SAND2020-9158C

Cyber Threat Modeling & Validation: Port Scanning & Detection

PRESENTED BY
Eric Vugrin, Jerry Cruz, Christian Reedy, Tom Tarman, & Ali Pinar
Sandia National Laboratories

7th Annual Hot Topics in the Science of Security
Lawrence, Kansas
September 22, 2020

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.

Outline

Cyber Security Modeling Options

A Scanning and Detection Scenario

Two Analysis Approaches

Results Comparison

Cyber Modeling & Analysis: A Spectrum of Platforms

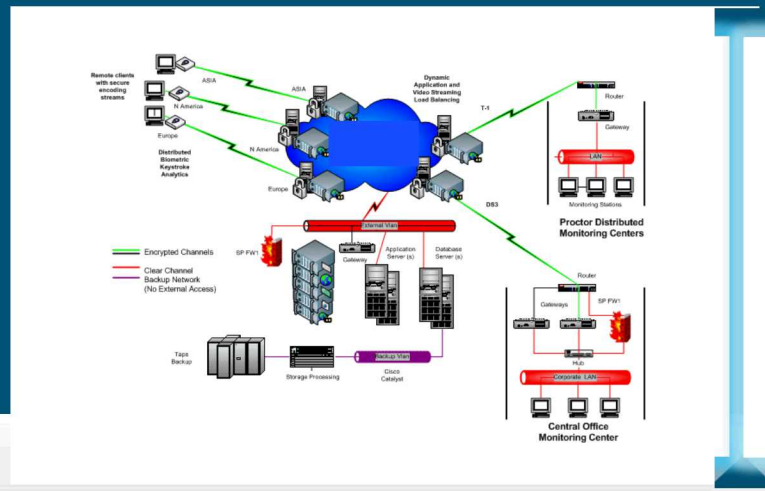

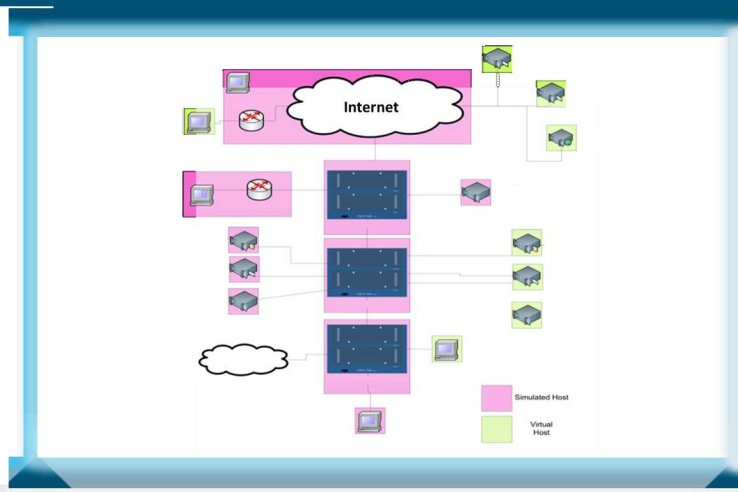

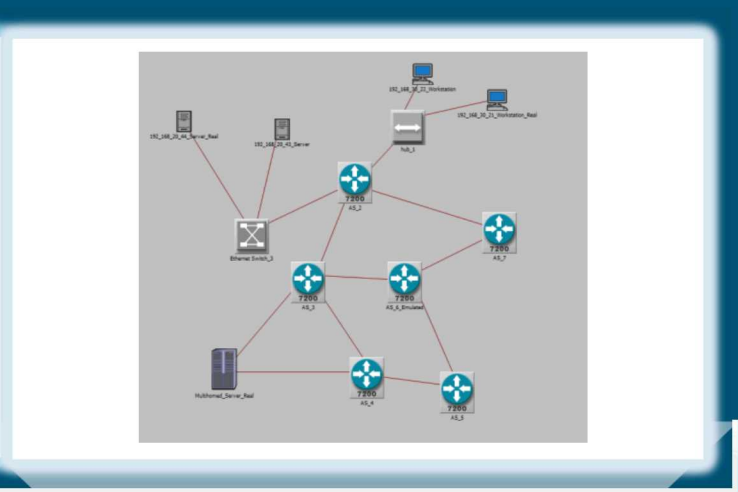

ACTUAL SYSTEM | **EMULATION TESTBED** | **MATHEMATICAL MODELING**

Interoperability in a single experiment

LIVE ← Increase Realism — Decrease Cost, Decrease Time → SIMULATED

REAL HARDWARE REAL SOFTWARE	ABSTRACT HARDWARE REAL SOFTWARE	ABSTRACT HARDWARE ABSTRACT SOFTWARE

Question: how can we use emulation test beds to develop and gain confidence in mathematical models of cyber systems?

Scenario: A Notional SCADA/ICS Network

8 substations, 24 remote terminal units (RTUs)

- Attacker scans network to find potential vulnerabilities
- Causes disruptions via RTU payloads

Defender monitors network traffic to detect attacks

Vulnerable RTUs not firewalled for maintenance

Assumptions: Attacker Tools

Nmap: Half-open SYN scan

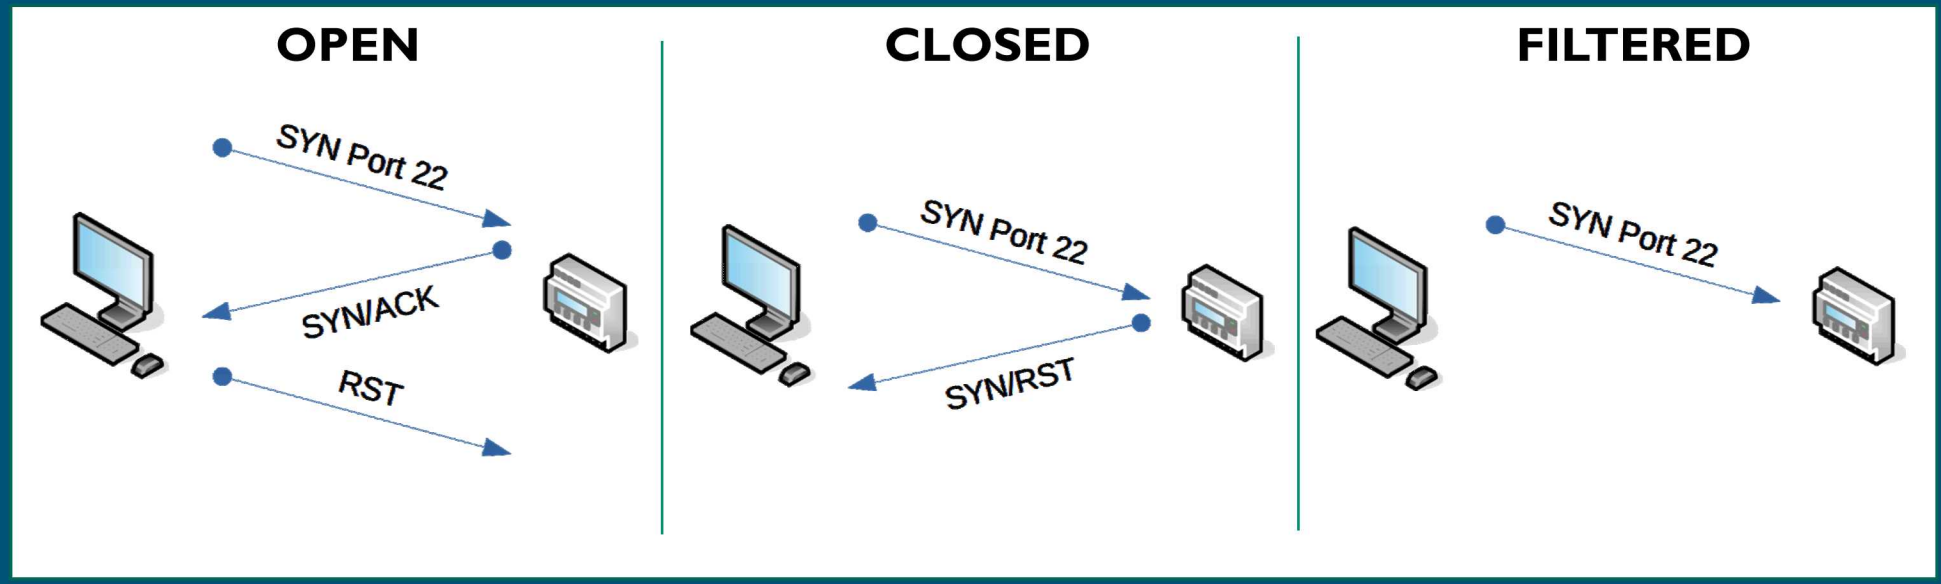

Key parameters

- **Host Group Size** – The number of hosts to scan in parallel
- **Delay** – The delay time between sequential probes

Stochastic features: ordering of addresses for scanning and time-outs

Assumption: Defender Tools

Snort: sfportscan (LOW setting)

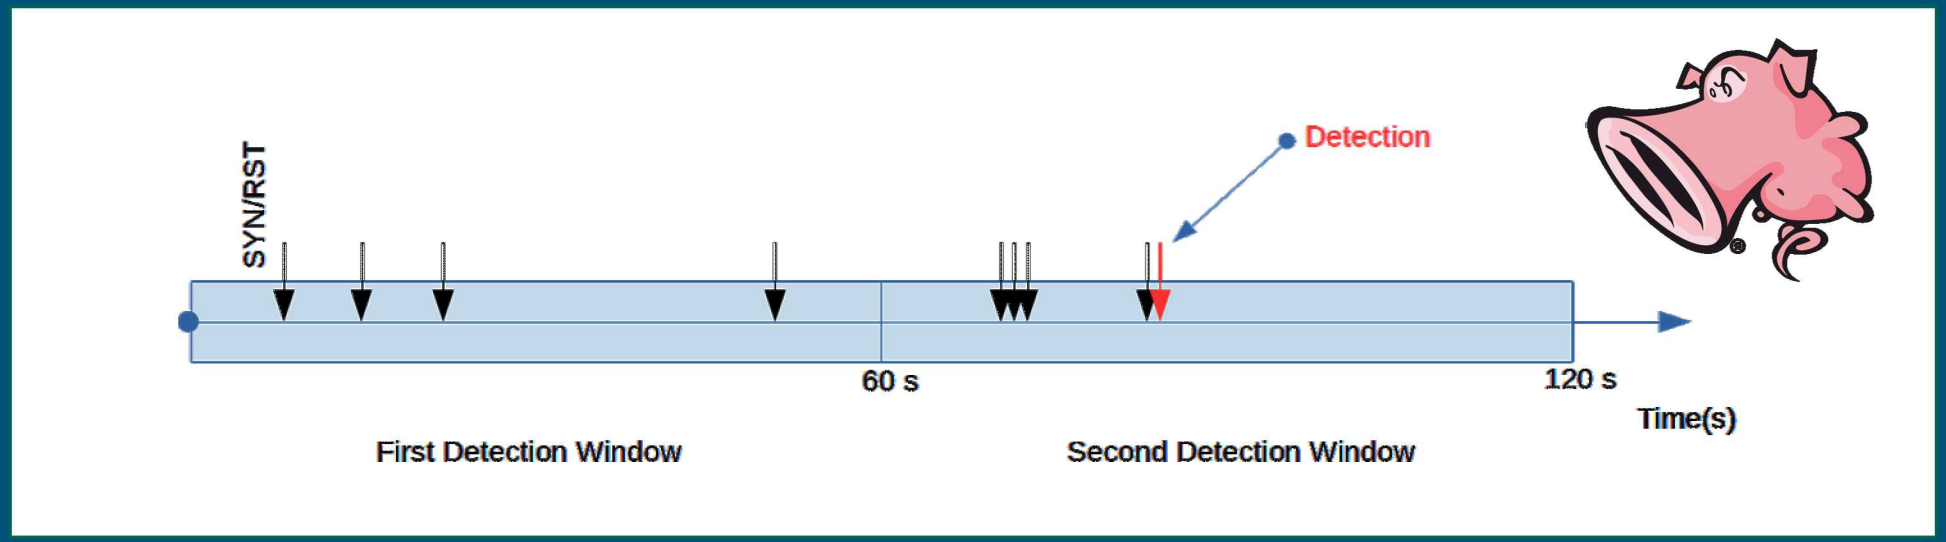

If Snort observes 5 or more TCP resets (during initial 3-way handshake) within a 60 second window, it creates an alert (i.e. detection)

An NMap probe to a closed port generates this kind of reset

Research Questions

For specified NMap and Snort settings,

- Can we estimate the rate at which the attacker identifies vulnerabilities?
- What is the probability (over time) that the attacker is detected?
- What are the associated uncertainties?
- Can we validate our estimates?

This effort developed emulations and mathematical models to analyze the scanning and detection scenario.

Virtual Testbed Set-up

Virtualization tool: minimega – launches and manages virtual machines

- Can scale to run on massive clusters
- Orchestrates Kernel-based Virtual Machines (KVM) to run unmodified OSes on emulated hardware
- Uses 802.1q VLAN tagging via Open vSwitch to support arbitrary network topologies

(In-experiment) Software

- Node OS: pared down Ubuntu 18.04
- Snort 2.9.13
- Nmap 7.60
- Router OS: VyOS 3.13.11

Host hardware

- Dual Socket Intel E5-2683v4 2.10GHz CPUs (32 total cores)
- 512 GB DDR3 Memory
- 100 GbE experiment network
- 10 GbE boot/storage network

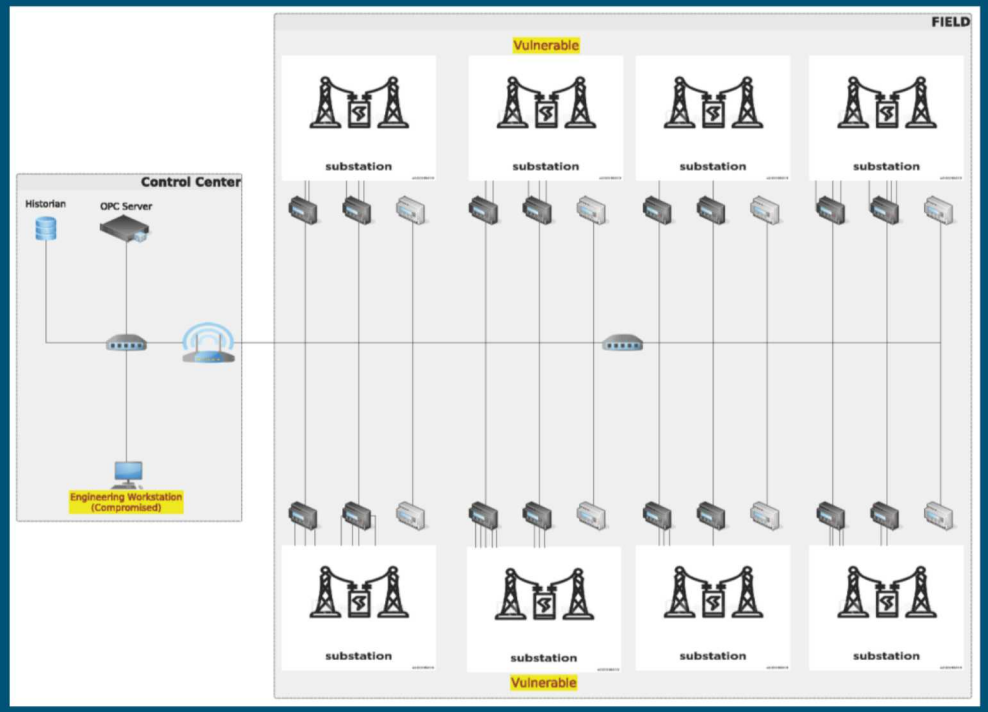

Mathematical Model

T=0

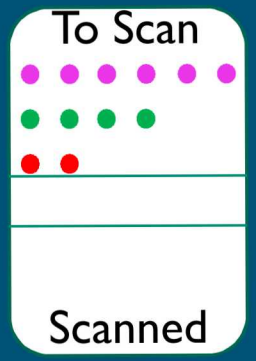

Step 1: initial conditions

Mathematical Model

T=0

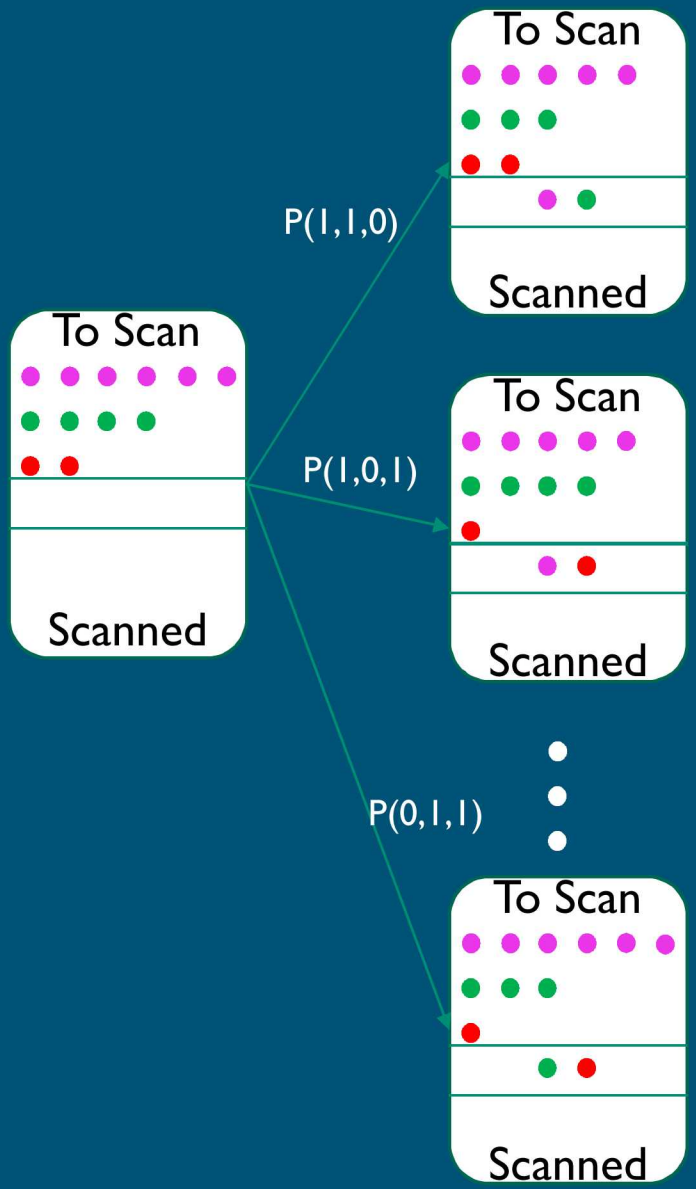

Step 2: select RTUs to scan

Mathematical Model

T=0

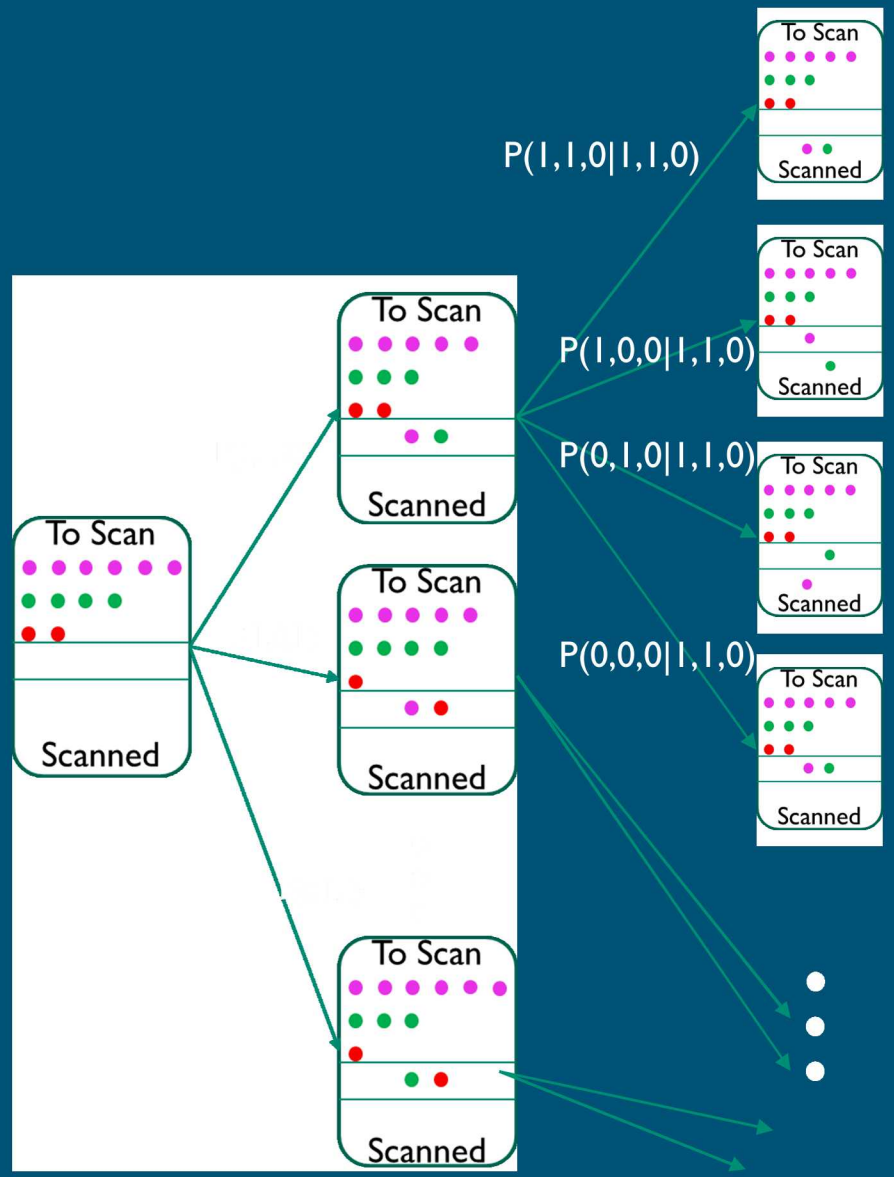

Step 3: determine if scan succeeds or times out

Mathematical Model

T=0

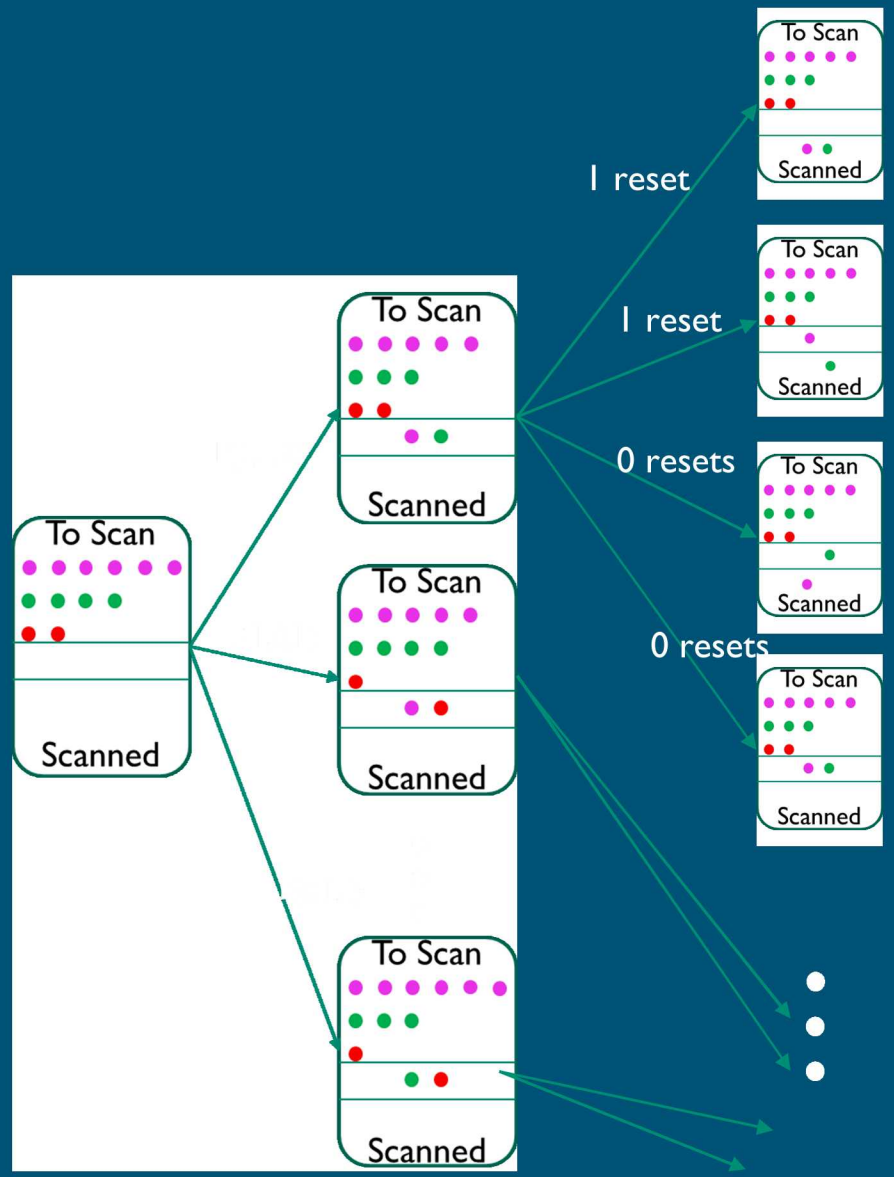

Step 4: determine if TCP resets occurred

Mathematical Model

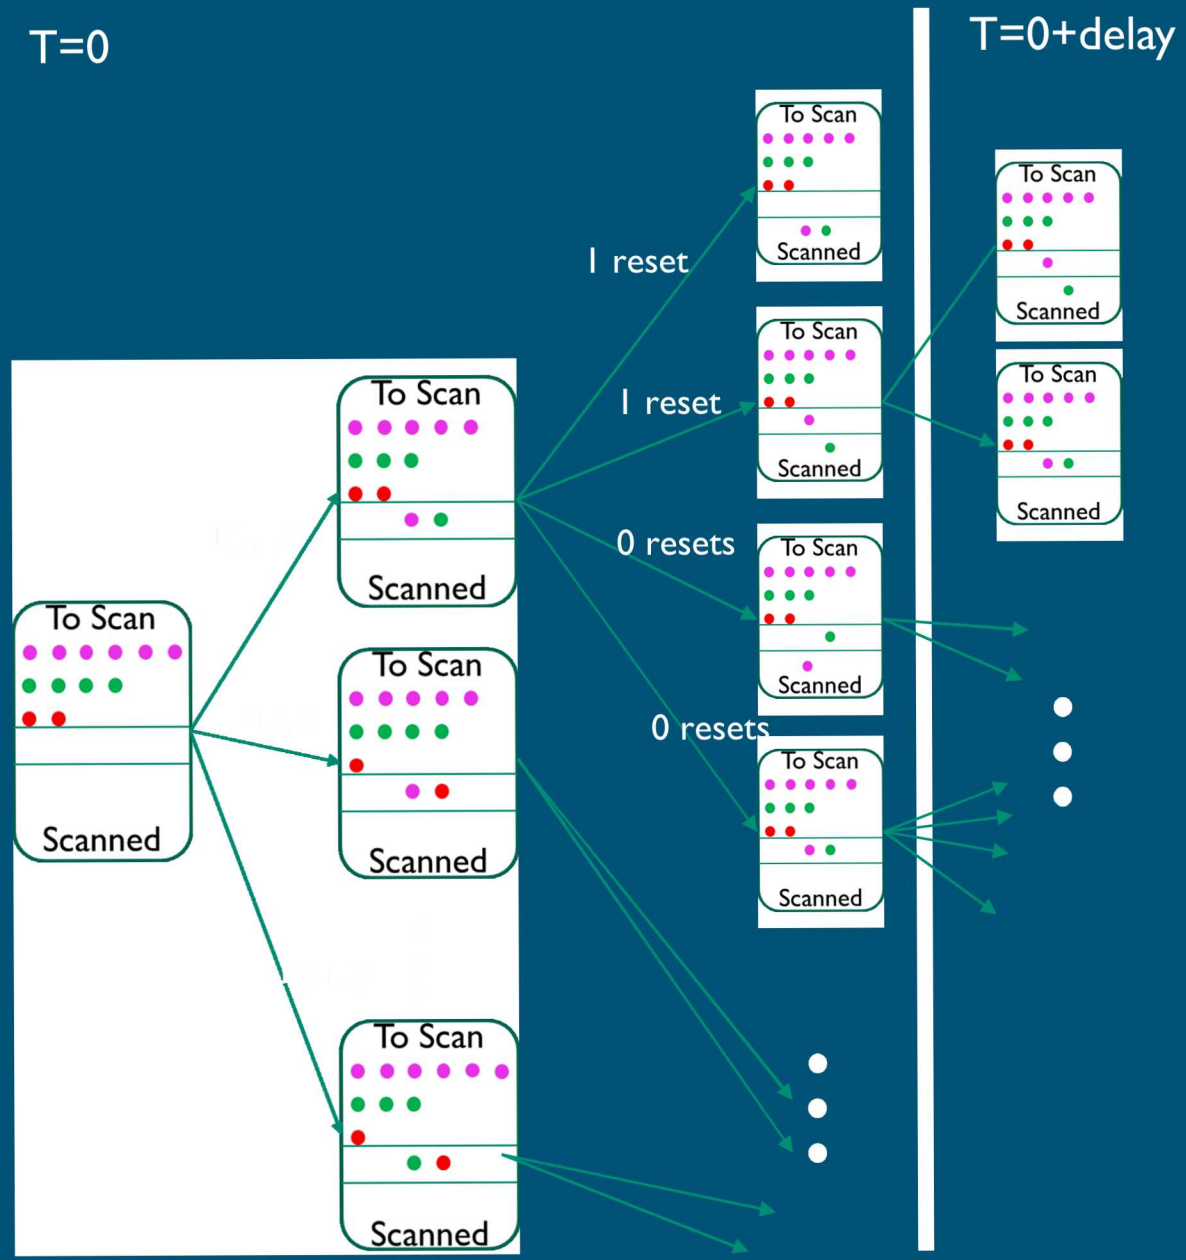

Step 5: if time outs occurred, if time outs occurred, repeat steps 2-4 for timed out RTUs

Mathematical Model

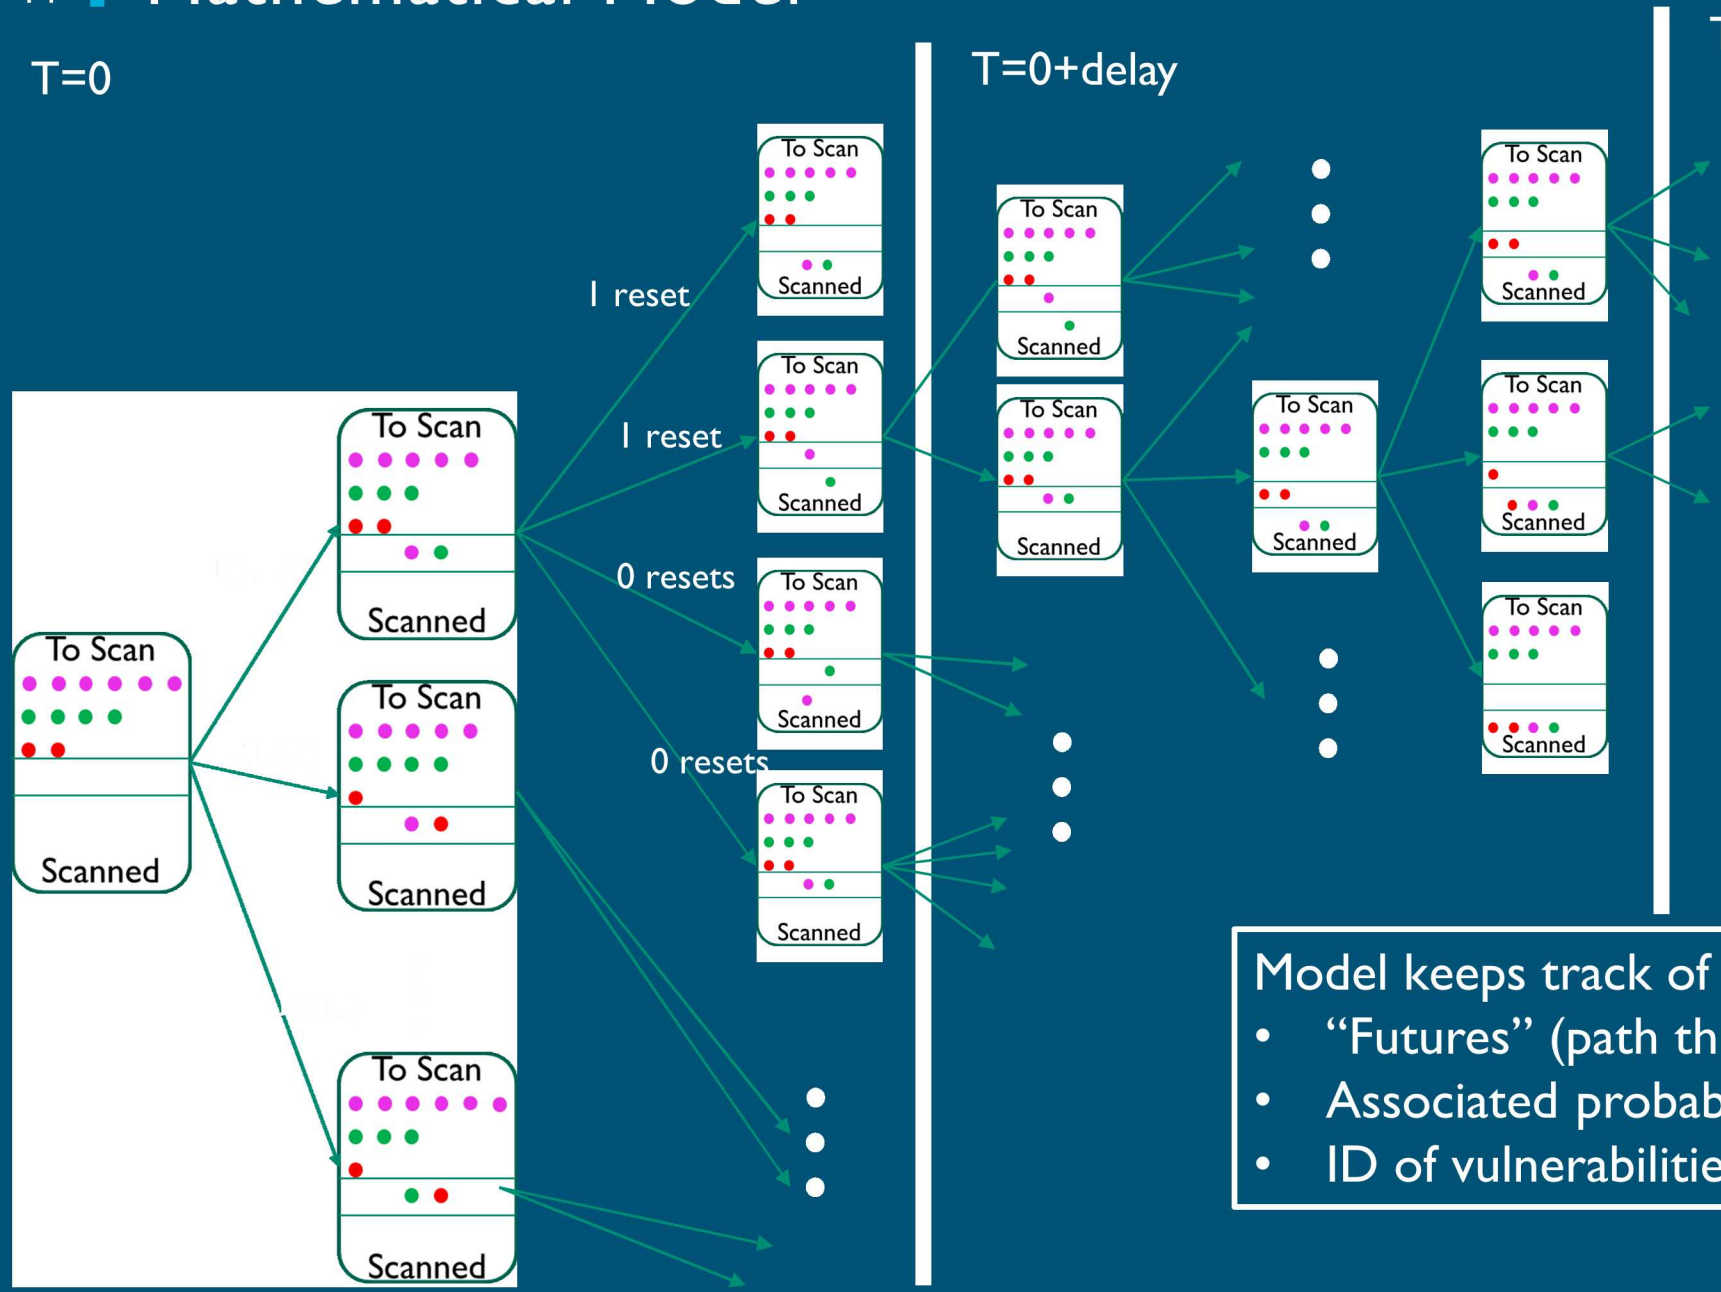

Model keeps track of

- "Futures" (path through the tree)
- Associated probabilities
- ID of vulnerabilities and TCP resets

Example Results

System settings
- 4 open (aka vulnerable) RTUs
- 8 closed RTUs
- 12 filtered RTUs
- Probability of probe time out = 0.1

NMap settings
- Host group: 4
- Scan delay: 10s
- Max # of retries: 1

Snort setting:
- Low sensitivity

Emulation experiments: 1000 trials

Results: Attacker Progress

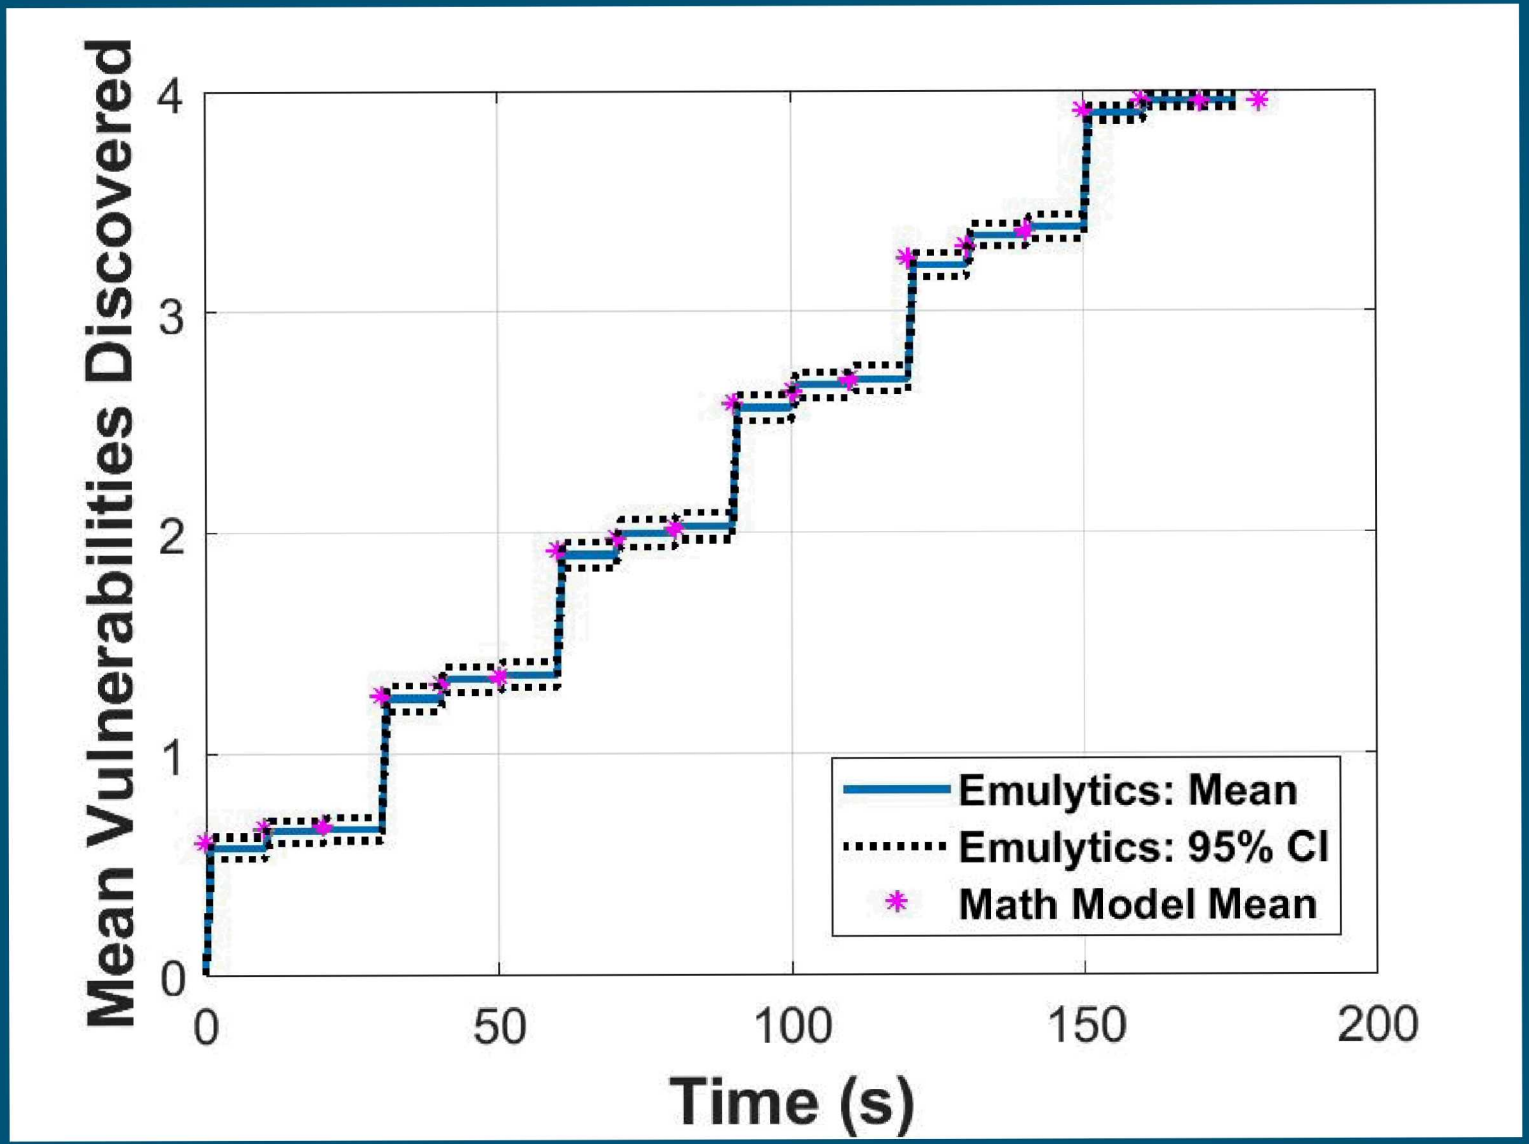

System Parameters
- 24 hosts up
- 4 open (susceptible to CRASH payload)
- 8 closed (inactive RTUs)
- 12 filtered (active but firewalled)
- Timeout prob: 0.1

Nmap setting
- Host group: 4
- Scan delay: 10s
- Max retries: 1

Snort sfportscan setting: low

1000 Emulytics Runs

Results: Detection Probabilities

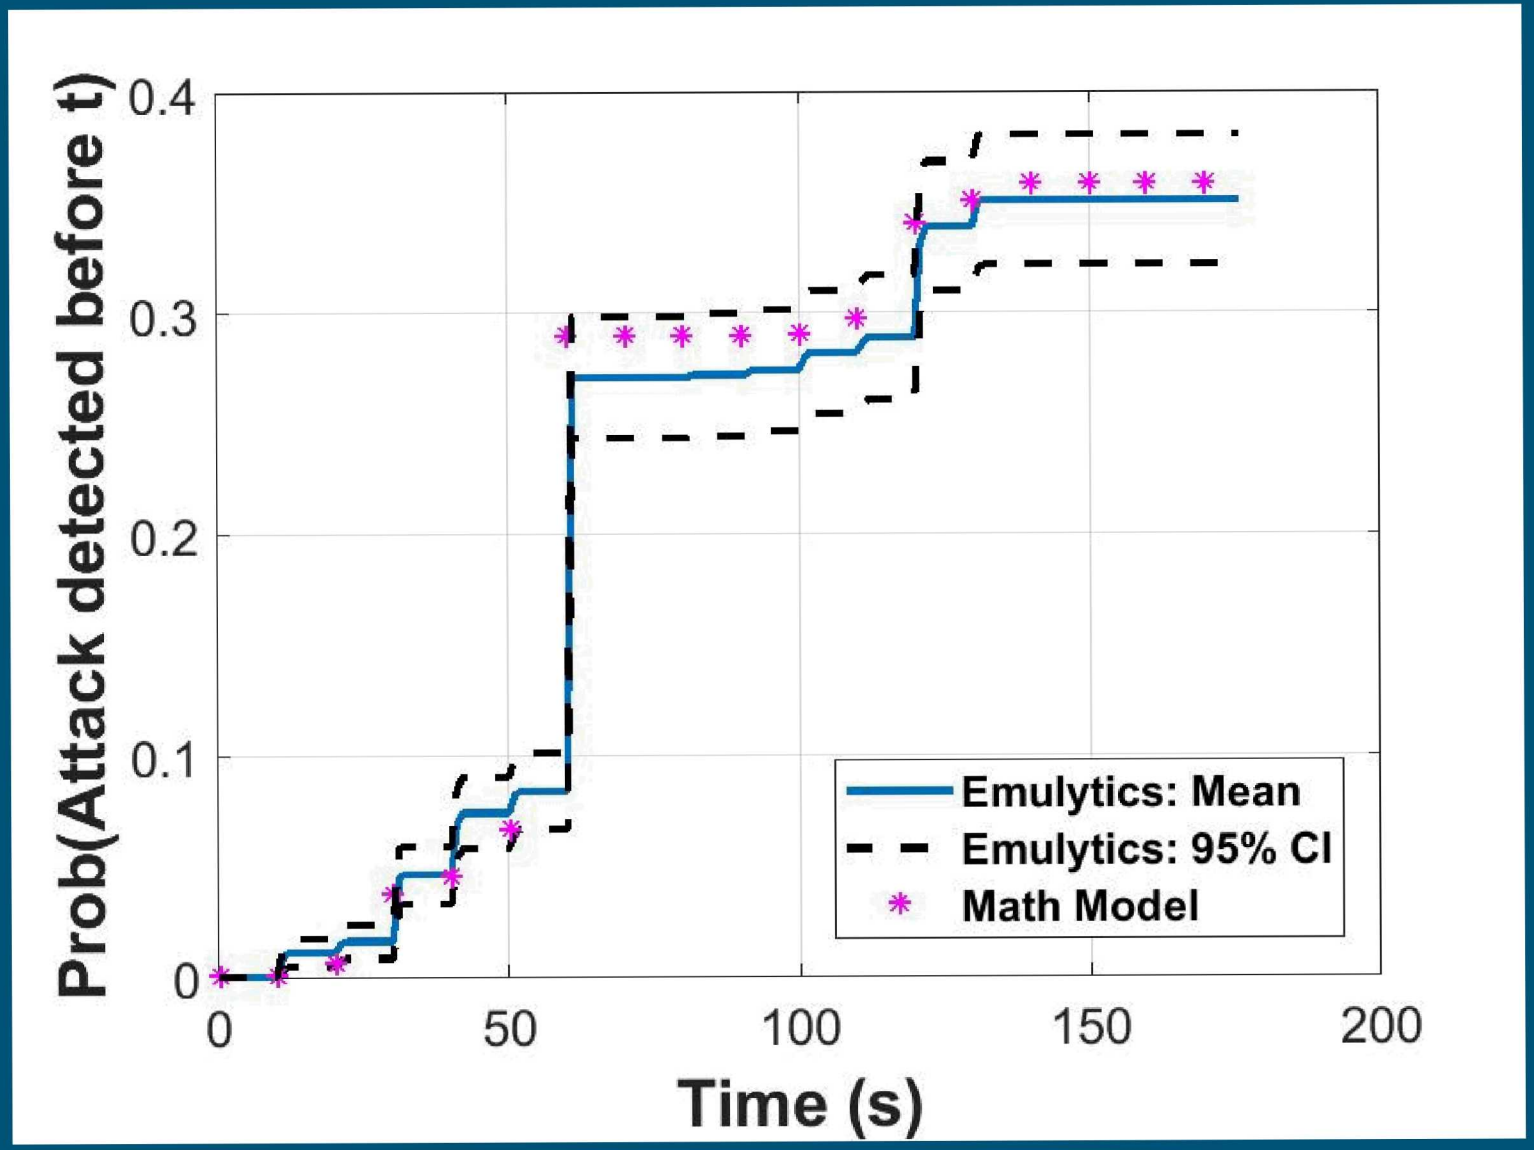

System Parameters
- 24 hosts up
- 4 open (susceptible to CRASH payload)
- 8 closed (inactive RTUs)
- 12 filtered (active but firewalled)
- Timeout prob: 0.1

Nmap setting
- Host group: 4
- Scan delay: 10s
- Max retries: 1

Snort sfportscan setting: low

1000 Emulytics Runs

Summary and Insights Gained

This effort modeled the reconnaissance portion of a hypothetical grid attack

- Developed mathematical model of model scanning and detection
- Emulation testbeds provided means of evaluating models, increasing confidence

Challenges:

- Discrete vs. continuous time comparisons
- Scale

Future extensions

- Include different scanning and detection tools
- Scale
- Physical Impacts
- Compare with "real" network