Leveraging Prefix Structure to Detect Volumetric DDoS Attack Signatures with Programmable Switches

Chris Misa
Ramakrishnan Durairajan
Reza Rejaie
University of Oregon

Arpit Gupta — UCSB

Walter Willinger — NIKSUN, Inc.

2024-05-21

Background: IP Addresses

$$\overbrace{\{0.0.0.0,\dots \qquad \dots,128.0.0.0,\dots \qquad \dots,255.255.255.255\}}^{\text{IPv4 “Address space”}}$$

Background: IP Address Prefixes

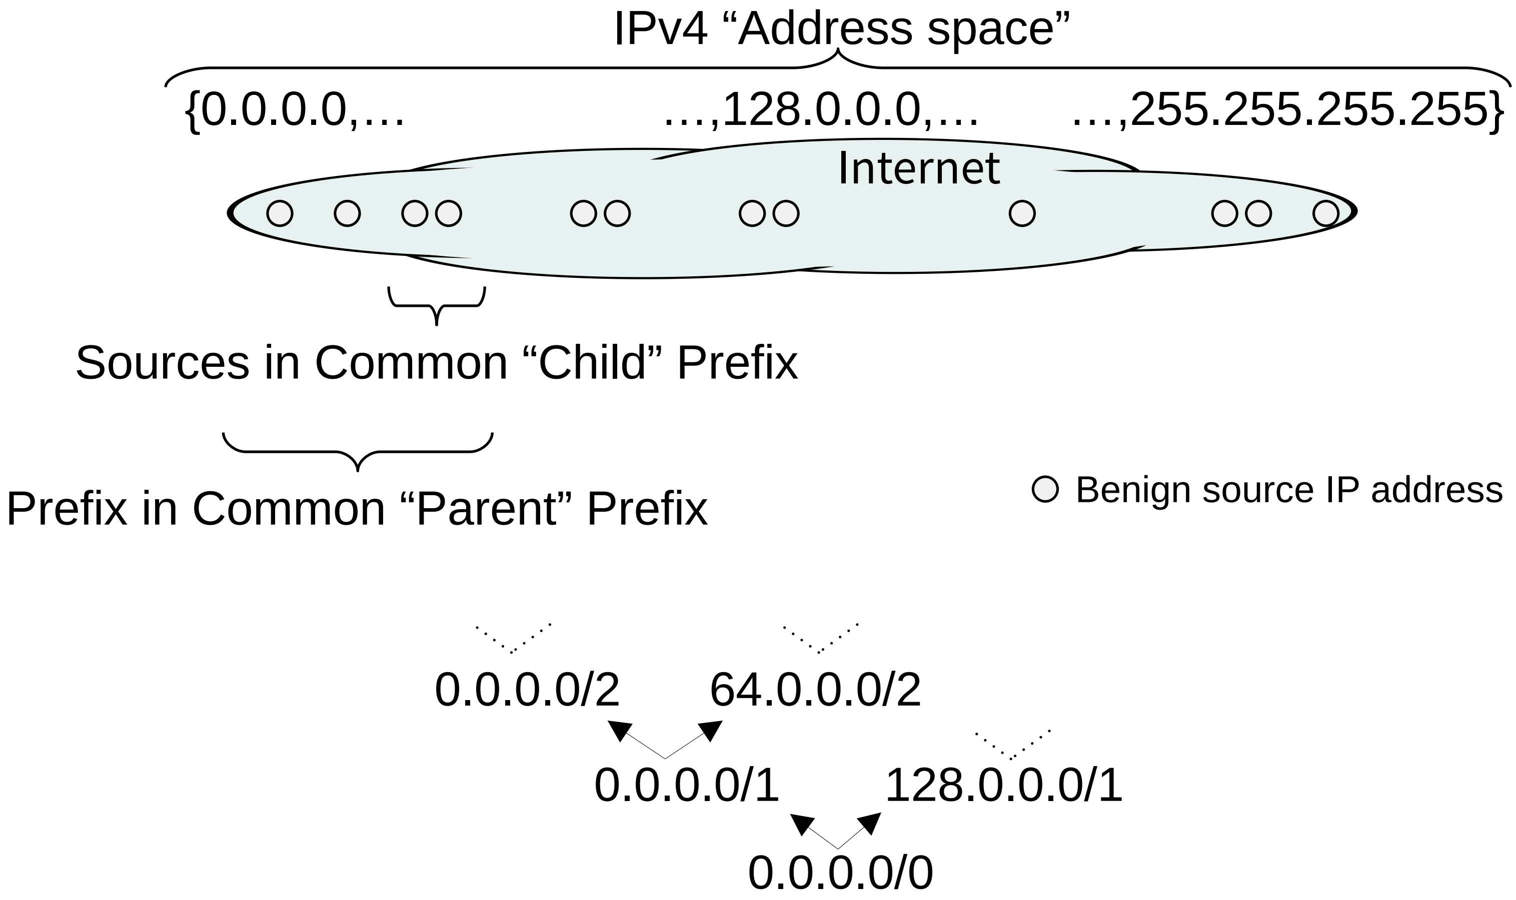

Setting: Edge Networks

External source
IP addresses

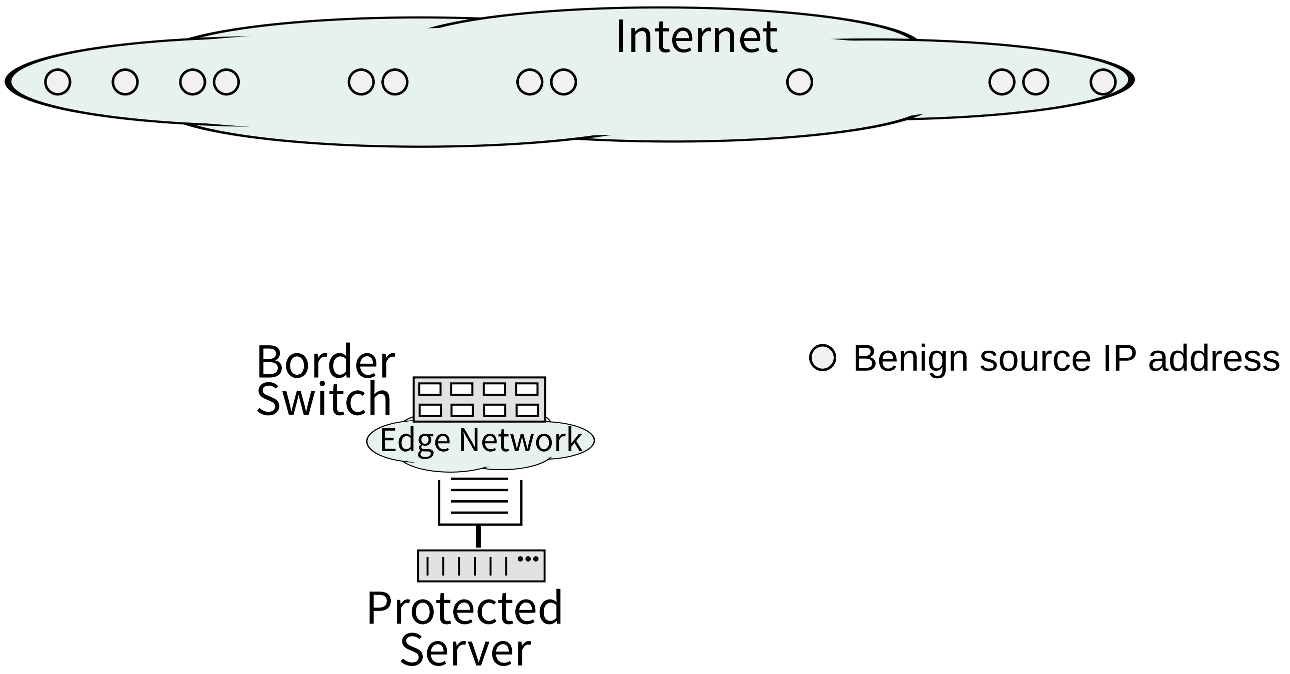

Setting: Edge Networks

External source
IP addresses

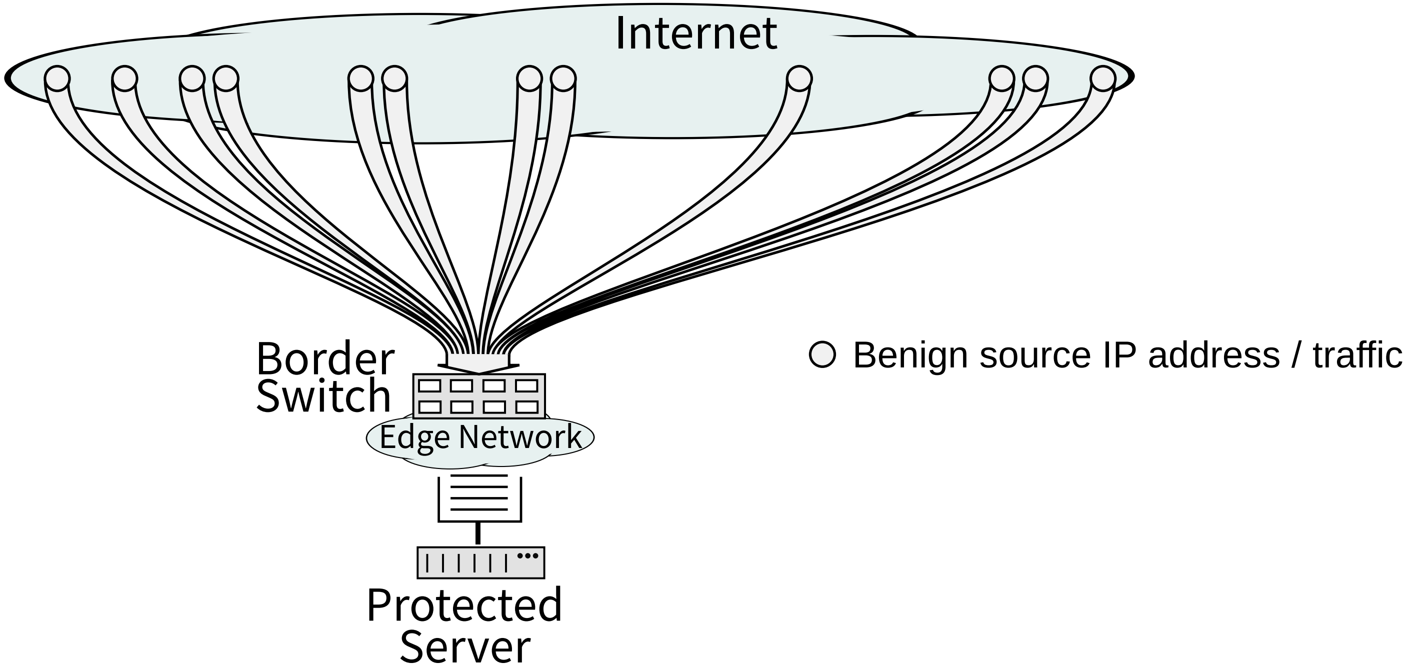

Setting: Volumetric DDoS Against Edge Networks

Benign traffic is affected!

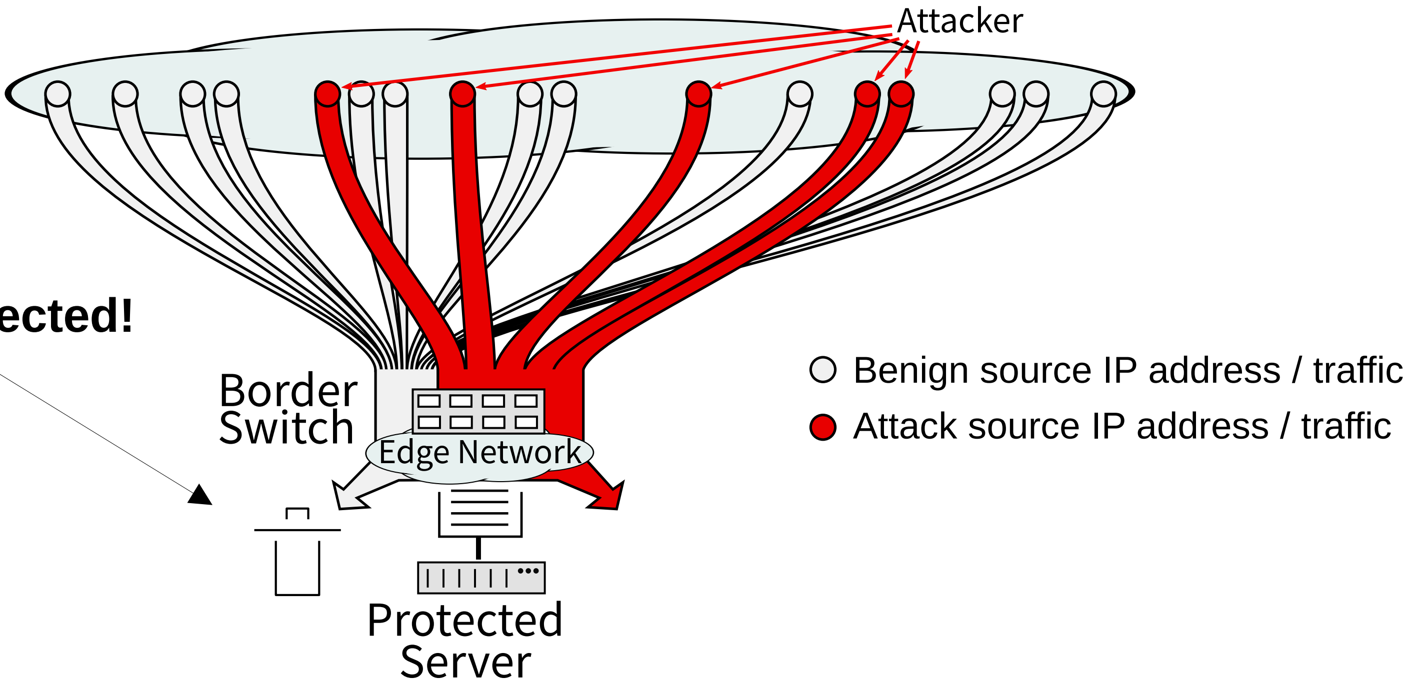

Setting: Dynamic Volumetric DDoS Against …

Attack sources can change!

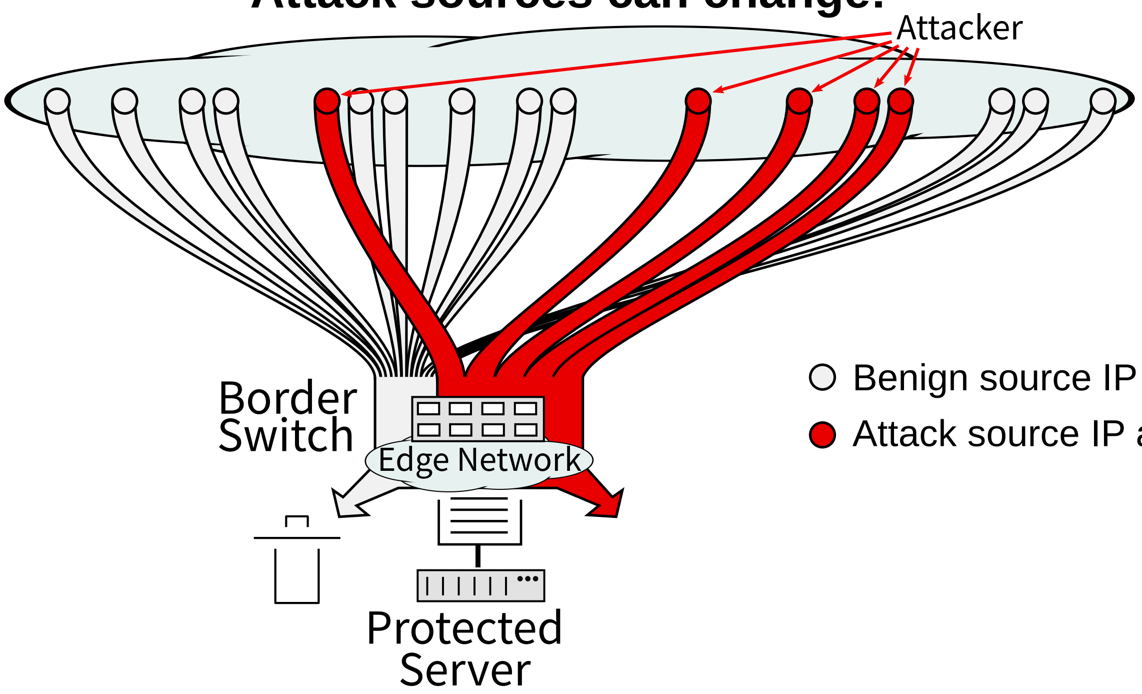

- Benign source IP address / traffic
- Attack source IP address / traffic

Need: Volumetric DDoS Signature Detection

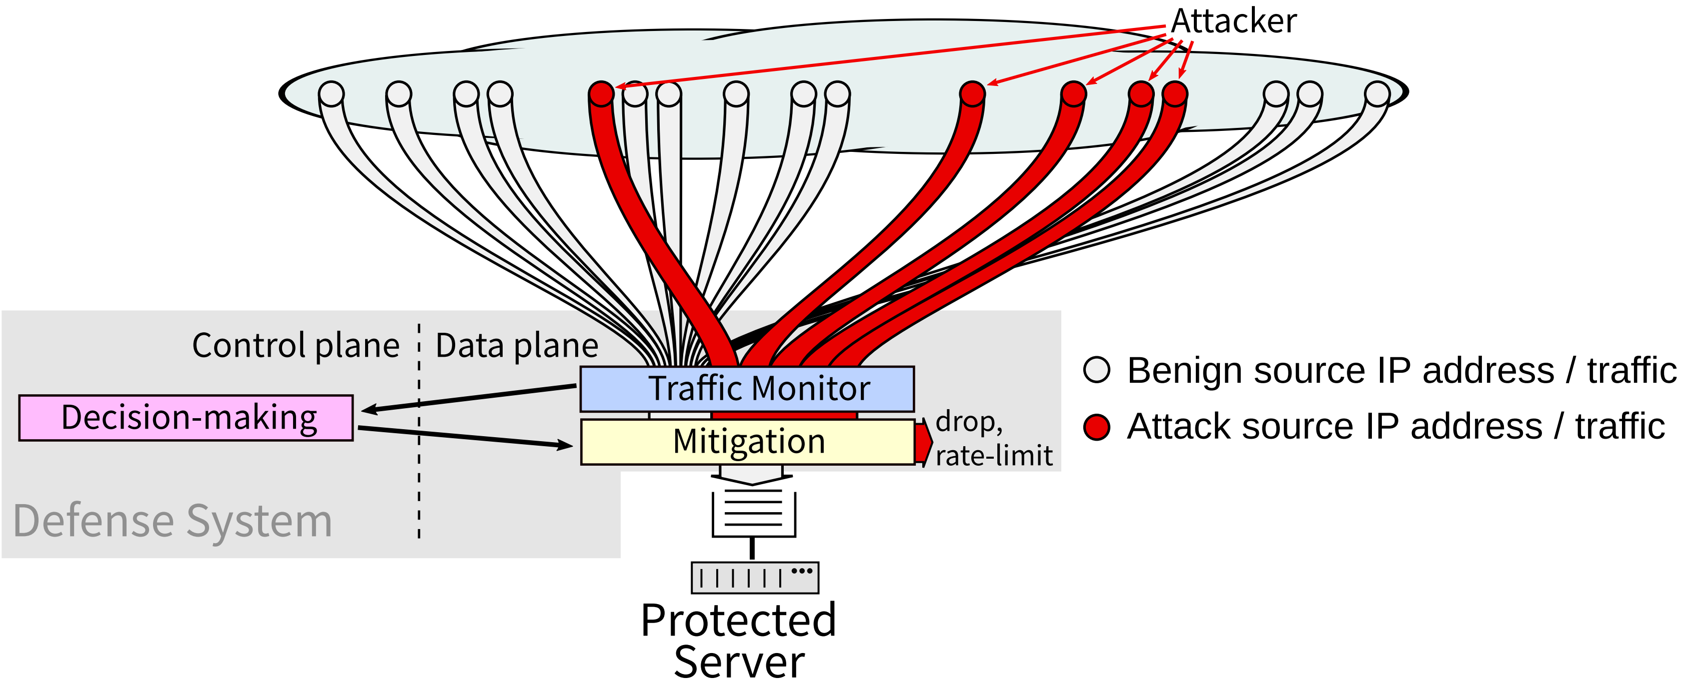

(Prior) Sketch-Based Detection has Random FPs

For example, Euclid and Jaqen get 4% - 50% false-positive rate in a realistic attack scenario with 50k attack source addresses.

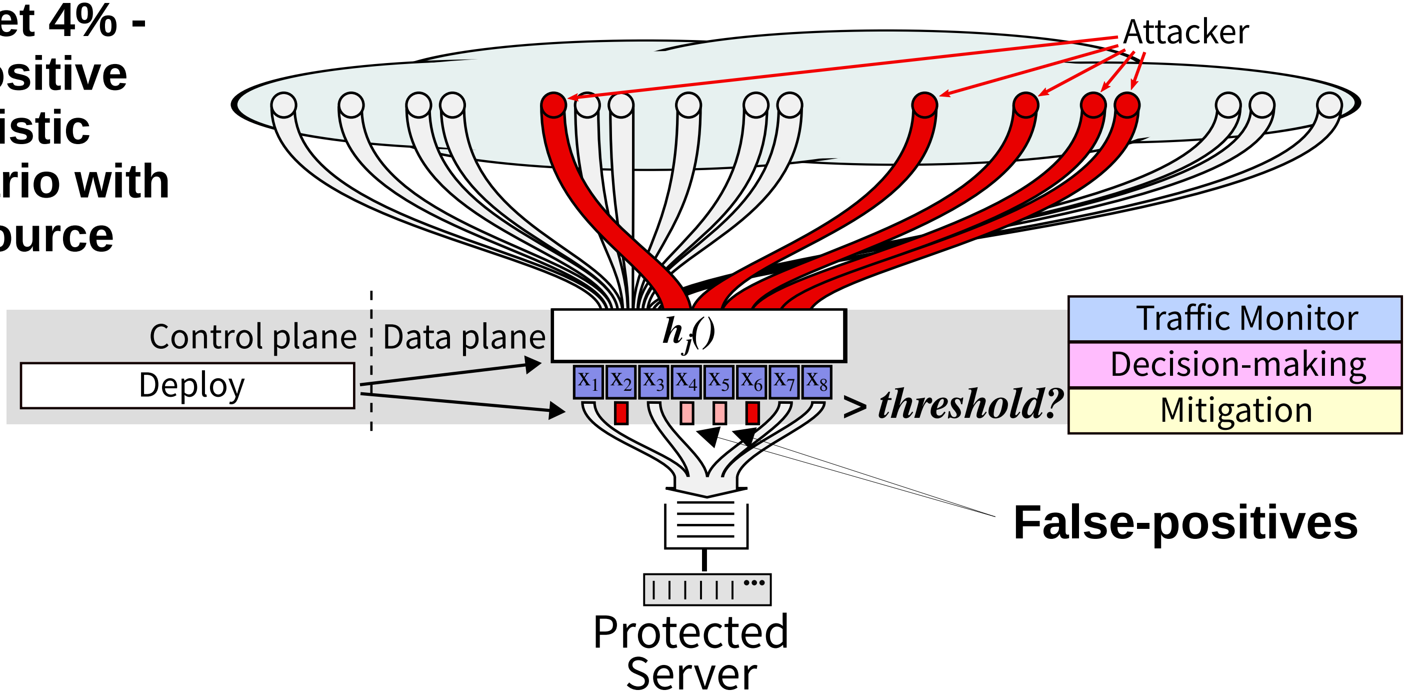

(Prior) ML-Based Detection has High Overheads

For example, LUCID would require ~80MB switch hardware memory.

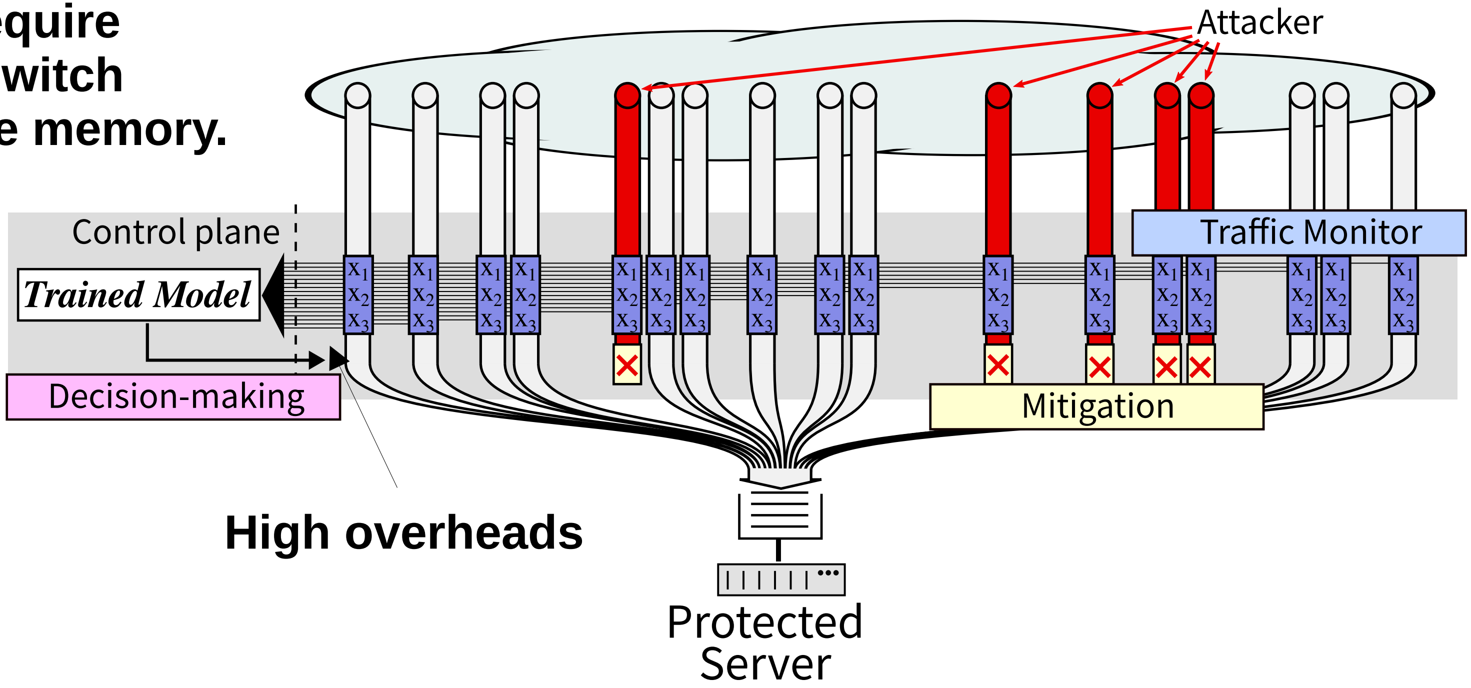

(Prior) Prefix-Level Detection has High FNs

For example, DREAM only allows counter-based detection.

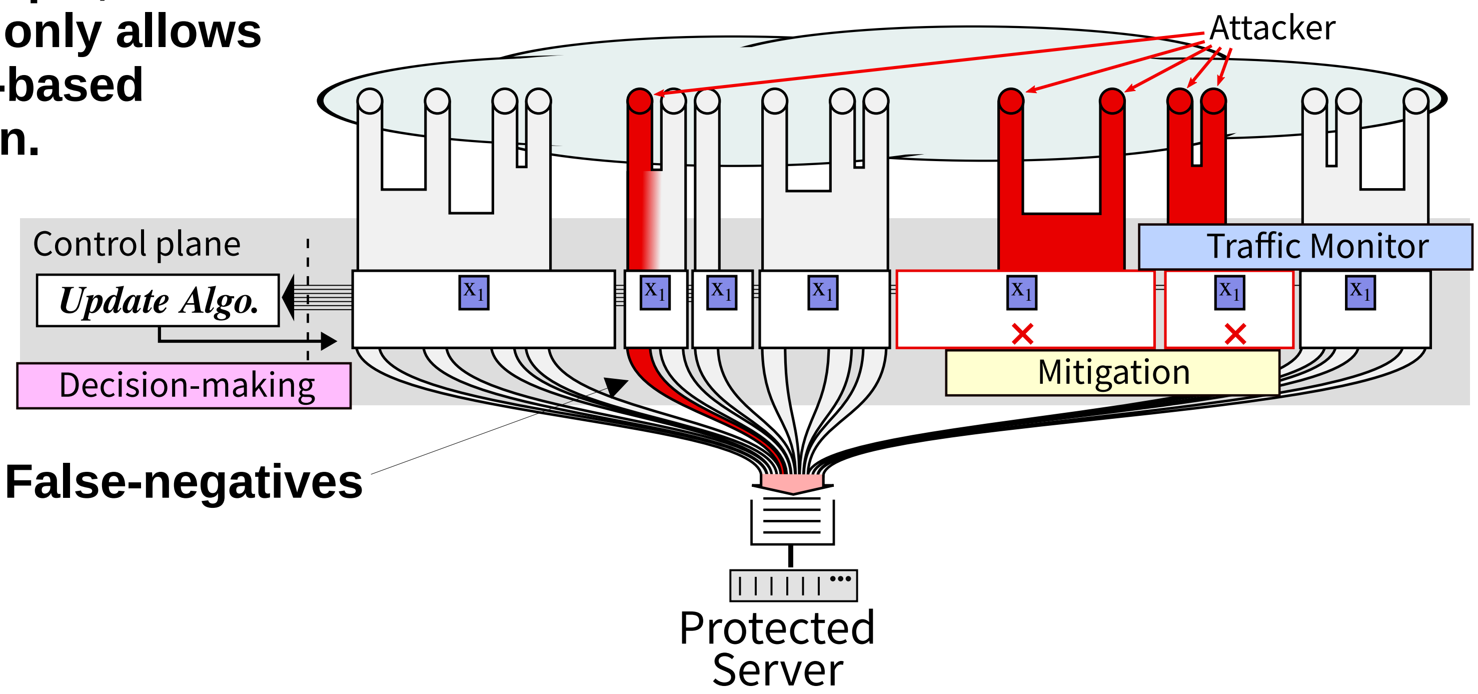

Prior Efforts Fail to Meet All Requirements

		Accurate	Scalable	Robust
Sketch	Jaqen, Euclid		**X**	
Flow-level ML	LUCID	**X**		
Prefix refinement	RADAR		**X**	

ZAPDOS: Prefix-Level ML + Iterative Refinement

Accurate →
Scalable →
Robust →

ZAPDOS: Prefix-Level ML + Iterative Refinement

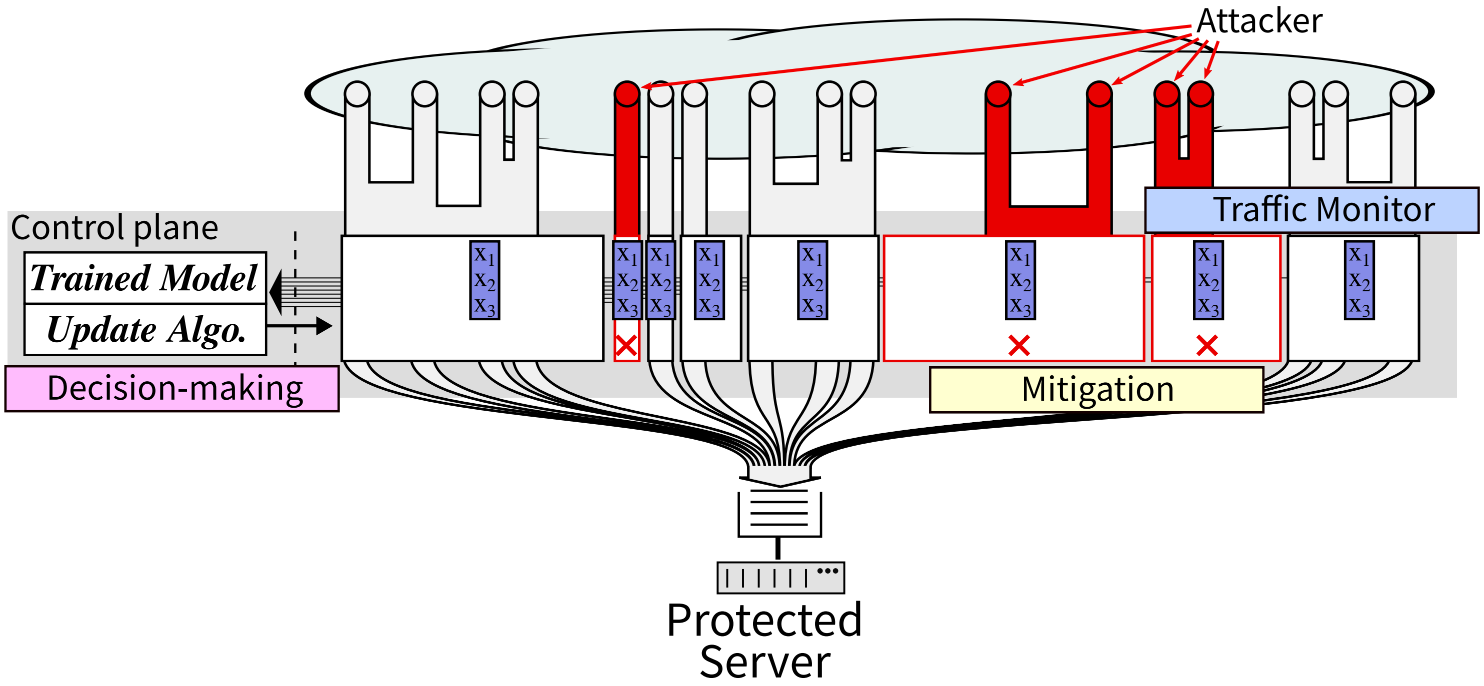

ZAPDOS: Prefix-Level ML

Accurate → prefix-level ML
Scalable →
Robust →

Prefix-level model can better “see through” ambiguous prefixes.

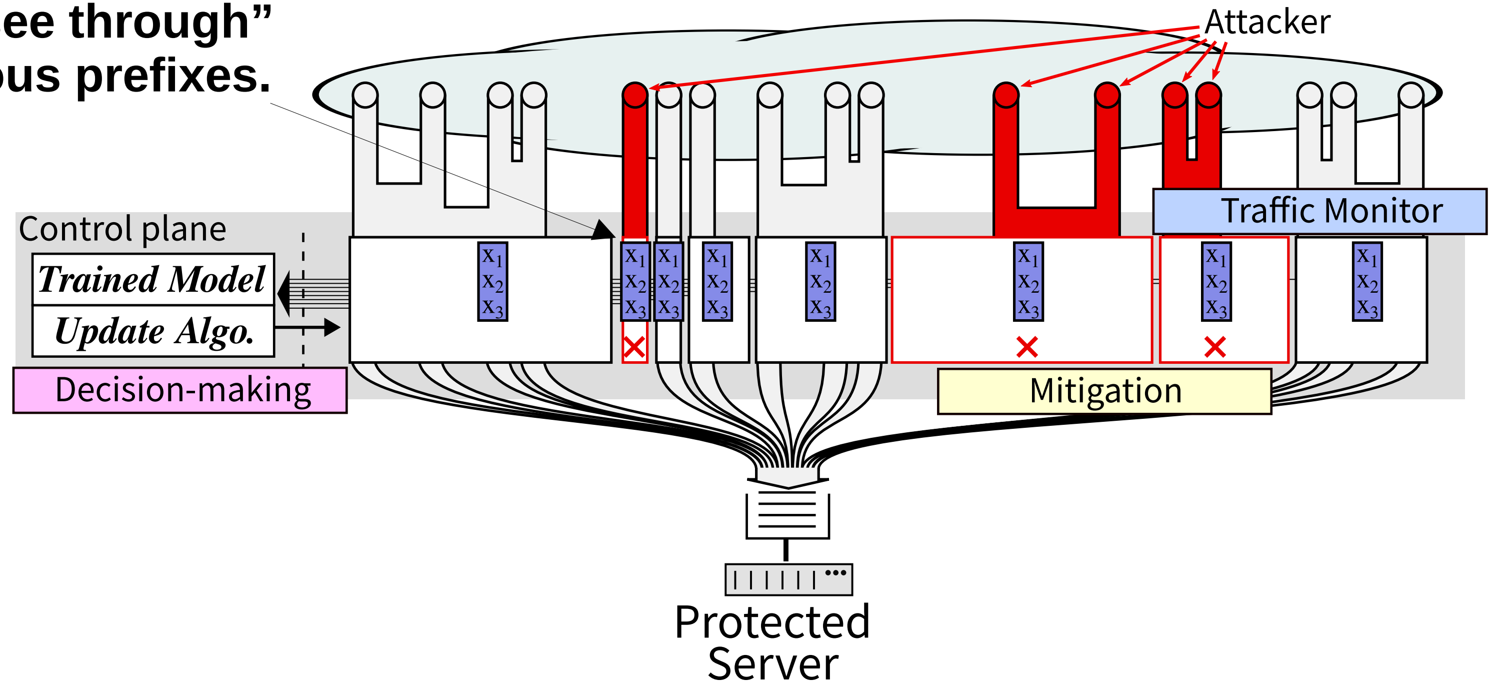

ZAPDOS: Iterative Refinement

Improved refinement
- Look-Ahead
- Look-Back

Accurate → prefix-level ML
Scalable →
Robust →

ZAPDOS: Refine with Look-Ahead

Accurate → prefix-level ML
Scalable → look-ahead
Robust →

Real addresses are clustered and sparse at long prefix lengths.

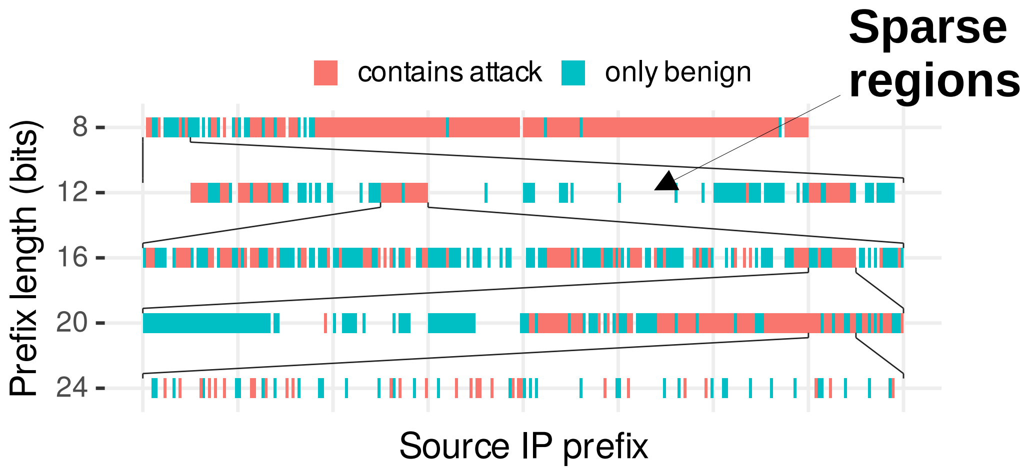

Structure of 50k mirai sources (attack) and ~5M MAWILab sources.

With look-ahead only non-empty children are monitored.

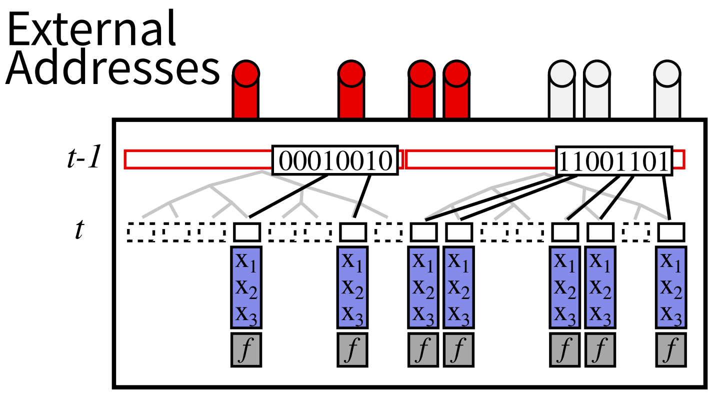

ZAPDOS: Refine with Look-Ahead

Accurate → prefix-level ML
Scalable → look-ahead
Robust →

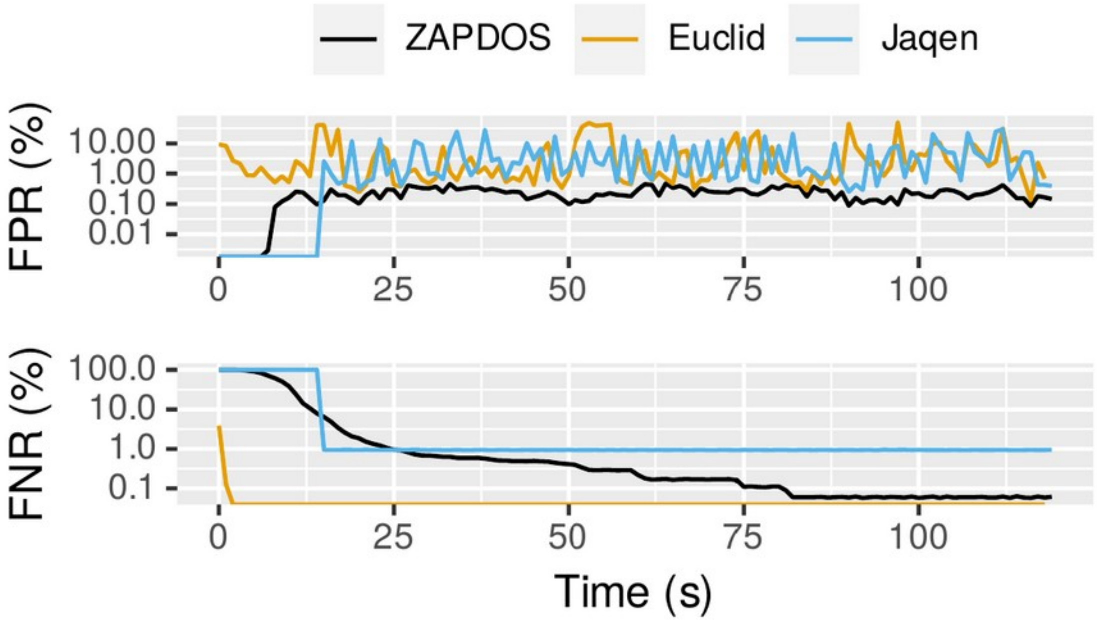

ZAPDOS achieves fast, accurate refinement using limited hardware resources.

ZAPDOS: Look-Back to Catch Changes

Accurate → prefix-level ML
Scalable → look-ahead
Robust → look-back

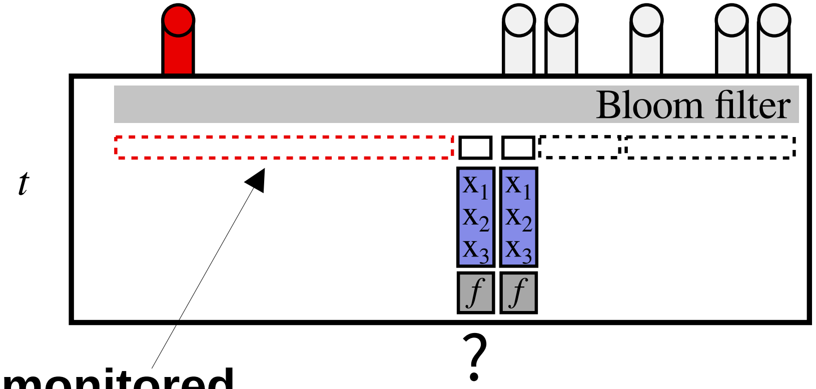

Un-monitored Prefix

cmisa@cs.uoregon.edu

ZAPDOS: Look-Back to Catch Changes

Accurate → prefix-level ML
Scalable → look-ahead
Robust → look-back

Look-back quickly re-focuses on the changed attack source address.

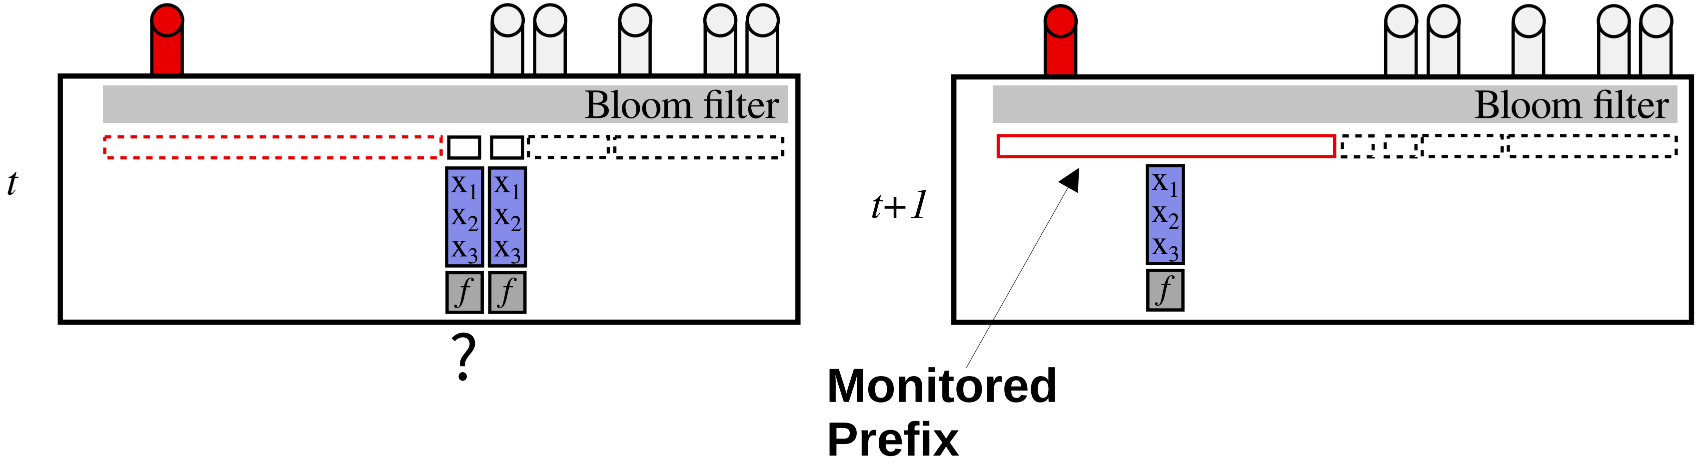

ZAPDOS: Look-Back to Catch Changes

Accurate → prefix-level ML
Scalable → look-ahead
Robust → look-back

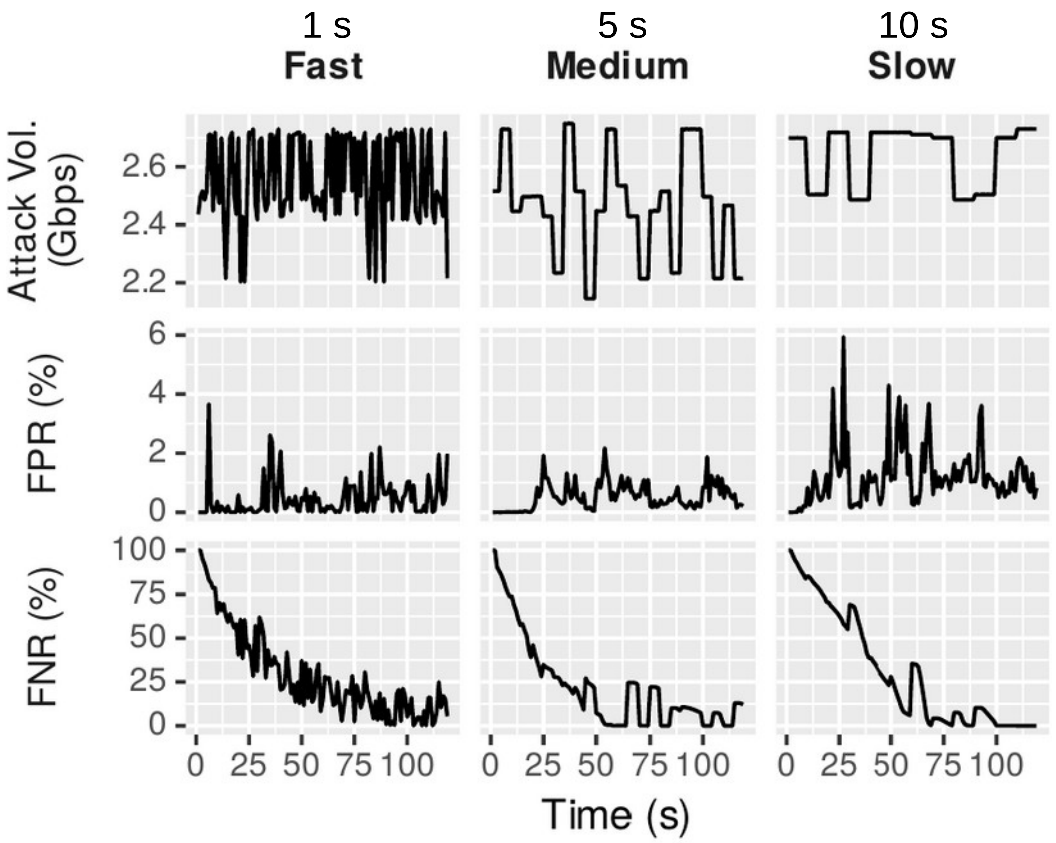

Look-back enables quickly re-focusing refinement when the attack changes.

ZAPDOS: See More in Our Paper

Data-fusion for training on **>100 attack scenarios** with realistic prefix-level distributions.

Accurate → prefix-level ML
Scalable → look-ahead
Robust → look-back

Implemented a **Tofino prototype**.

- **Batch-round-robin** updates.
- **Packet ferries** for feature collection.

Extended evaluation.

- <1% FPR and FNR across attack vectors.
- Works for **spoofed attack sources.**
- Works when the **upstream border link is flooded.**

Chris Misa
cmisa@cs.uoregon.edu

See our project page:
onrg.gitlab.io/projects/zapdos/

Thanks!