This paper describes objective technical results and analysis. Any subjective views or opinions that might be expressed in the paper do not necessarily represent the views of the U.S. Department of Energy or the United States Government.

SAND2019-4392C

STREAM B

How many printers are on my network?

Dr. David Fritz, Sandia National Laboratories

#CYBERUK19

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.

Emulytics

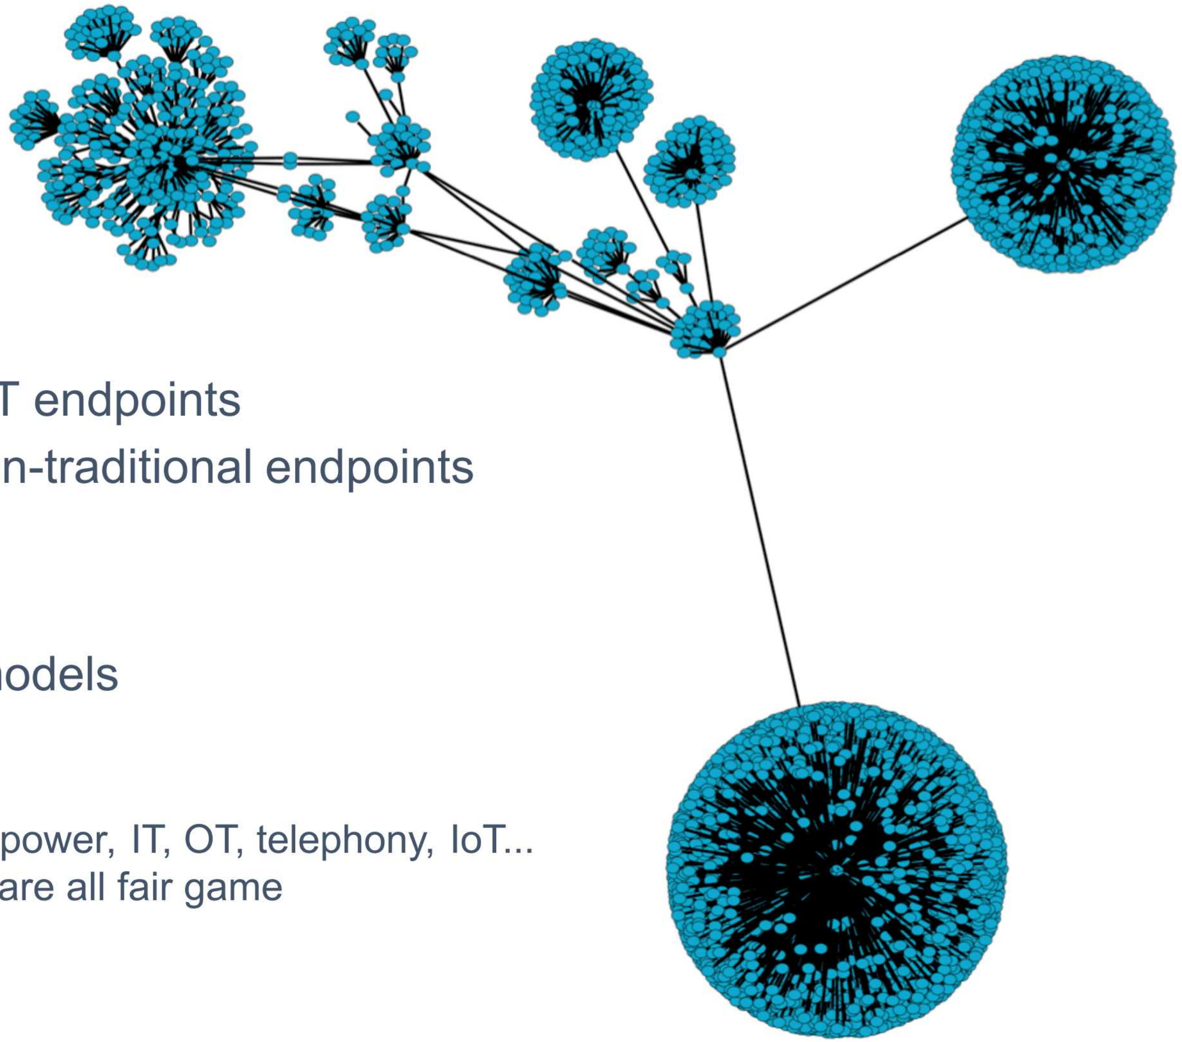

- Emulation + Analytics
- We're working to boot millions of IT endpoints
- And couple that with millions of non-traditional endpoints
  - Mobile devices
  - IoT
  - SCADA/ICS
- And couple **that** with behavioral models
  - Sociology, cognitive science, ...
- Focus on national scale problems
  - Interdependency studies of electric power, IT, OT, telephony, IoT...
  - Cyber and cyber-physical domains are all fair game

Some history as motivation

- Sandia National Laboratories
- Nuclear weapons laboratory, began as the Z division of Los Alamos at the end of the Manhattan Project
- Ordinance design, testing, and assembly
- Gained expertise in red-teaming NW systems
- Carried that expertise into new domains as the lab grew, including cyber
  - ~$250M in cyber-related funding annually
- Lots of cyber focus areas, including modeling and simulation for critical infrastructure, etc.

As a national laboratory…

- *DevOps*: Can we pre-flight new hardware, software, services, to ensure operation in high consequence environments? Can we conduct predictive analysis to detect malfunctions, misconfigurations, malicious consequences?
- *Malware*: Can we gain new understanding of malware through pseudo-in situ execution? How will these 1 million samples impact *my* network specifically?
- *ICS/SCADA*: What are the best countermeasures for my IT-connected ICS systems, despite not having certainty about the threat? Can I detect attacks on ICS systems from the IT-connected perspective? Can I prove resiliency solutions for IT-control over entire power grids?
- ***Nuclear Weapons*: Can we ensure the President will always be able to communicate with a weapon regardless of network state and threat?**

Research and Development in Emulytics

- These are great research questions!
  - And we won't be looking at any of them today!
  - But they do prompt a number of R&D activities in Emulytics itself

So where do we start?

- A few interesting things happened in 2007
  - Cyberattacks in Estonia
  - KVM gets merged into Linux 2.6.20
  - iPhone is released
  - Worst European heat wave in a century (probably unrelated)
- Fast forward to 2008
  - We boot 4 million KVM VMs on Jaguar at Oak Ridge National Laboratory
  - Lessons learned: Switch forwarding tables are still very much vulnerable to MAC flooding

Who cares?

- Can we make use of VM-based sandboxes to model enough of a country to study national-scale attacks?
  - Is my house on fire or is London on fire?
  - How much detail needs to be in the model?
- We certainly have resources
  - Titan - ~300k core / ~18k node AMD cluster
  - Sequoia - ~1.5M core / 98k node PPC cluster
  - And all those little 10k core systems out there

The current situation

- Take a botnet as a “simple” example
  - Reaper
    - Millions of devices
    - *Over a million organizations*
    - Based on MARAI
  - MARAI
    - 600k devices
    - 623Gbps peak!

The current situation

- Amazon has millions of servers
  - Probably in the 3.5 million range, based on the number of availability zones
  - 1.4 million in 2014, the latest datapoint I can find
- What would it take to cycle 1/10th of the VMs hosted on AWS every 90 seconds?
  - With network convergence
  - And C^2
  - And …
- What would it take to measure/instrument all of those VMs in real time?

At scale

- A DHCP file for this is at least 350MBytes
- If all nodes talk to all nodes, kernel tables and lookups dominate runtime!
- Even efforts to implement hierarchy are hard
- … this is a large world

Enter Emulytics

- Suite of tools to support VM-based models
  - IT
  - Electric power
  - IoT
  - Mobile
  - Human behavior
- Much of it is open source
  - minimega.org

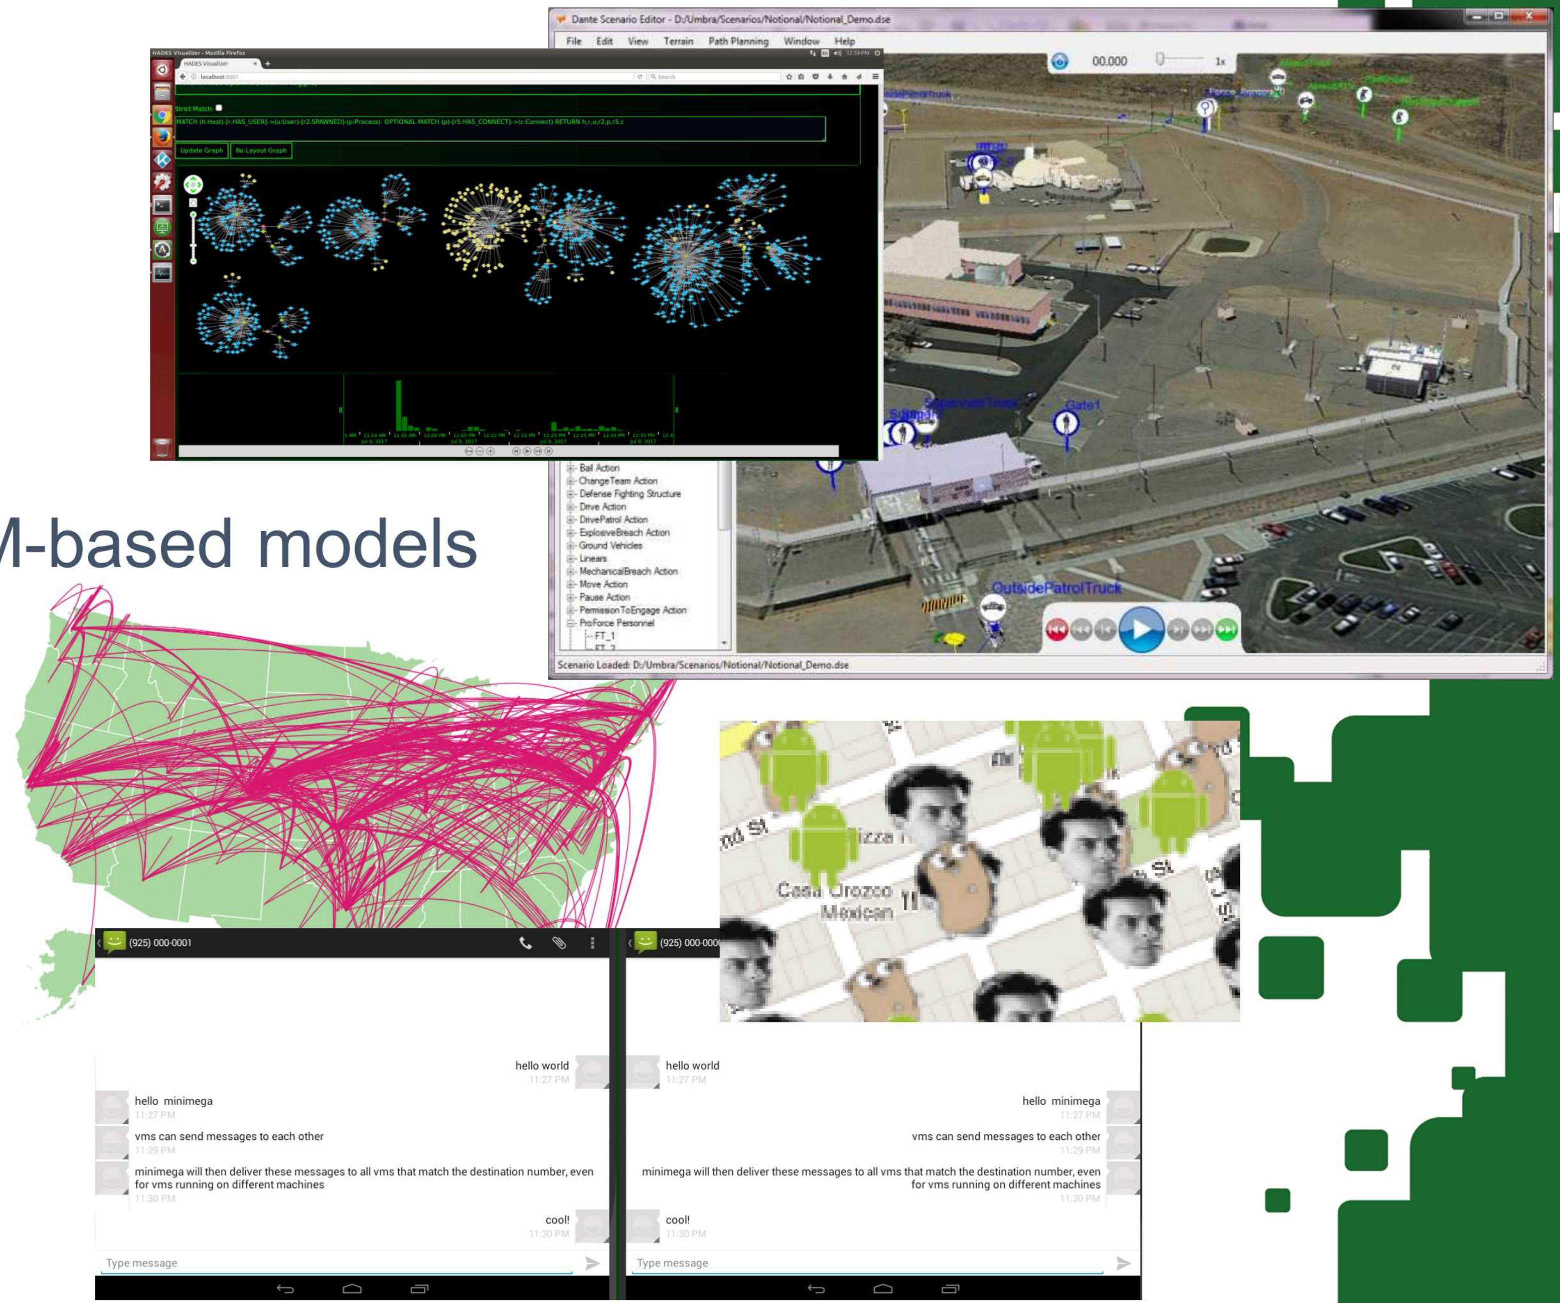

```
MATCH (h:Host)-[r:HAS_USER]->(u:User)-[r2:SPAWNED]-(p:Process)  OPTIONAL MATCH (p)-[r5:HAS_CONNECT]->(c:Connect) RETURN h,r,u,r2,p,r5,c
```

hello world
11:27 PM

hello minimega
11:27 PM

vms can send messages to each other
11:29 PM

minimega will then deliver these messages to all vms that match the destination number, even for vms running on different machines
11:30 PM

cool!
11:30 PM

Modeling *Texas*

- **Texas is a well known global adversary**
- ~28M people
  - Only 17M internet users (wow…)
  - Only 11M facebook users
  - Let’s just take that as the number of devices we’ll model (easily off by an order of magnitude…)
- What does it mean to build a cyber model of Texas?

IT Infrastructure

- Oncor Electric Delivery Company
  - AS-26958
  - 3+ million customers, 105 counties
- Graph shows closest 1000 routers within Texas
  - Only peered with ZAYO
- About 32k nodes

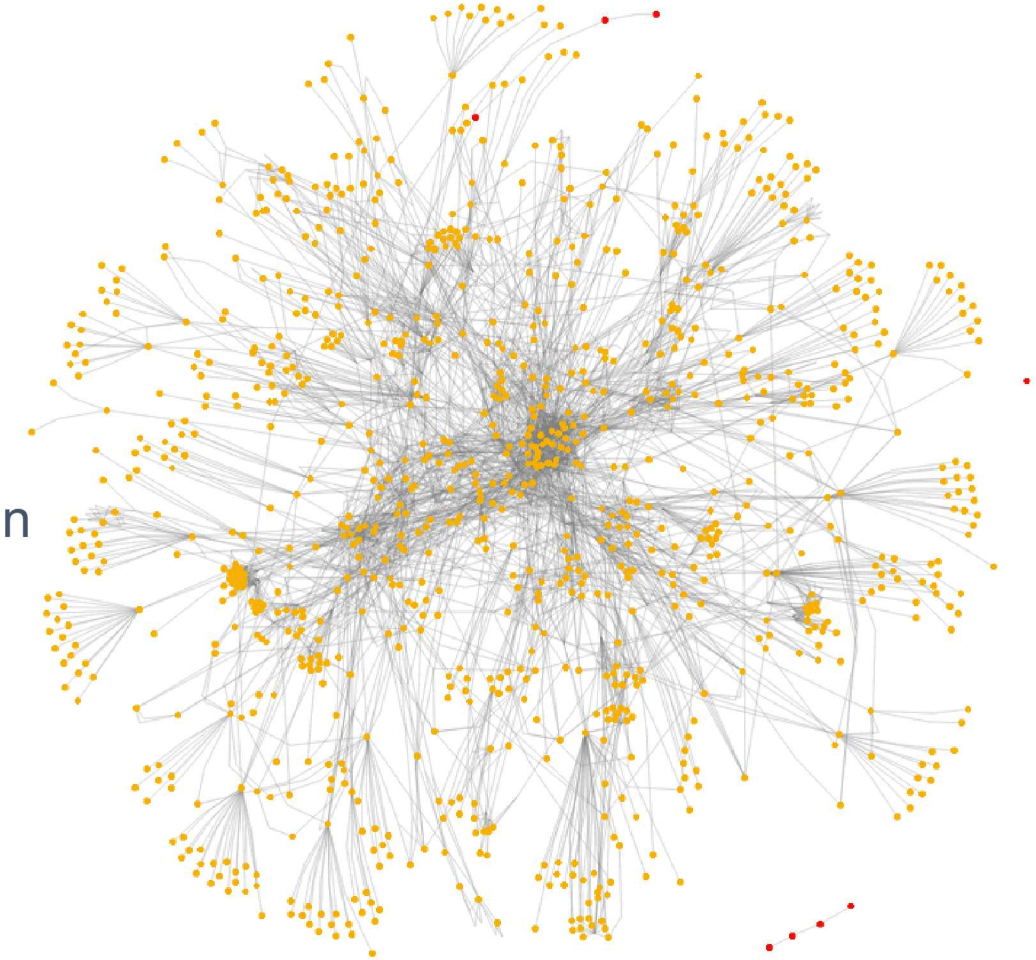

Electric power

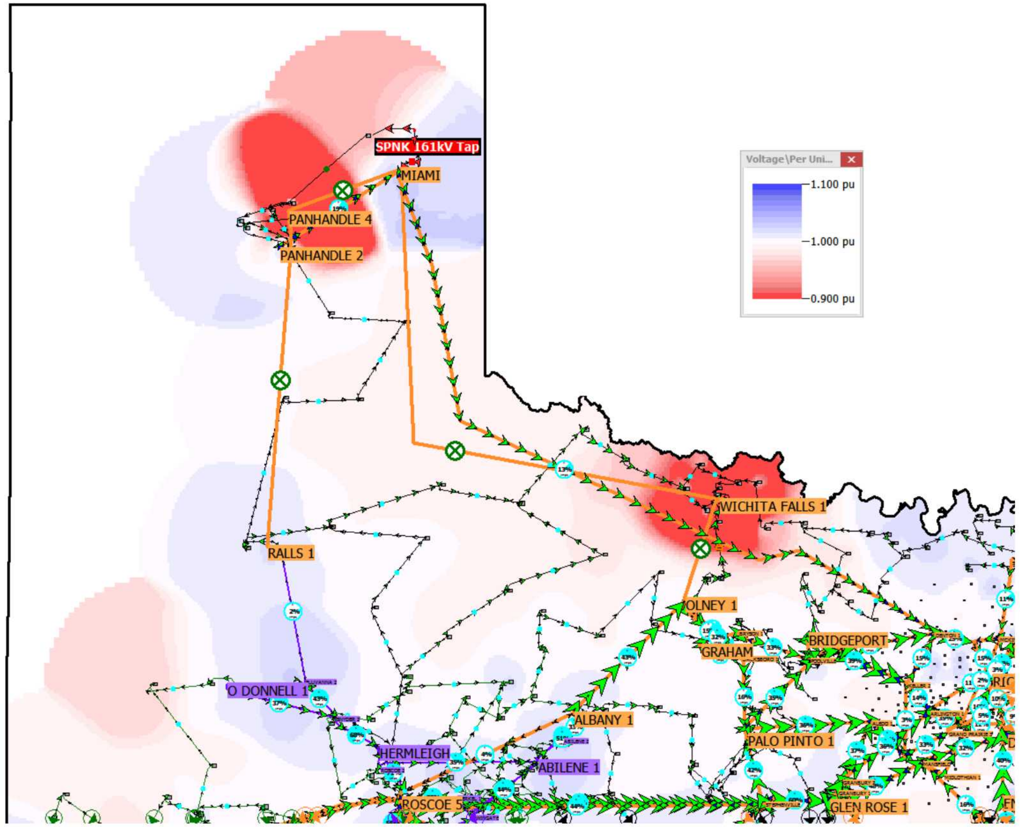

Electric power

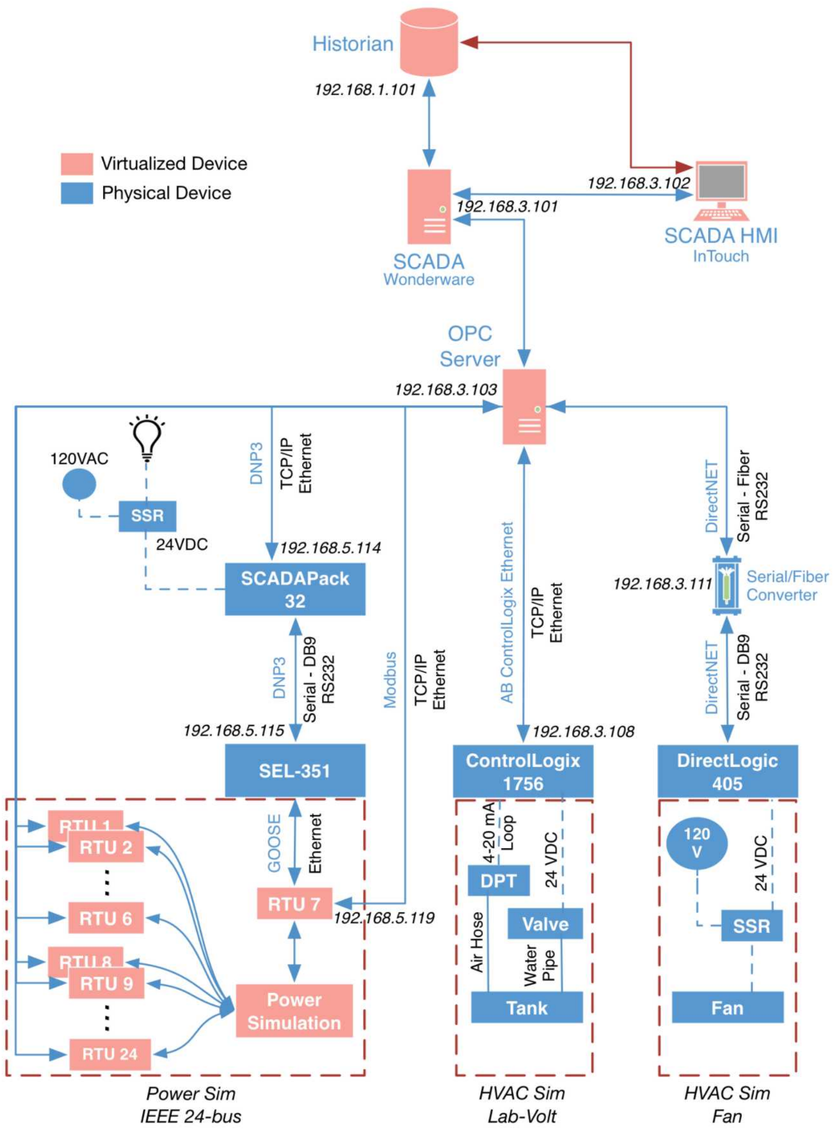

National Cyber Security Centre
a part of GCHQ

CYBERUK 2019

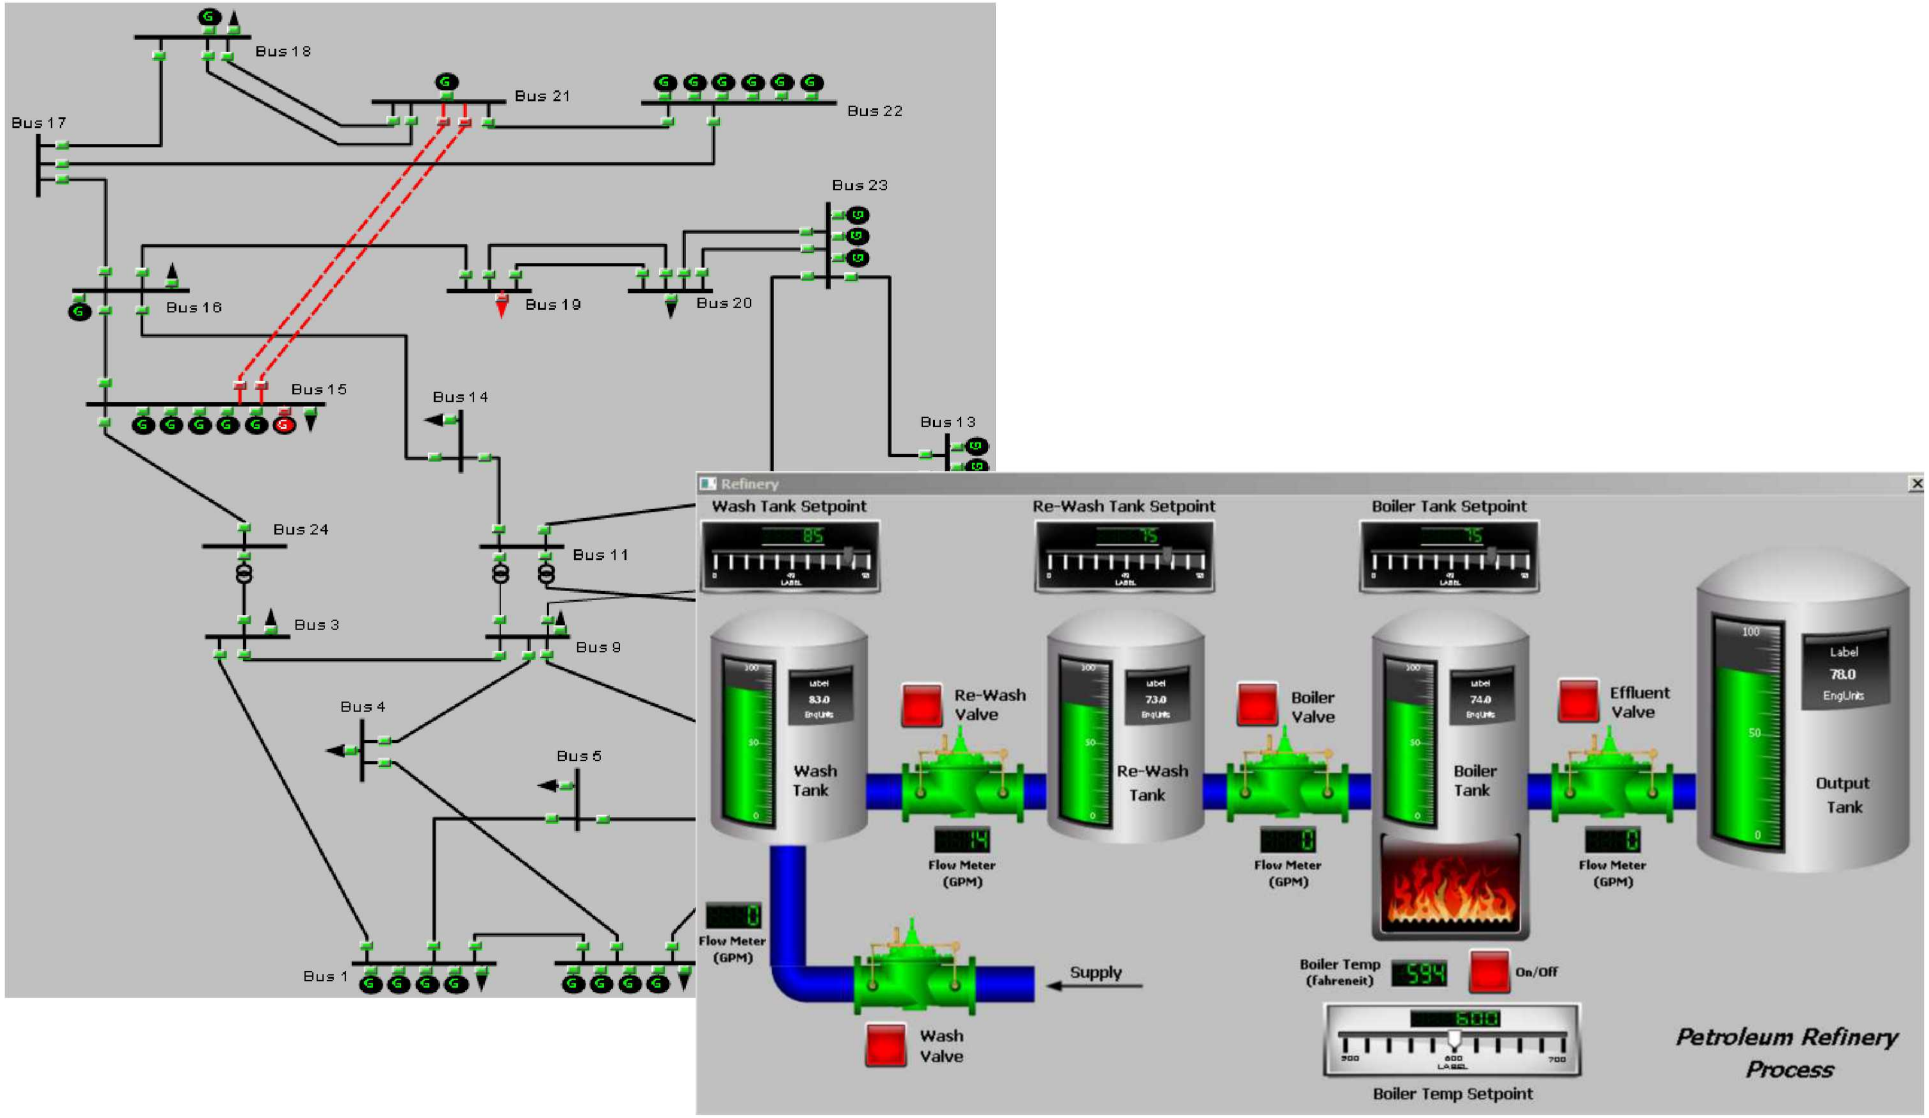

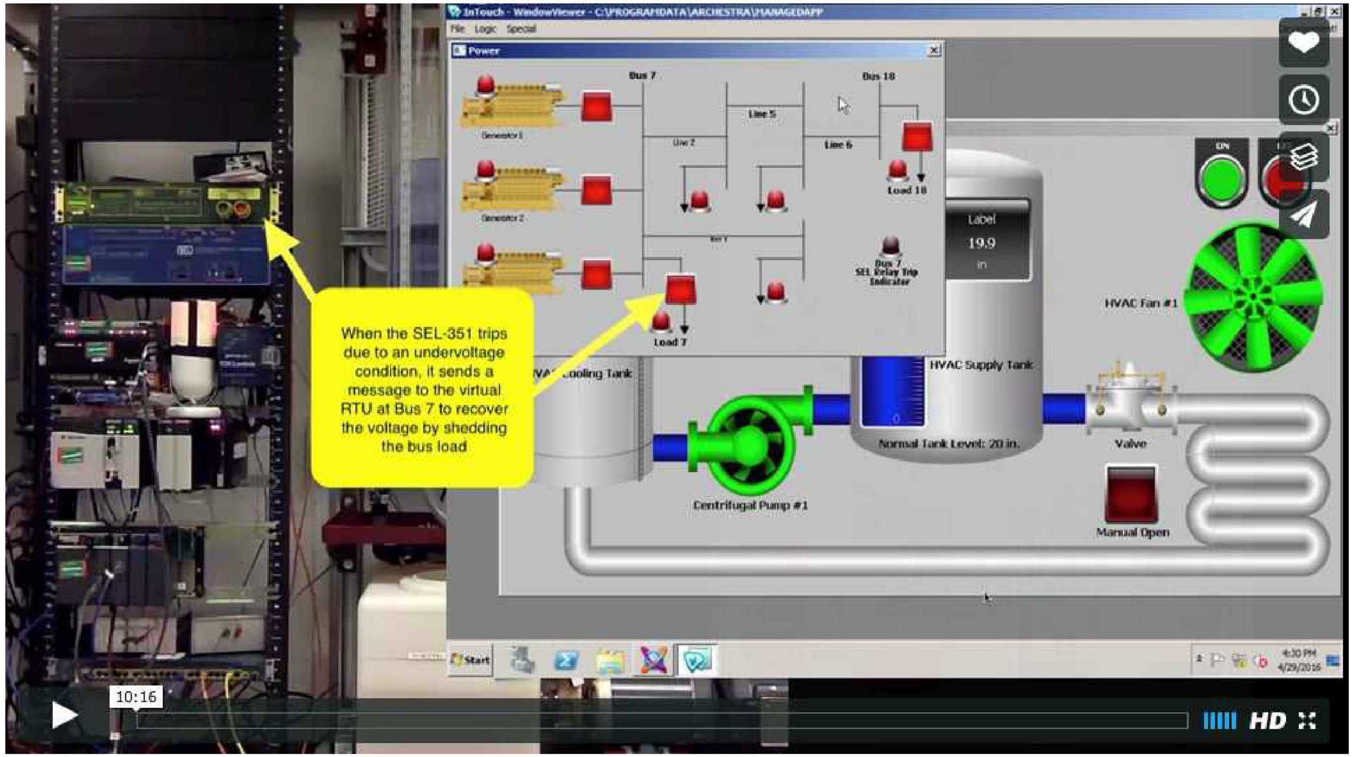

When the SEL-351 trips due to an undervoltage condition, it sends a message to the virtual RTU at Bus 7 to recover the voltage by shedding the bus load

https://vimeo.com/178492617

Electric power

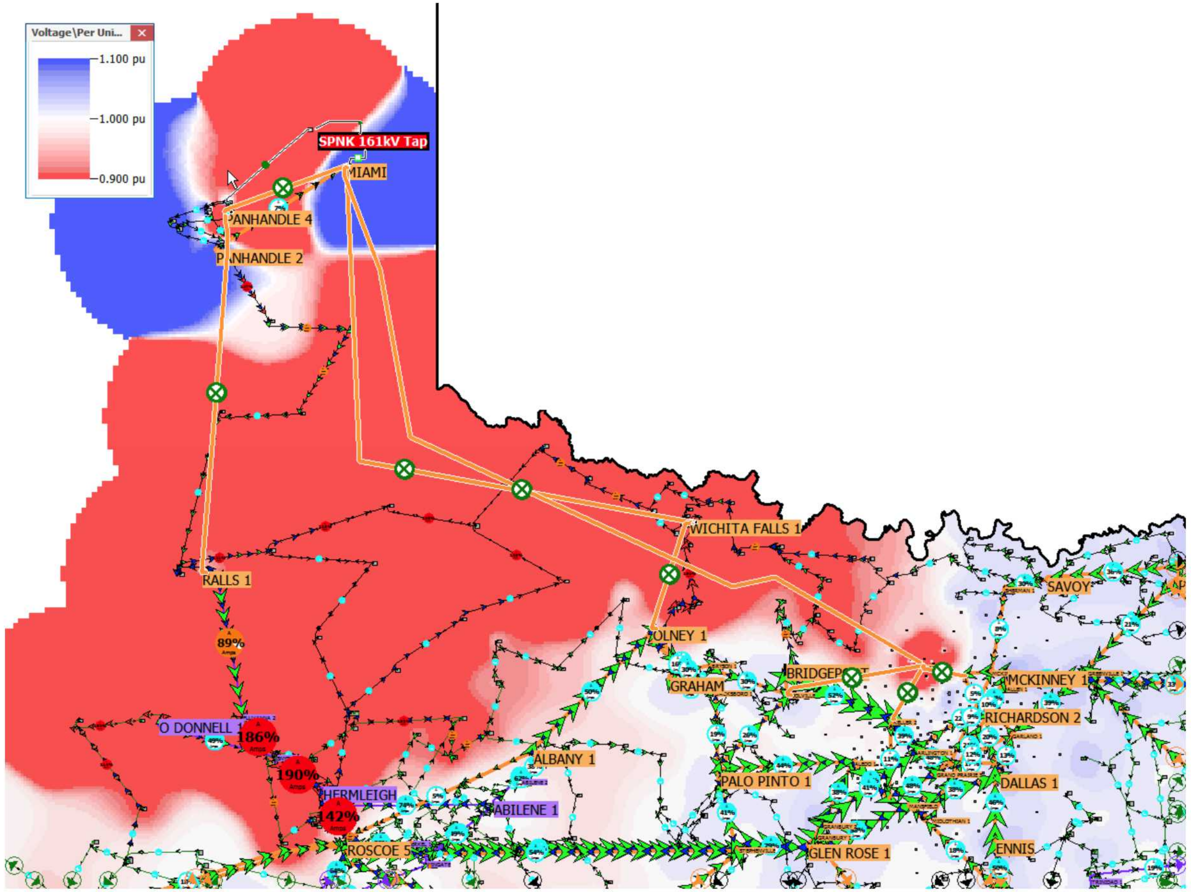

NCSC/Sandia Collaboration

- NCSC and Sandia share a lot of cyber mission space
  - Novel approaches to cybersecurity
  - R&D and application in protecting critical infrastructure
  - National security applications
- Major intersection in modeling activities
- *Can we combine efforts to create the world's foremost Emulation platform?*
  - What new inter-government research opportunities can be created out of this collaboration?

Find this interesting?

- minimega.org
- Workshop later today!