



LANCER

Cornell Site Visit
November 30, 2023

Outline

- [15 minutes] Introduction (Nate & Wen)
 - Team Introductions
 - Technical Approach
- [10 minutes] Progress Since Kick-Off (Nate & Wen)
 - Executive Summary
 - Planned Trajectory for end of Phase I
- [15 minutes] Collaboration Efforts (Nate & Rebecca)
 - CAGE
 - Talking to Kryptowire
 - Network Action Space
- [60 minutes] Early Results
 - [20 minutes] NetKAT (Jules & Nate)
 - [30 minutes] Inverse RL (Nico/Rebecca & Wen)
 - [10 minutes] Aether: Pronto + OnRamp (Hussain & Nate)
- [20 minutes] Response to Crawl Questions (Everyone)
- [30 minutes] Budget & Contracting (Shailja & Nate)

Introduction



Progress




Progress

- Got going with Kryptowire TA1 Platform
- Started development using CAGE 2
- Started Modeling Red Agents Using Inverse RL
- Fast NetKAT implementation
- Standing Up Aether OnRamp

Trajectory

- Crawl (6 month)
- Walk (6 month)
- Run (6 month)

Collaboration Efforts



Cage Challenge: Overview

- Scenario of a network attack
 - **Red Agent** (malicious): infiltrates network
 - **Blue Agent** (defensive): protects the network
 - **Green Agents** (neutral users): generate noise
- Integrated with CybORG, a reinforcement learning gym




Figure 1: Network Topology (Cage Challenge 2)

Cage Challenge: Red Agent Actions




Figure 2: Effect of actions on host state (Cage Challenge 2)

Early Results

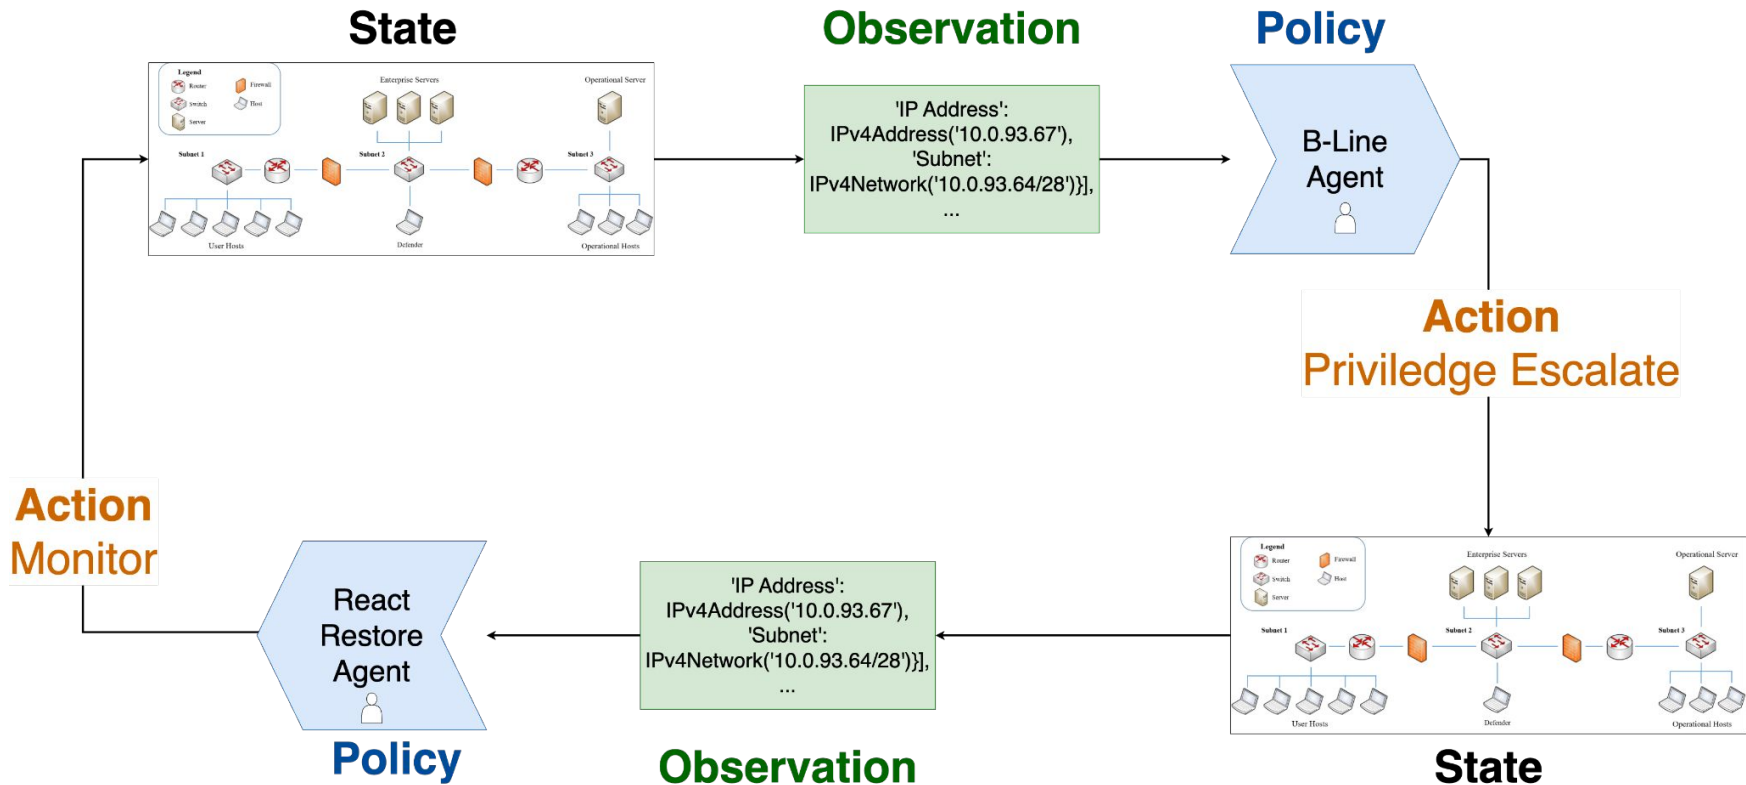


Learning Approach for Modeling Red Agents

Imitation Learning: learn red agents' behavior from their traces

- Real-world scenario: only have examples (data) of network exploit (i.e. Red agent infiltration)
 - No access to novel Red agents for simulation
- Once red agents are learned: train blue agents against them
 - Targeted RL training
 - Adversarial RL training: train Blue and Red to fight each other
 - Often results in very conservative behaviors

Reinforcement Learning Terminology




Reinforcement Learning Terminology

States: configuration of the environment

Observation: environment information observed by an agent

Policy: how an agent decides what action to take

Rollout: a sequence of states, actions, and associated reward



Behavior Cloning (BC)



Behavior Cloning (BC)

1. Collect data from environment with Blue, Green, Red agents
 - (Blue agent observation, Red agent action)
2. Train neural network on collected data
 - Blue agent observation \rightarrow *predicted* Red agent action
3. Created a learned Red agent: used trained neural network as policy
4. Collected reward during rollout: environment with Blue, Green, and learned Red agent
 - Measure of learned Red agent's quality: reward collected during rollout

BC: 1 Input Observation

Red Agent	Blue Agent	Training Metrics			Learned Red Agent		True Red Agent	
		Train Loss	Train Accuracy	Validation Accuracy	Reward	Standard Deviation	Reward	Standard Deviation
B-Line	React Remove	0.16	0.95	0.93	556	361	947	193
B-Line	React Restore	0.64	0.77	0.77	-10.0	0.0	508	366
Meander	React Remove	0.71	0.72	0.67	11.1	39.5	630	259
Meander	React Restore	1.10	0.56	0.53	3.55	7.77	185	210


BC: 4 Input Observations

Red Agent	Blue Agent	Training Metrics			Learned Red Agent		True Red Agent	
		Train Loss	Train Accuracy	Validation Accuracy	Reward	Standard Deviation	Reward	Standard Deviation
B_Line	React Remove	0.038	0.986	0.967	694	305	947	193
B-Line	React Restore	0.0372	0.987	0.965	484	336	508	366
Meander	React Remove	0.327	0.870	0.710	255	246	630	259
Meander	React Restore	0.615	0.762	0.587	77	141	185	210

BC Plot: Reward vs. Number of Input Observations

BC Plot: Reward vs. Dataset Size


Results




Issue of Behavior Cloning: Distribution Mismatch

Learning to Drive


Compounding error makes learner deviate from the expert track quickly



Inverse RL to the Rescue



1. Inverse RL aims to learn a reward model from the data (e.g., red agent's reward function when they plan attacks)
2. It then learns a policy to optimize the learned reward
3. The learned policy acts as the predictive model for the red agent



IRL Results

Next Steps

1. New IRL algorithms for improving modeling red agents;
2. Training RL agents against the learned red agents



KATch

A Fast Symbolic Verifier for NetKAT

Mark Moeller, Jules Jacobs, Nate Foster, Alexandra Silva (Cornell), Olivier Savary Belanger, David Darais, Cole Schlesinger (Galois), Steffen Smolka (Google)

The Control Plane and Network Defense Agents

Control Plane

- Computes routing tables
- Ensures network connectivity
- Enforces network policies

Network Defense Agents

- Detects and responds to network attacks
- Example: Security breach containment
- Example: DDoS mitigation
- **Action space?**
- **Modify routing tables?**

Neural Strengths

- General pattern recognition
- Learns from experience
- Adaptability to new situations
- Ideal when explicit programming is difficult

Symbolic Strengths

- Domain specific reasoning
- Guarantees correctness
- Verifiable and explainable
- Ideal when strict compliance with rules is required

Neural+Symbolic AI in Network Defense: Idea

Neural

- Utilizes deep learning for real-time attack detection and response
- Adapts to evolving network threats
- Modifies routing tables dynamically
- Example: Detecting and rerouting traffic to mitigate DDoS attacks
- Example: Detecting and isolating compromised hosts

Symbolic

- Computes consequences of routing changes
- Ensures correctness of routing tables
- Verifies adherence to network policies and security rules
- Example: Validating routing paths for security compliance
- Example: Verifying reachability of critical network services

NetKAT: network specification language for SDN

- Network topology
- Routing tables
- Network-wide policies

Verification of network policies

- Security properties, e.g. slice isolation
- Operational properties, e.g. reachability
- Verified in a common framework

Problem: NetKAT verification is slow
Not suitable for real-time network defense




KATch

A Fast Symbolic Verifier for NetKAT

A new NetKAT verifier that is

- **Fast:** $1000\times$ faster
- **Symbolic:** explains verification failures
- **Scalable:** handles larger networks




Full Reachability



Detailed comparison: (un)reachability and slice isolation

Name	Size (atoms)	Reachability		Unreachability		Slicing		Min Speedup
		KATch	Frenetic	KATch	Frenetic	KATch	Frenetic	
Layer42	135	0.00	0.04	0.00	0.04	0.01	0.07	7×
Compuserv	539	0.01	0.36	0.01	0.38	0.01	0.85	36×
Airtel	785	0.01	0.83	0.01	0.84	0.02	2.08	83×
Belnet	1388	0.01	3.17	0.01	3.16	0.04	7.99	200×
Shentel	1865	0.02	4.01	0.02	4.00	0.04	9.80	200×
Arpa	1964	0.01	4.32	0.02	4.32	0.05	10.99	216×
Sanet	4100	0.04	23.46	0.03	25.23	0.12	62.70	522×
Unet	5456	0.04	81.54	0.04	81.92	0.15	204.85	1366×
Missouri	9680	0.11	161.28	0.10	165.85	0.27	519.46	1658×
Telcove	10720	0.09	464.15	0.08	465.27	0.28	1274.24	4551×
Deltacom	27092	0.31	2392.56	0.30	2523.03	0.75	7069.54	7718×
Cogentco	79682	0.97	22581.39	0.88	23300.87	1.78	53066.82	23280×

Synthetic combinatorial benchmarks



Conclusion

NetKAT verification can be fast

Can we combine neural and symbolic AI?


To reach challenging locations

Ultra-low energy
10+ years of battery life

**Massive
Internet of
Things**

Ultra-low complexity
10s of bits per second

Ultra-high density
1 million nodes per km²



To reach challenging locations

Strong security

e.g., Health / government / financial trusted

Ultra-high reliability

$< 10^{-5}$ per 1 millisecond

Ultra-low latency

As low as 1 millisecond

Extreme user mobility

Up to 500 km/h

Mission-critical control

Massive Internet of Things

Ultra-low energy

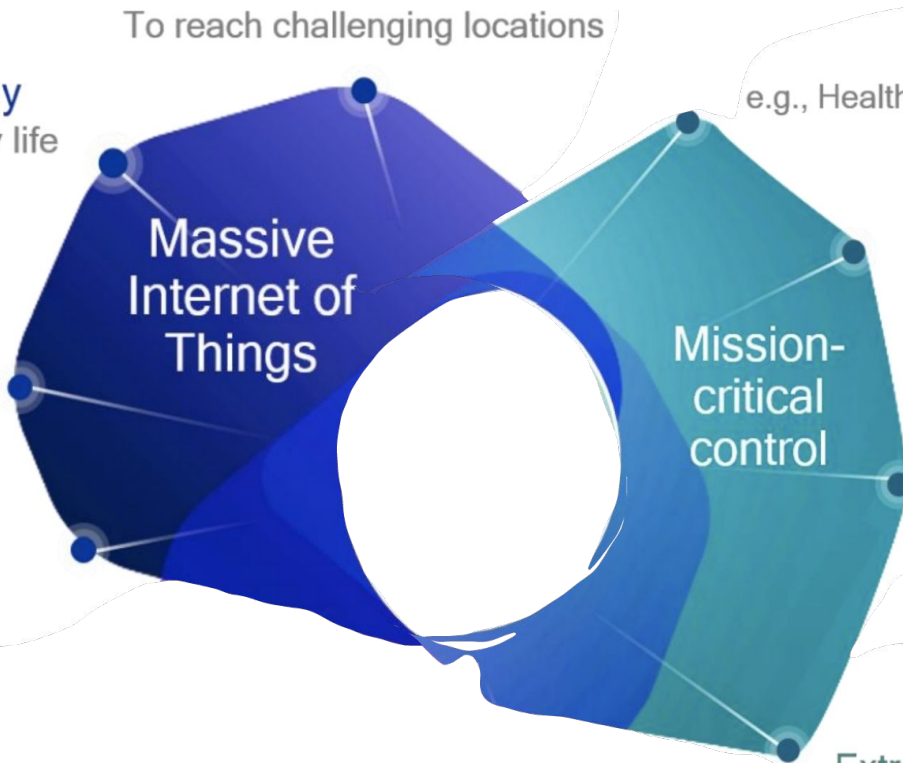
10+ years of battery life

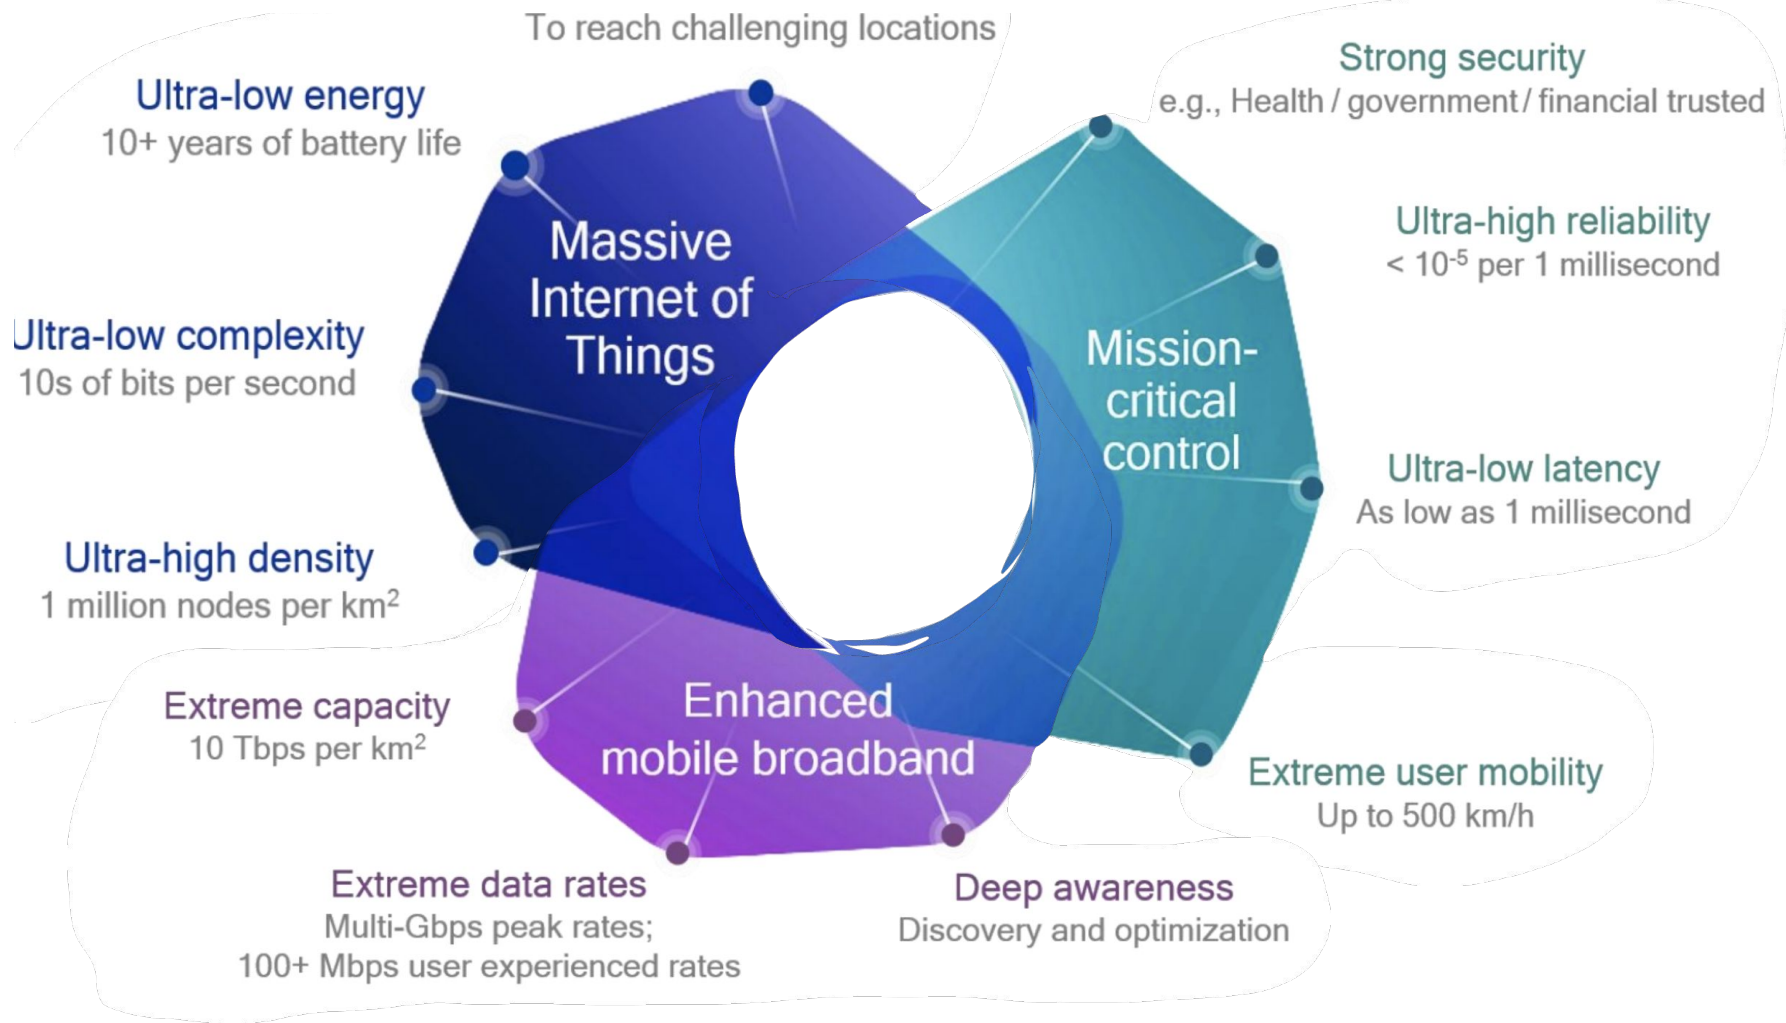
Ultra-low complexity


10s of bits per second

Ultra-high density

1 million nodes per km²







5G networks -

5G Mobile Network two main subsystems :

1. RAN - manages radio resources(spectrum)
2. Mobile Core - provide packet data network to mobile subscribers

AetheronRamp - Private Enterprise 5G network

- operational cluster that is capable of running 24/7 and supports live 5G workloads.
- Cluster containerizing subsystems components, can scale horizontally with dynamic workloads.

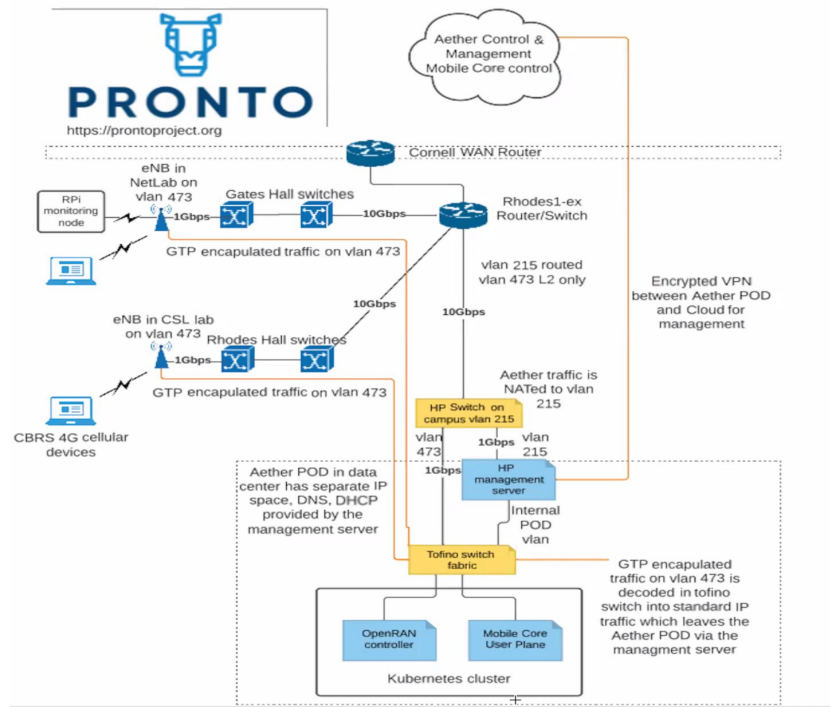
Pronto & AetherOnRamp Demo

Pronto 4G network

- Current testbed located at Gates Lab and Robotics Lab.
- Supports both direct access 4G connectivity, extended with APN connectivity

AetherOnRamp 5G network

- Work in Progress, currently emulate UEs(mobile devices) control and data plane connectivity



Crawl Questions

- Learn about one another's approaches, find integration points, and collaborate on shared infrastructure
- What network should we model first and what workflows should be present?
- What agent actions will be simulated and executed?
- What is a 'good' resiliency criteria and how will we judge whether your approach is successful?
- What data types are needed for each performer and what data can be provided by each performer?
 - Data for attackers
 - Reward function for defenders (domain knowledge, Inverse RL)
- How do we collaborate on API design and code interfaces?
- What open-source technology can enable an end-to-end integration demo quickly?
- Who is the intended operator of your approach and what is the desired impact/benefit to their job?