

LANTERN: Layered Adaptive Network Telemetry Collection for Programmable Dataplanes

[EuroP4' 23]

Kaiyu Hou
Alibaba Cloud

Dhiraj Saharia
Georgetown
University

Vinod Yegneswaran
SRI International

Phillip Porras
SRI International




GEORGETOWN UNIVERSITY


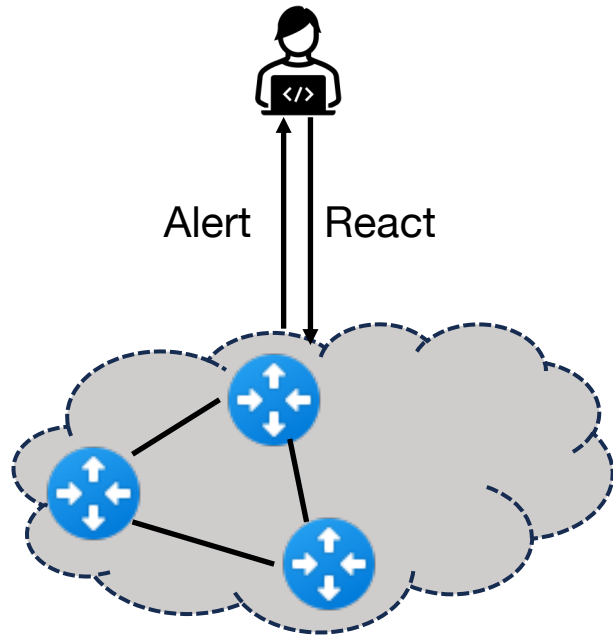


Motivation

DDoS is a well studied area in many different forms



P4 as a Solution?



Enabling the dataplane to dynamically adjust telemetry collection based on emergent network traffic patterns.

- High line rate processing (e.g. – Tbps)
- Programmability using DSLs such as P4

Background – State-of-the-Art



Method	Approach	Controller	Adaptive Telemetry	Data Plane Alert	Mitigating Triggers
Stats101	Online Computation	Yes	By Controller	Yes	No
Sonata	Reactive Query	Yes	Reactive by Admin	No	By Admin
DynATOS	Reactive Query	Yes	Reactive by Admin	No	By Admin
Poseidon	Predefined DP Policy	Yes	Policy-based	Yes	DDoS only
Jaqen	Sketch-based	Yes	No	Yes	DDoS only

Programmable Data Planes for Threat Mitigation

Background – State-of-the-Art



Method	Approach	Controller	Adaptive Telemetry	Data Plane Alert	Mitigating Triggers
Stats101	Online Computation	Yes	By Controller	Yes	No
Sonata	Reactive Query	Yes	Reactive by Admin	No	By Admin
DynATOS	Reactive Query	Yes	Reactive by Admin	No	By Admin
Poseidon	Predefined DP Policy	Yes	Policy-based	Yes	DDoS only
Jaqen	Sketch-based	Yes	No	Yes	DDoS only

Programmable Data Planes for Threat Mitigation

Background – State-of-the-Art



Method	Approach	Controller	Adaptive Telemetry	Data Plane Alert	Mitigating Triggers
Stats101	Online Computation	Yes	By Controller	Yes	No
Sonata	Reactive Query	Yes	Reactive by Admin	No	By Admin
DynATOS	Reactive Query	Yes	Reactive by Admin	No	By Admin
Poseidon	Predefined DP Policy	Yes	Policy-based	Yes	DDoS only
Jaqen	Sketch-based	Yes	No	Yes	DDoS only

Programmable Data Planes for Threat Mitigation

Background – State-of-the-Art




Method	Approach	Controller	Adaptive Telemetry	Data Plane Alert	Mitigating Triggers
Stats101	Online Computation	Yes	By Controller	Yes	No
Sonata	Reactive Query	Yes	Reactive by Admin	No	By Admin
DynATOS	Reactive Query	Yes	Reactive by Admin	No	By Admin
Poseidon	Predefined DP Policy	Yes	Policy-based	Yes	DDoS only
Jaqen	Sketch-based	Yes	No	Yes	DDoS only
LANTERN	Layer-based ML	Yes	Layer-based	Yes	Diverse attacks

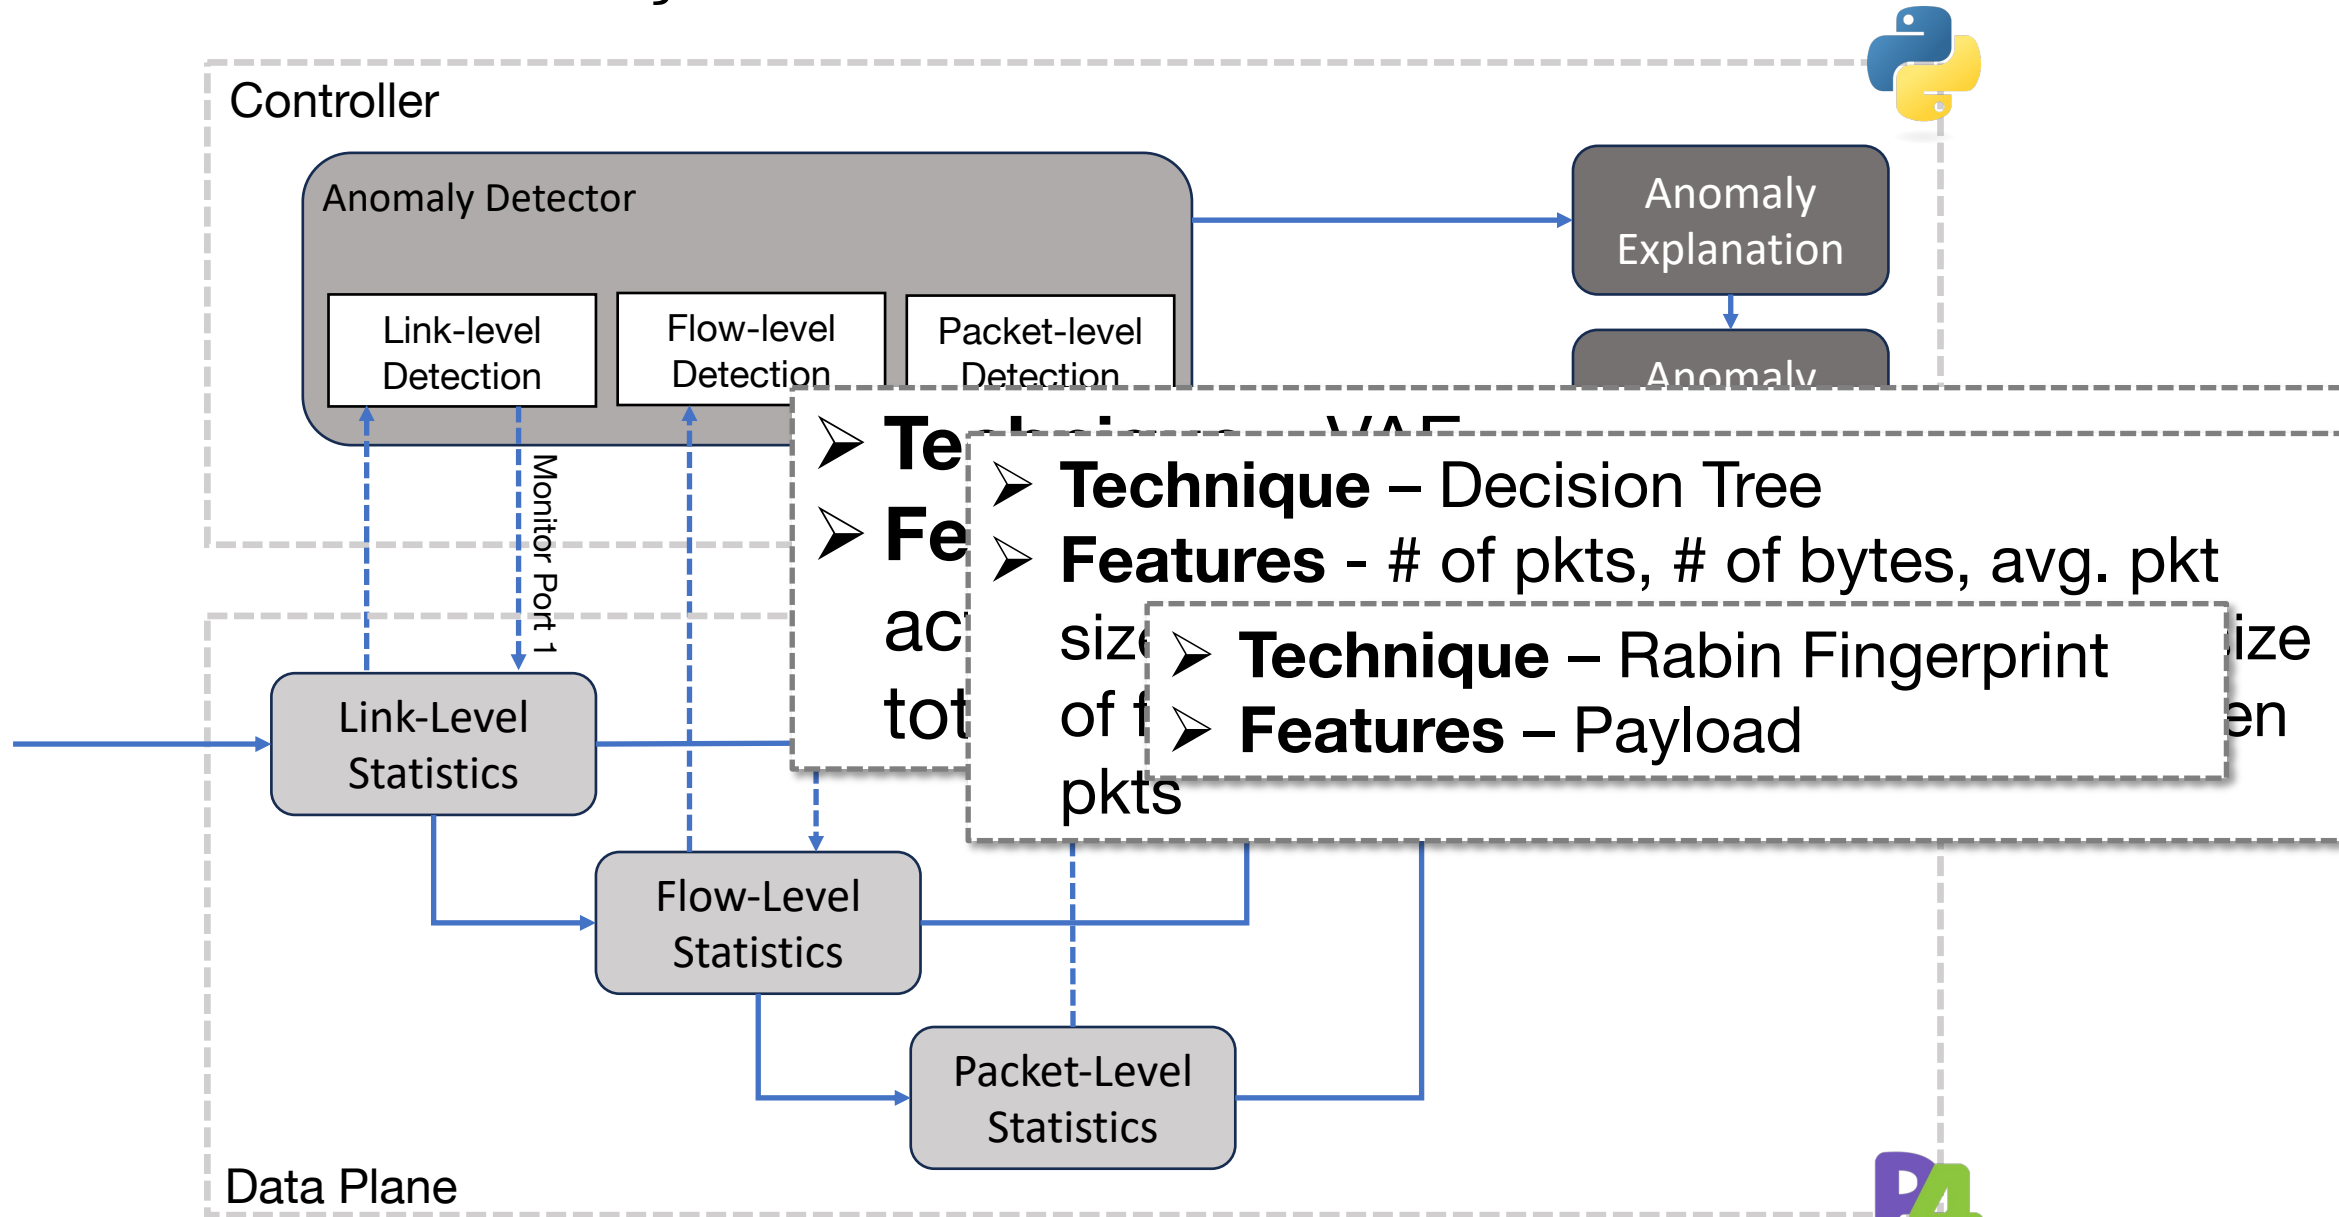
Programmable Data Planes for Threat Mitigation

Vision

Scenario - Internal → External IoT DDoS Attack



LANTERN*: System Overview




* Layered Adaptive Network TElemetry Collection for Programmable Dataplanes

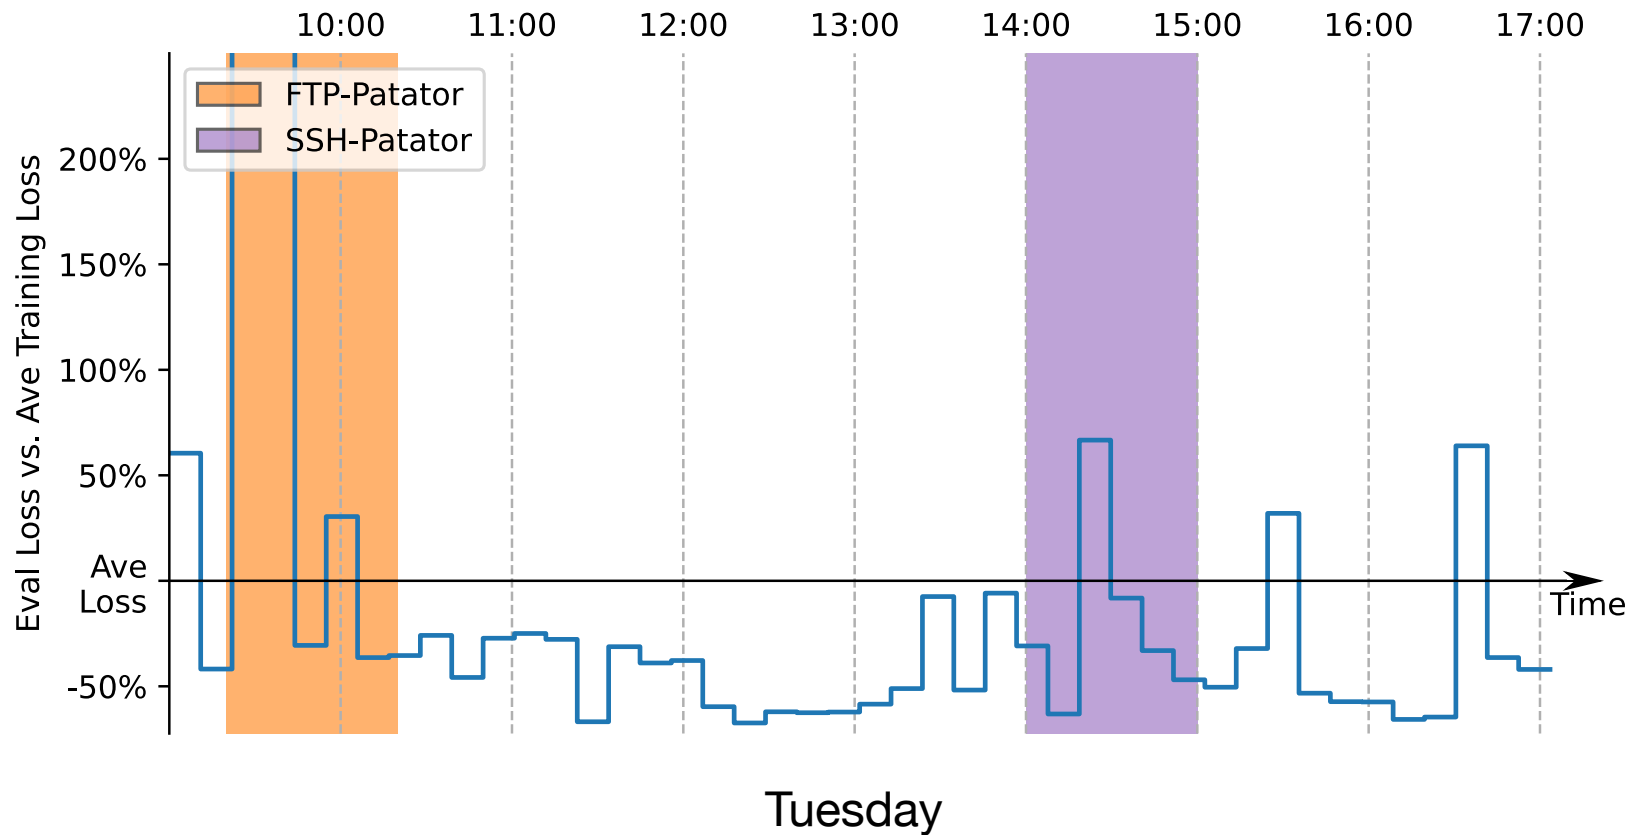


Evaluation Setup

- Dataset – CIC-IDS – [[LINK](#)]
 - Model trained on Monday benign traffic
 - Tested on Tuesday, Wednesday, and Friday traffic
- Testbed
 - 1 switch (bmv2, Tofino)
 - 2 end hosts (mininet)
 - Python based controller




Evaluation: Link-Level Anomaly Detection



Key Takeaways

- VAEs can detect vast majority of attacks with limited telemetry
- False positives can be corrected at the next layer

Evaluation: Flow-Level Attack Classification



Key Takeaway

- Decision Tree can perform attack classification using flow features with high accuracy
- Unbalanced dataset

Evaluation: Packet Level Mitigation

- Sample packets from attack flows
- Background flows → Benign
- Eliminate from candidate signatures

Algorithm 1: Fingerprint-based attack mitigation

Input : *attack_flows*, *benign_flows*
Output: *mitigation_rules*

- 1 **for** flow *in* *attack_flows* **do**
- 2 $cand_sigs \leftarrow cand_sigs \cup \{\text{Rabin Fingerprint}(flow[1:4])\}$
- 3 **for** flow *in* *background_flows* **do**
- 4 $cand_sigs \leftarrow cand_sigs - \{\text{Rabin Fingerprint}(flow[1:4])\}$
- 5 **for** sig *in* *cand_sigs* **do**
- 6 $rules \leftarrow rules \cup \{\text{Match:raw}(sig), \text{Action:reject}(src_ip)\}$

```
'\nAccept-Encoding: gzip, deflate\r\nAccept-  
'on: keep-alive\r\nAccept-Encoding: gzip,'  
'08.1 HTTP/1.1\r\nHost: 205.174.165.73:808'
```

Exemplar Benign Signature

```
't: */*\r\nUser-Agent: python-requests/2.1'  
' : */*\r\nUser-Agent: python-requests/2.14'  
'*/*\r\nUser-Agent: python-requests/2.14.2'  
...  
'api/pop?botid=mitacs-pc6&sysinfo=Window'  
'api/pop?botid=mitacs-pc4&sysinfo=Window'  
...  
'botid=mitacs-pc6&sysinfo=Windows%207 HT'  
'botid=mitacs-pc4&sysinfo=Windows%2010 H'  
..
```

Exemplar Attack Signature

Key Takeaway

- Our detection algorithm captured all 187 attack flows without false positives.

Evaluation: Resource / Performance

- Hardware Utilization

Resource	Usage
TCAM	0
SRAM	8.75%
VLIW	6.77%
Map RAM	13.89%

- Latency

- Compared against basic forwarding program
- ~7ns overhead on average across 2000 packets

Conclusion and Future Work

- ✓ Proposed a layered approach for dynamically adjustable tunable telemetry collection to make it easier to integrate different types of ML algorithms to defend against security threats
 - ✓ VAEs, DTs, Rabin Fingerprints
- ✓ Validated our layered approach using the CIC-IDS dataset and conducted a performance evaluation of the system.

Ongoing/Future Work

- Offloading ML algorithms to data plane + switch native implementation of mitigating triggers
- Scalability testing on diverse datasets

Questions?

Thank You

Dhiraj Saharia

Email: ds1849@georgetown.edu

Web: <https://dsaharia.com>

Lab: <https://seclab.cs.georgetown.edu/>