monitored by computers.

In these systems, the physical and software components interact, sometimes in subtle ways.

Known by many names:

- Operational Technology
- Control Systems
- Critical Infrastructure
- Supervisory Control and Data Acquisition (SCADA)

Cyber-Physical is the New Frontier For Attackers

Night Dragon Operation, 2006
- Attacks hit at least 71 organisations
- Included recon of Oil & Gas
- Continued for several years

Sandworm Team, 2009
- Attacks against:
  - NATO
  - Ukraine
  - Poland Energy
  - European Telecom
  - US Academic Orgs
- Spearfishing with trojanized
- Office documents
- Modified BlackEnergy 2/3

Stuxnet, 2010
- Thumb drive installation to jump air-gaps
- Well written Windows malware
- Payload installs code onto a PLC
- Prevented engineering software from reading PLC logic

Energetic Bear / Dragonfly, 2011
- Attacks ICS asset owners
- SpearPhish Email, Trojanized Software, Watering Hole
- Uses Havex: plug-in framework

Shamoon, 2012
- Attacked Saudi Aramco
- Virus that wipes hard drives
- No evidence that it aimed at SCADA
- very damaging to operations

Ukraine Electric Power Outages, 2015
- Attacks against electric grid
- User lockout followed by on-line attacks (VNC)
- Outages lasted hours

Ukraine Electric Power Outages, 2016
- Attacks against electric grid
- Industroyer / CrashOverride
- Modular framework used to gather intel and script attacks
- Outages lasted hours

TRISIS, 2017
- Malware targeted a safety instrumented system
- Modified code on embedded system

2006 2012 2018

Data Analysis is A Promising Method For Cyber-Physical Security

The underlying systems are based in physics

Polling tends to be regular

Few protocols are necessary to monitor and control any one system

We should be able to detect cyber attacks by watching the control traffic

But there are some problems to be solved…

Problem: Collecting Good Data For Analysis

Ideally, analysis data:

1. Is captured during the whole range of states that can occur in a normally-functioning control system

2. Includes some events or data that are outside of 'normally-functioning' parameters

Problem: How Do We Get Good Data?

Captures of real-world systems often:

1. capture few of the allowable, normal system states
2. capture few anomalous conditions (or none at all)

It could require months to capture the range of states

Few asset owners are willing to let us 'fiddle with' the process to capture system states or introduce anomalies

Problem: Analysis Problems To Be Solved, Given Good Data

What, that occurred in this capture, should be labelled an 'event'?

How do we categorize/label these events?

Which of these events are anomalies?

Which sensors and actuators map onto which physical changes in this event?

Problem: Managing Good Data

Did we make this system less secure by providing false confidence?

Do we have the physical process expertise to understand and label the data from this process?

- Does anyone?

Do our nice-looking results reflect anything real?

- Are they over-fitted?
- Did we detect things that don't matter?
- Did we fail to detect things that matter?

Are we creating false alarms?

- Operators getting false alarms learn to ignore all alarms

What is Needed to Solve The Data Problems?

Time

Access

Bring the right expertise to the problem

Educate users about the limits of data analysis

Rigorously test both inputs and results

Alerts may have to be given with caveats and partial confidence

“The first principle is that you must not fool yourself – and you are the easiest person to fool.”

Richard Feynman