IoT Device Fingerprinting on Commodity Switches

Carson Kuzniar, Miguel Neves, Vladimir Gurevich, Israat Haque

DALHOUSIE UNIVERSITY

PINet
Programmable and Intelligent Networking

intel

NOMS

IoT Landscape

- 12.3 billion connected IoT devices
- Roughly $160 Billion in IoT enterprise spending

- IoT botnet-based DDoS attacks reaching over 1 Tbps
- Hundreds of thousands of compromised devices
  - Above 200K for Mirai botnet

IoT Fingerprinting

Critical task for network administrators to

- Check for known vulnerabilities
- Set access/firewall rules
- Configure intrusion detection systems

State of the Art

- Machine Learning (WiSec'20)
- Labeled Traffic from Manufacture/User (NOMS'20)
- Length and direction-based signatures (NDSS'20, IM'21)

Solutions face obstacles at scale or with network wide view
Require mirroring traffic to dedicated hardware

Device Signatures

- Use packet length and direction to create signatures for events (e.g., On/Off)
- Reliably fingerprint devices using these signatures

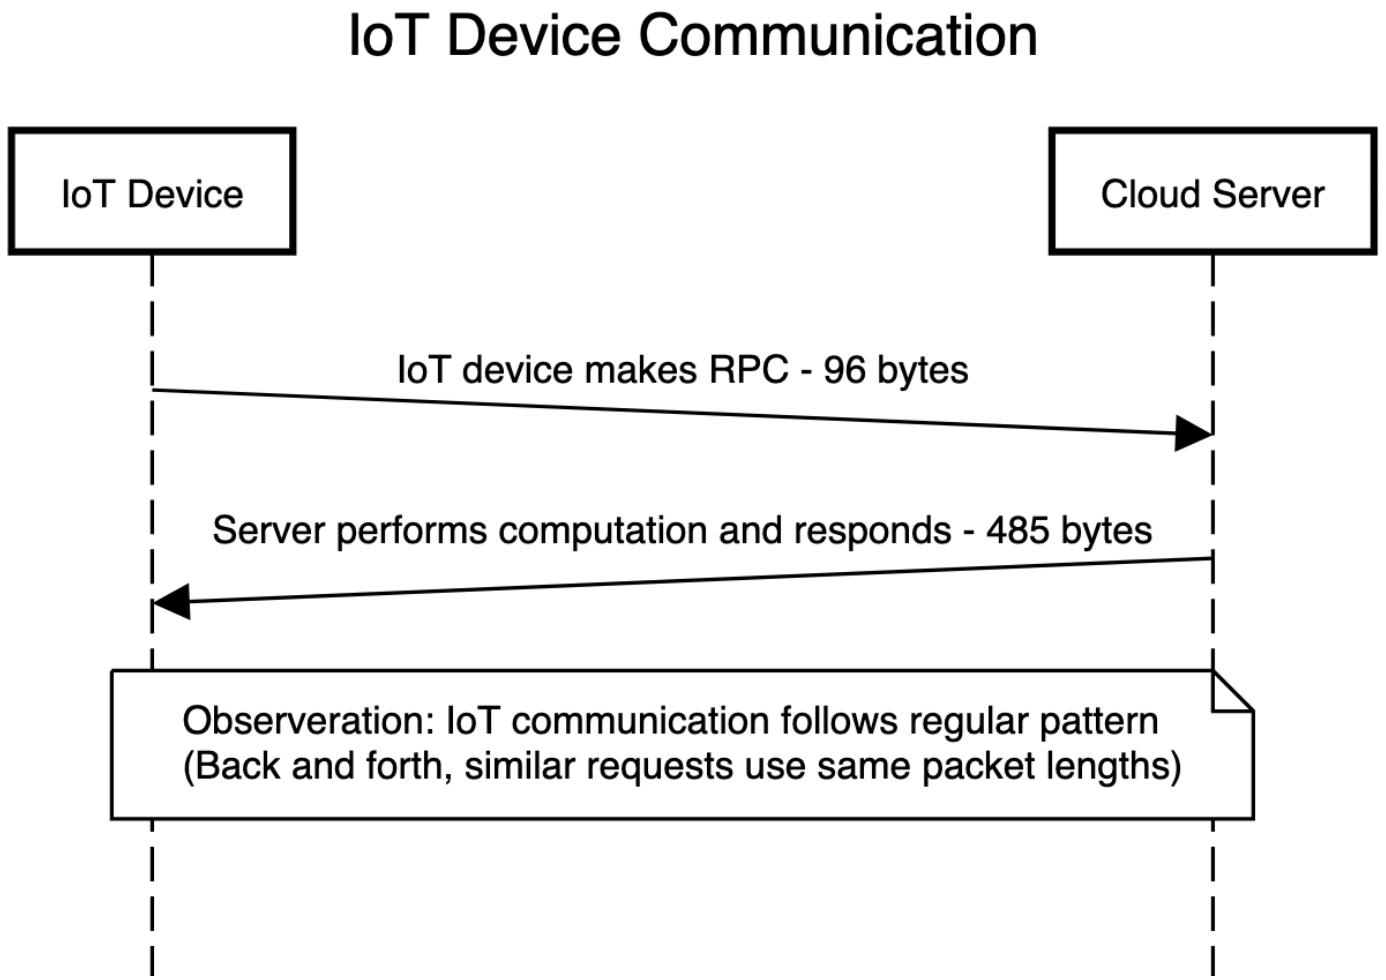

R. Trimananda, J. Varmarken, A. Markopoulou, and B. Demsky, "Packet-Level Signatures for Smart Home Devices," Proceedings of the 2020 Network and Distributed System Security (NDSS) Symposium, February 2020.

Challenges

Volume

- Large amounts of traffic on high-speed links
- Drop accuracy or add significant delay

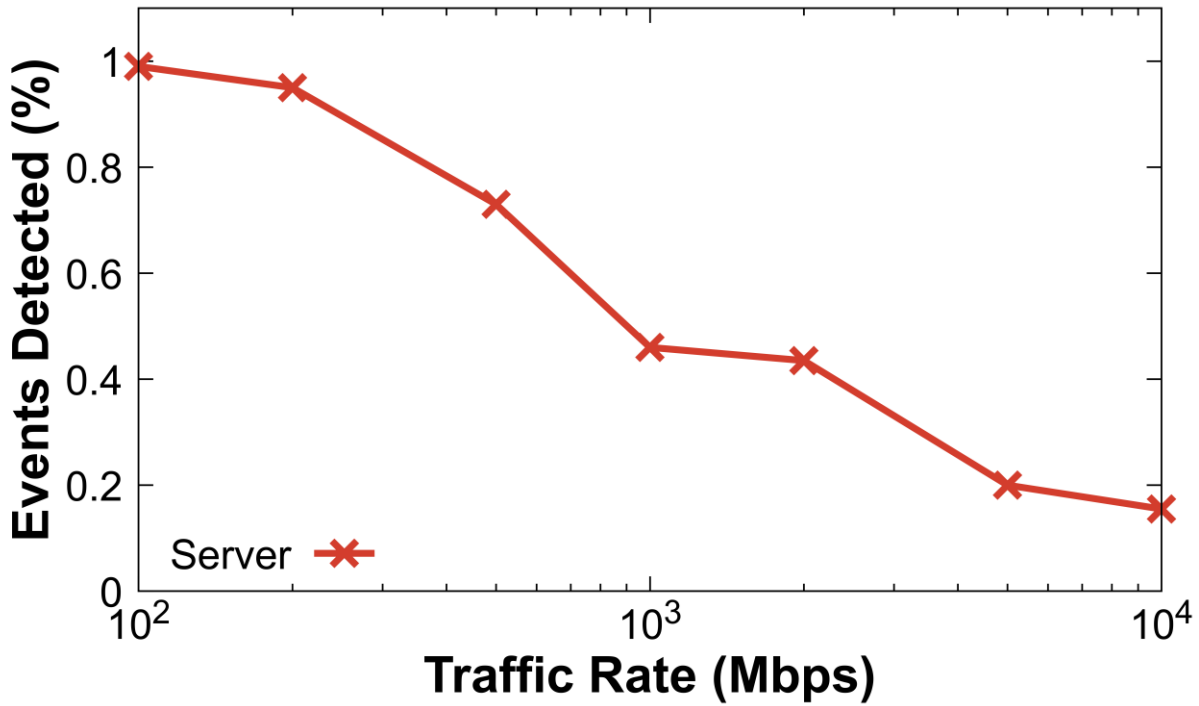

Challenges

Volume

- Large amounts of traffic on high-speed links
- Drop accuracy or add significant delay

Granularity

- Sampling and aggregation miss quiet device

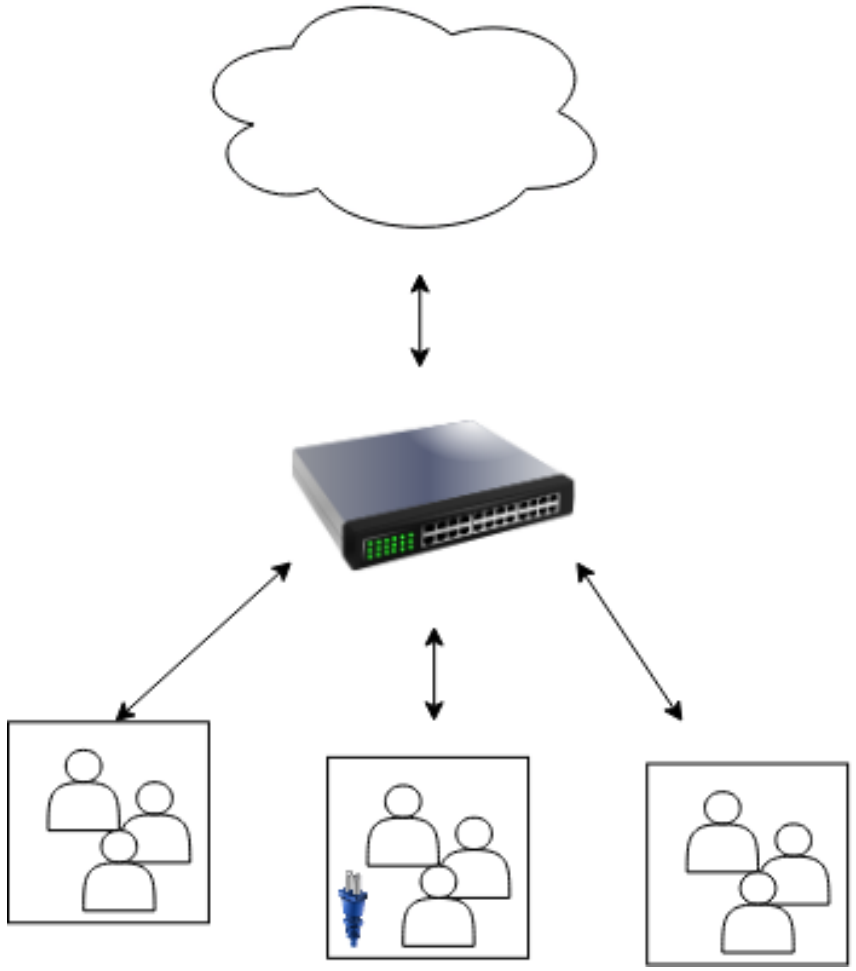

Introducing PoirIoT

Programmable Data-Plane based Fingerprinting:

- High-speed : Tofino ASIC Line Rate (Tbps)
- High granularity: Inspects every packet as part of its forwarding process
- Modular : Efficient use of switch resources

PoirIoT Architecture

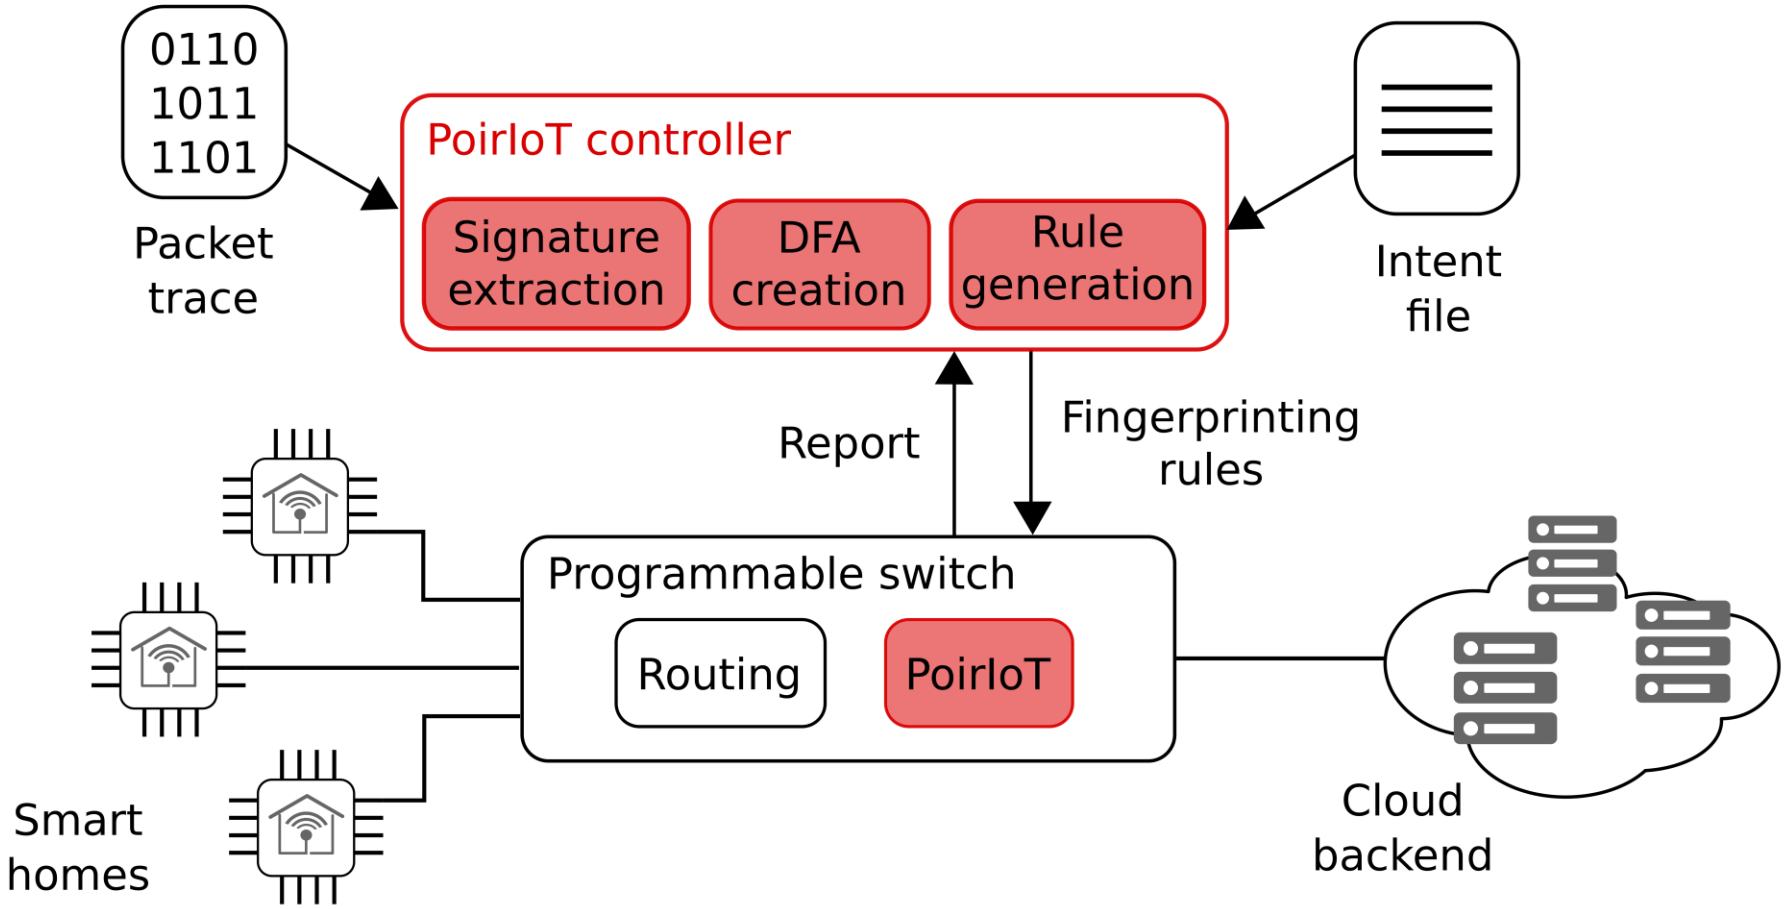

PoirIoT Architecture

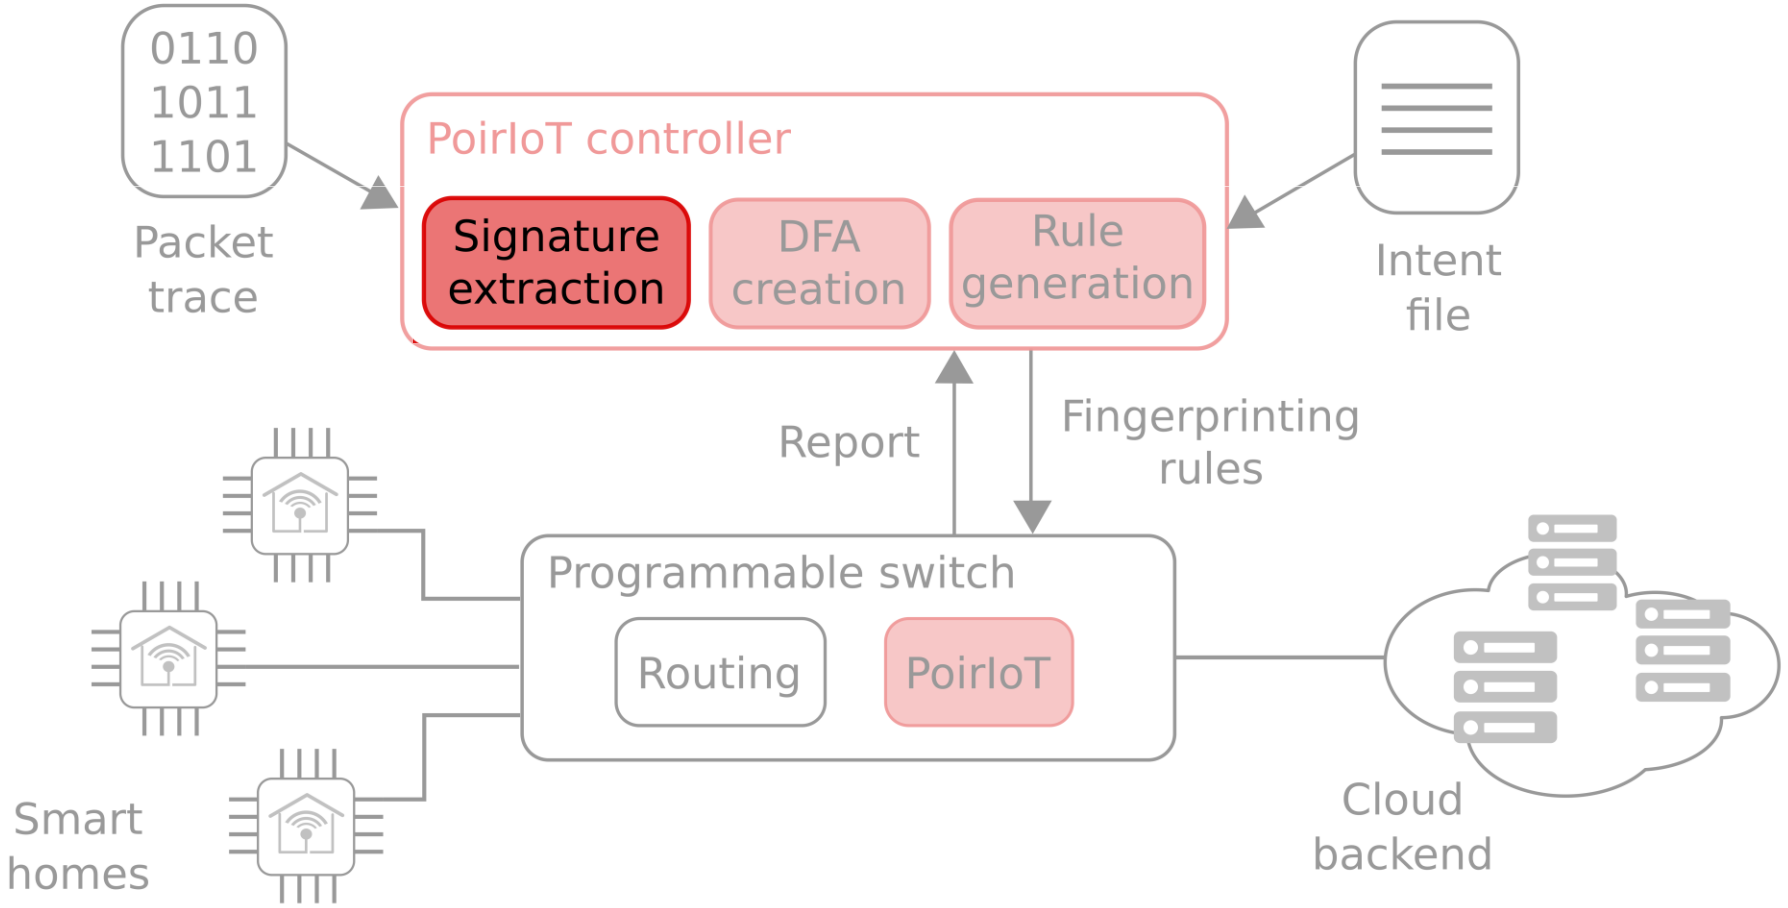

Signature Extraction

F1: ... C-211 S-1063 S-998 S-1276 ...

F2: ... C-211 S-1063 S-783 S-1277 ...

Signature Extraction

F1: ... C-211 S-1063 S-998 S-1276 ...

F2: ... C-211 S-1063 S-783 S-1277 ...

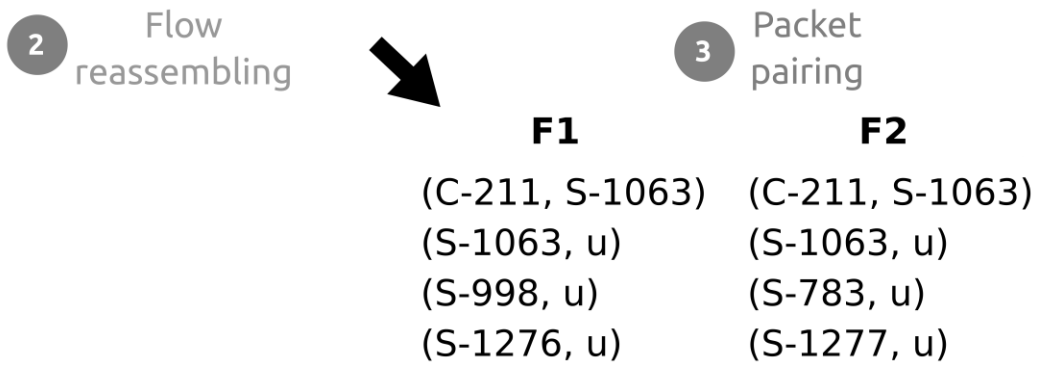

Signature Extraction

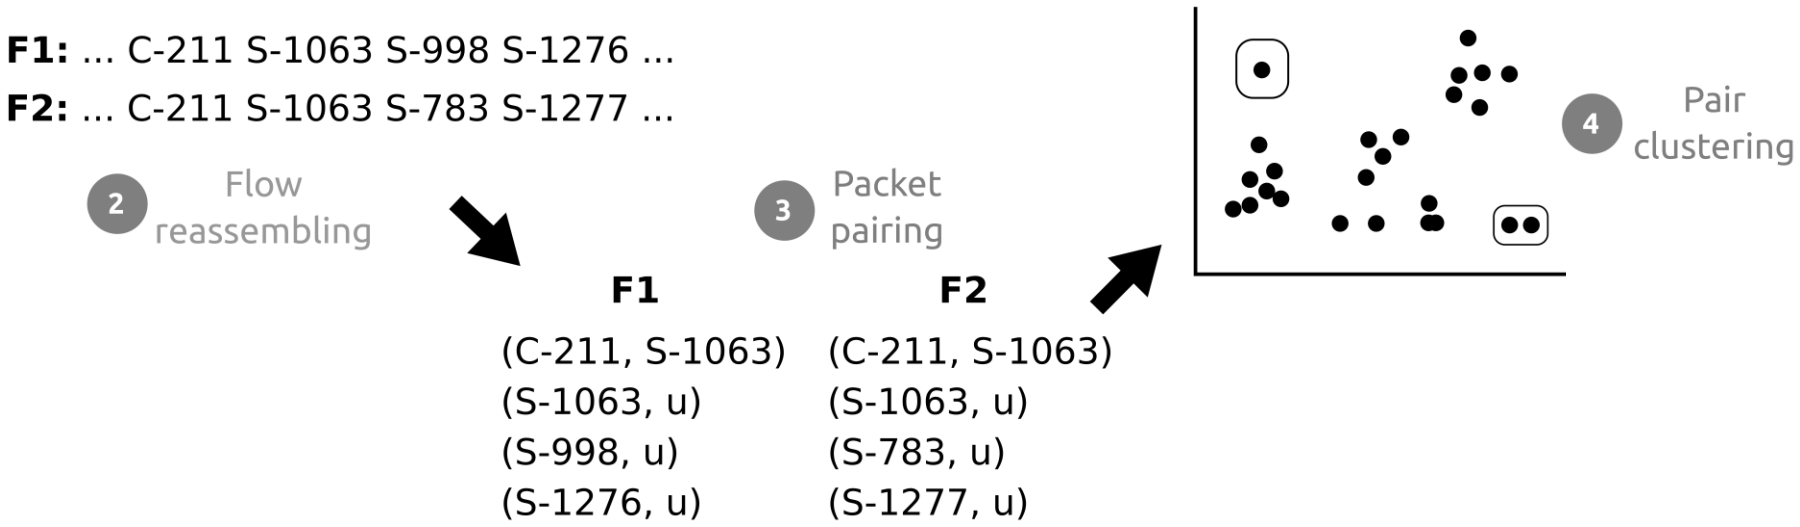

Signature Extraction

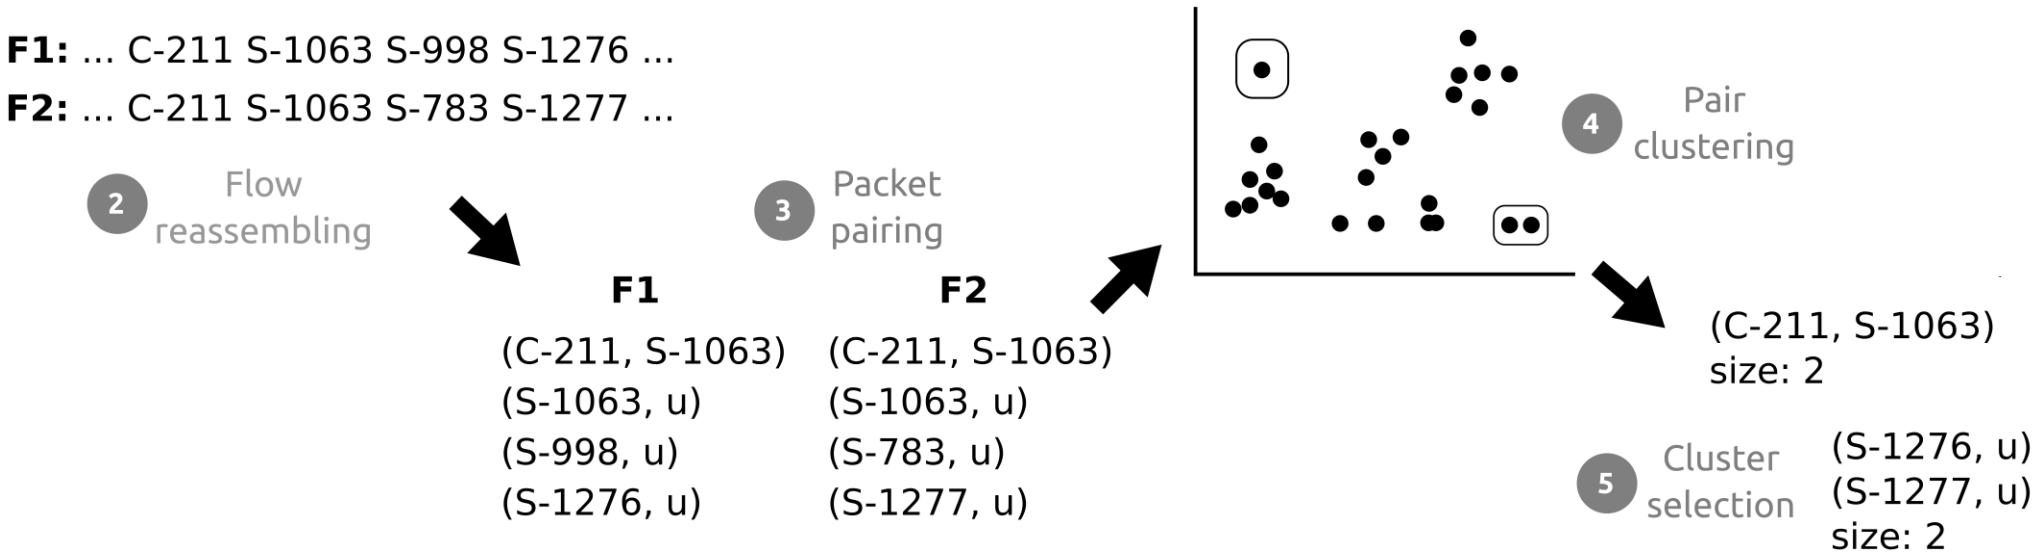

Signature Extraction

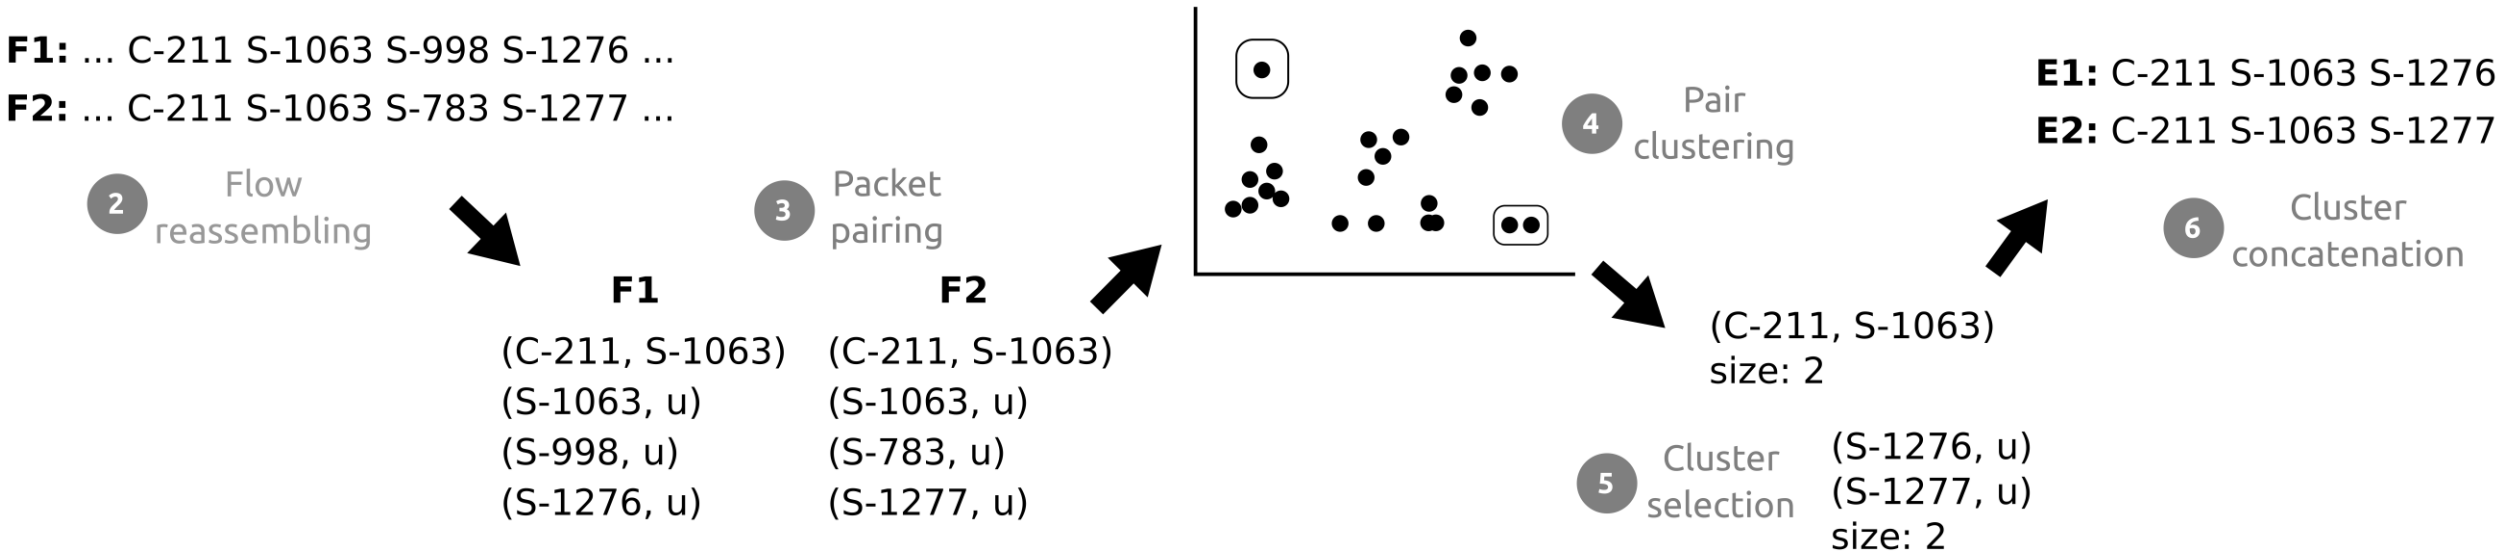

PoirIoT Architecture

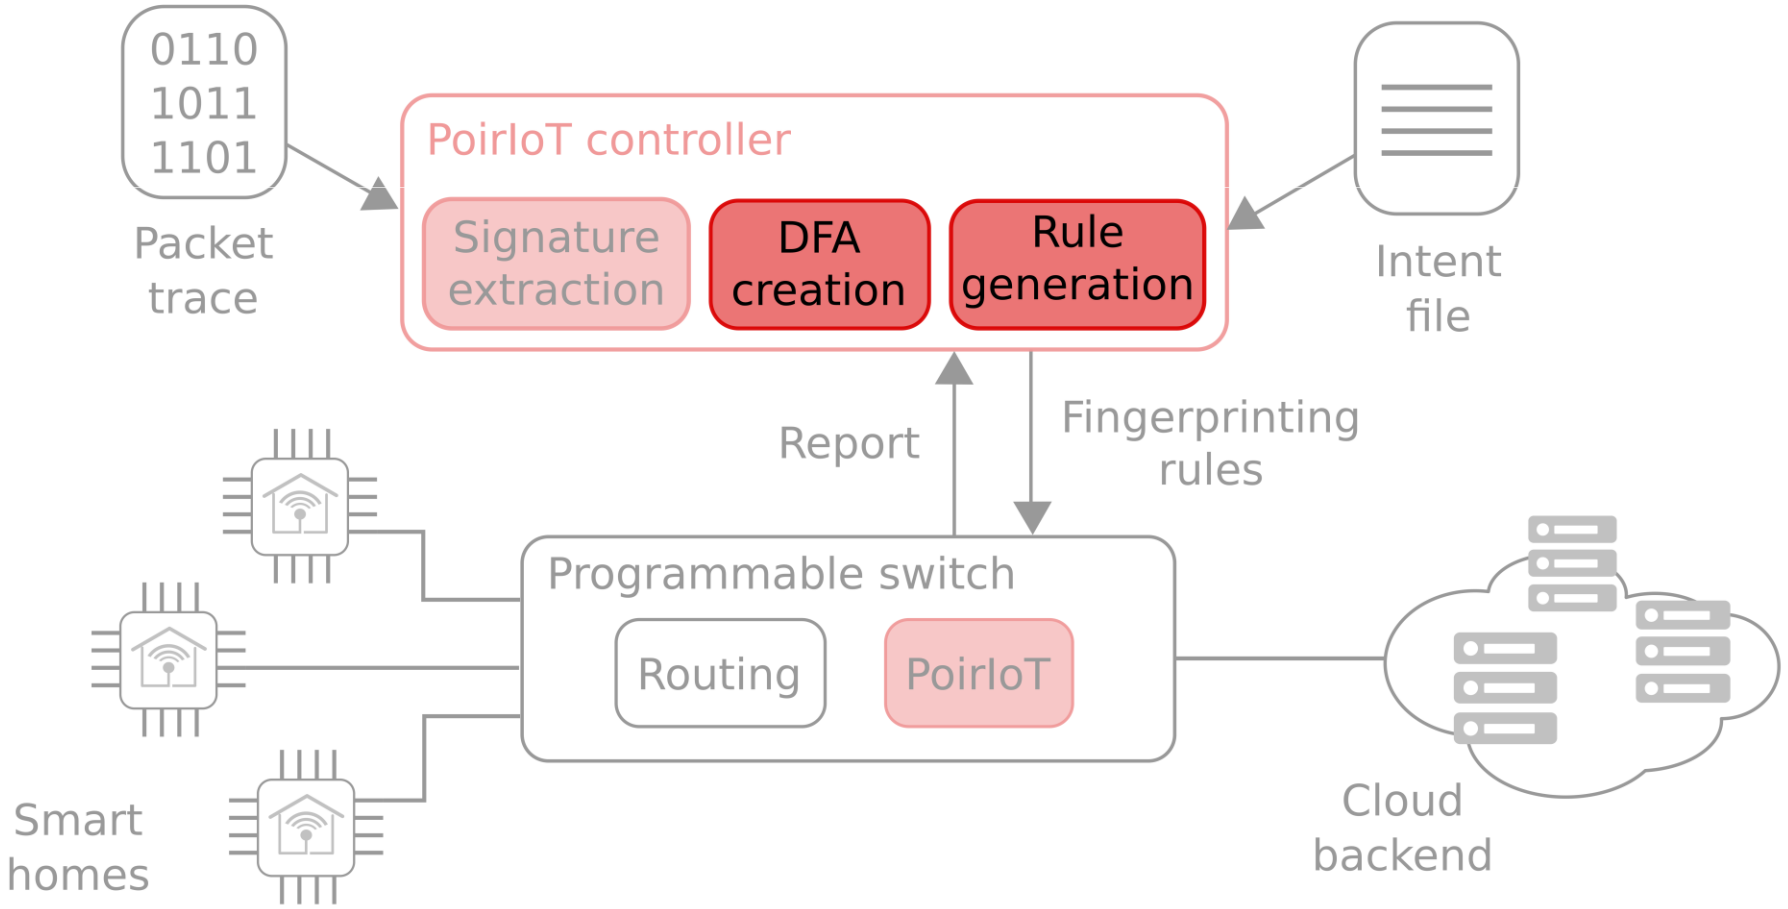

DFA and Rule Creation

E1:	**E2**:
C-211	C-211
S-1063	S-1063
S-1276	S-1277

DFA and Rule Creation

E1:	**E2**:
C-211	C-211
S-1063	S-1063
S-1276	S-1277

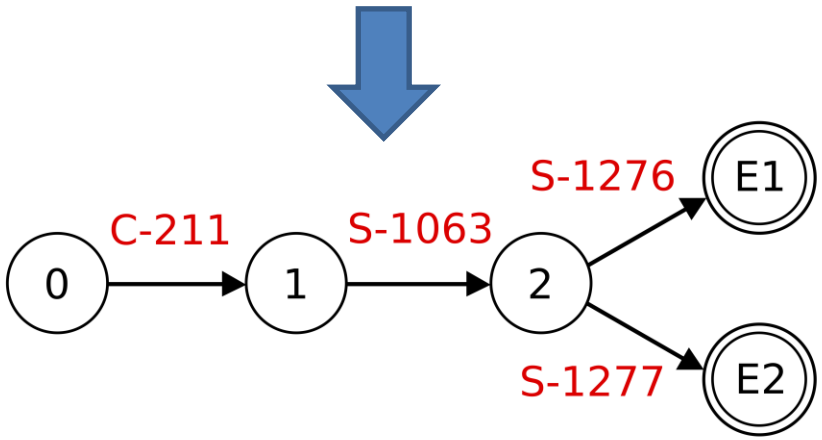

DFA and Rule Creation

E1:
C-211
S-1063
S-1276

E2:
C-211
S-1063
S-1277

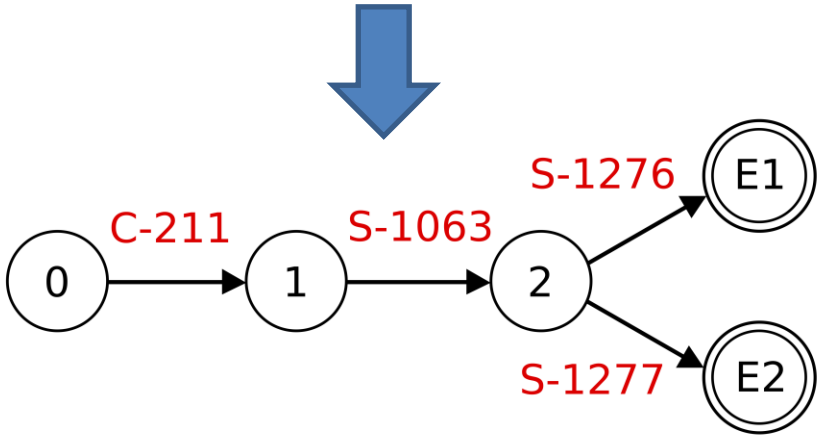

Match			Action
State	Dir	Length	
0	C	211	set_state(1)
1	S	1063	set_state(2)
2	S	1276	report_event(E1)
2	S	1277	report_event(E2)

PoirIoT Architecture

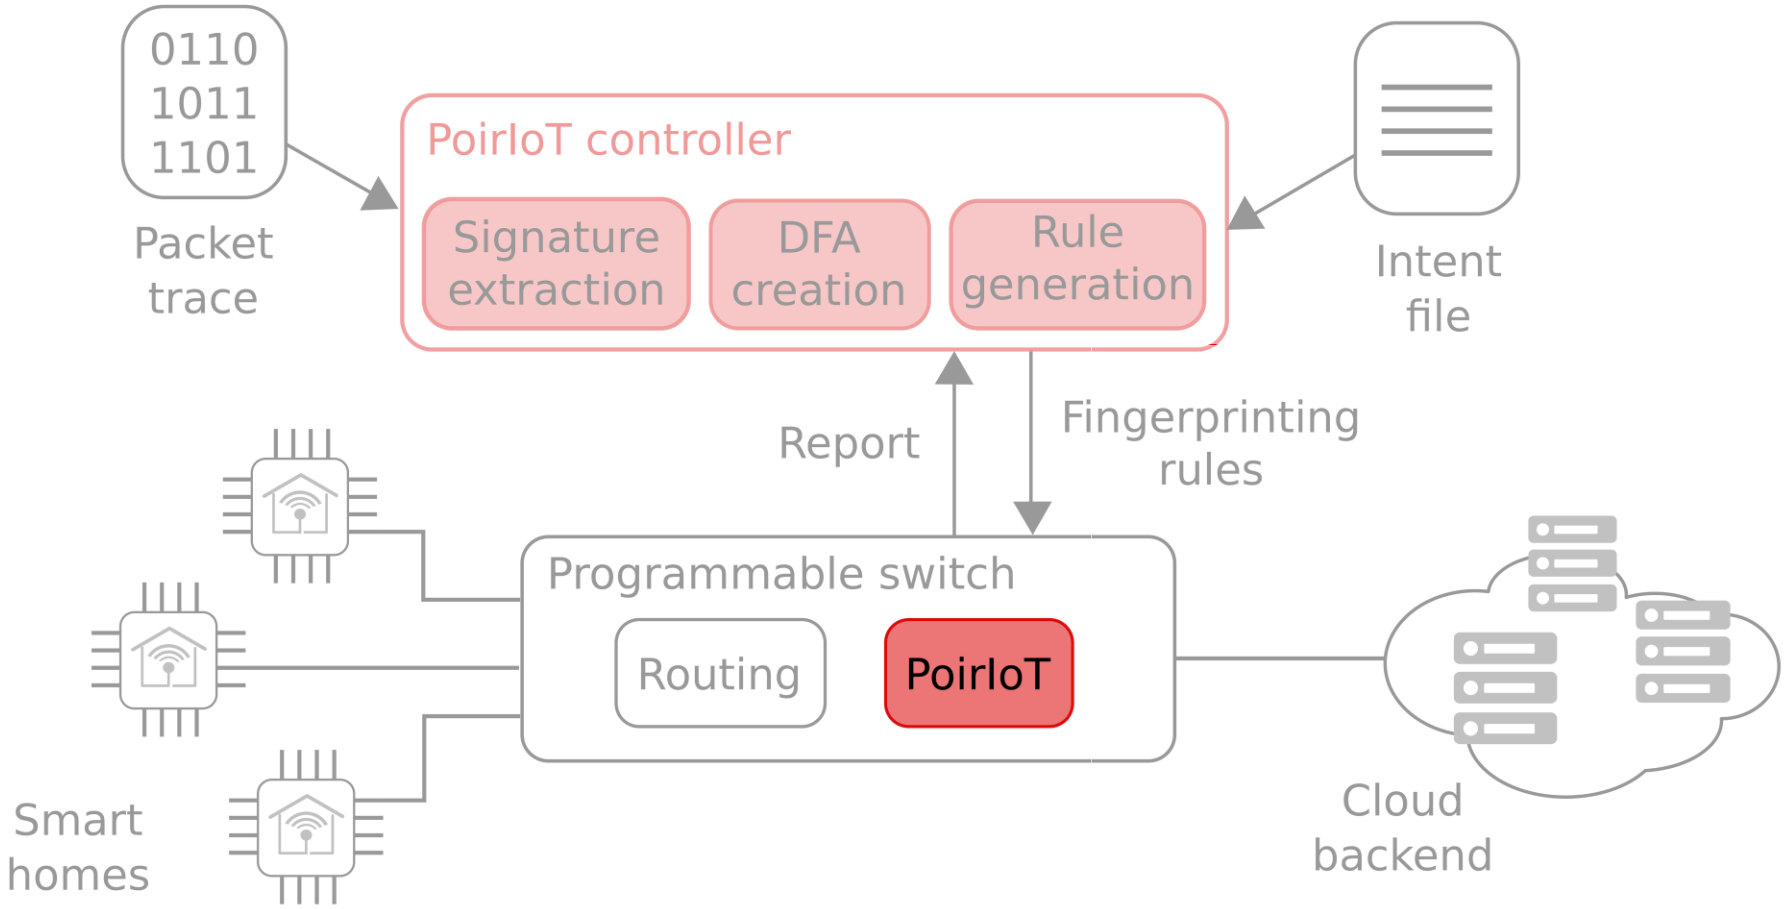

PoirIoT Data Plane

PoirIoT Data Plane

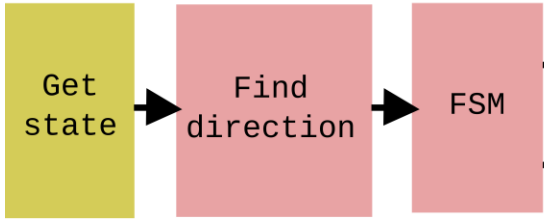

- State and Timer info stored per flow
- Virtual FSM per flow while requiring only a single table

PoirIoT Data Plane

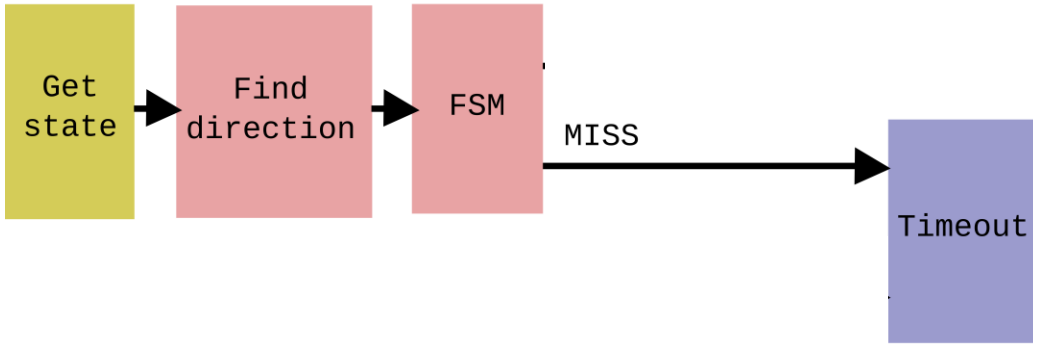

- State and Timer info stored per flow
- Virtual FSM per flow while requiring only a single table

PoirIoT Data Plane

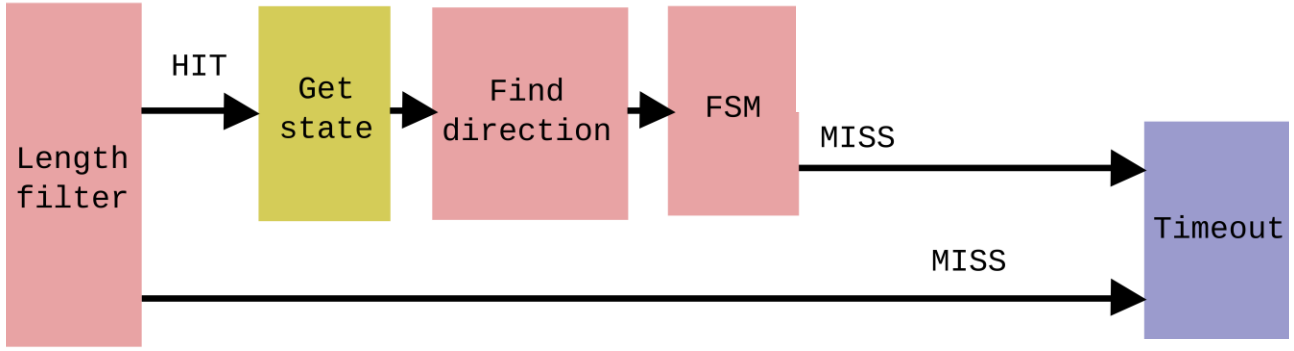

- Length filter prevents resource allocation to flows with no packets of interest

PoirIoT Data Plane

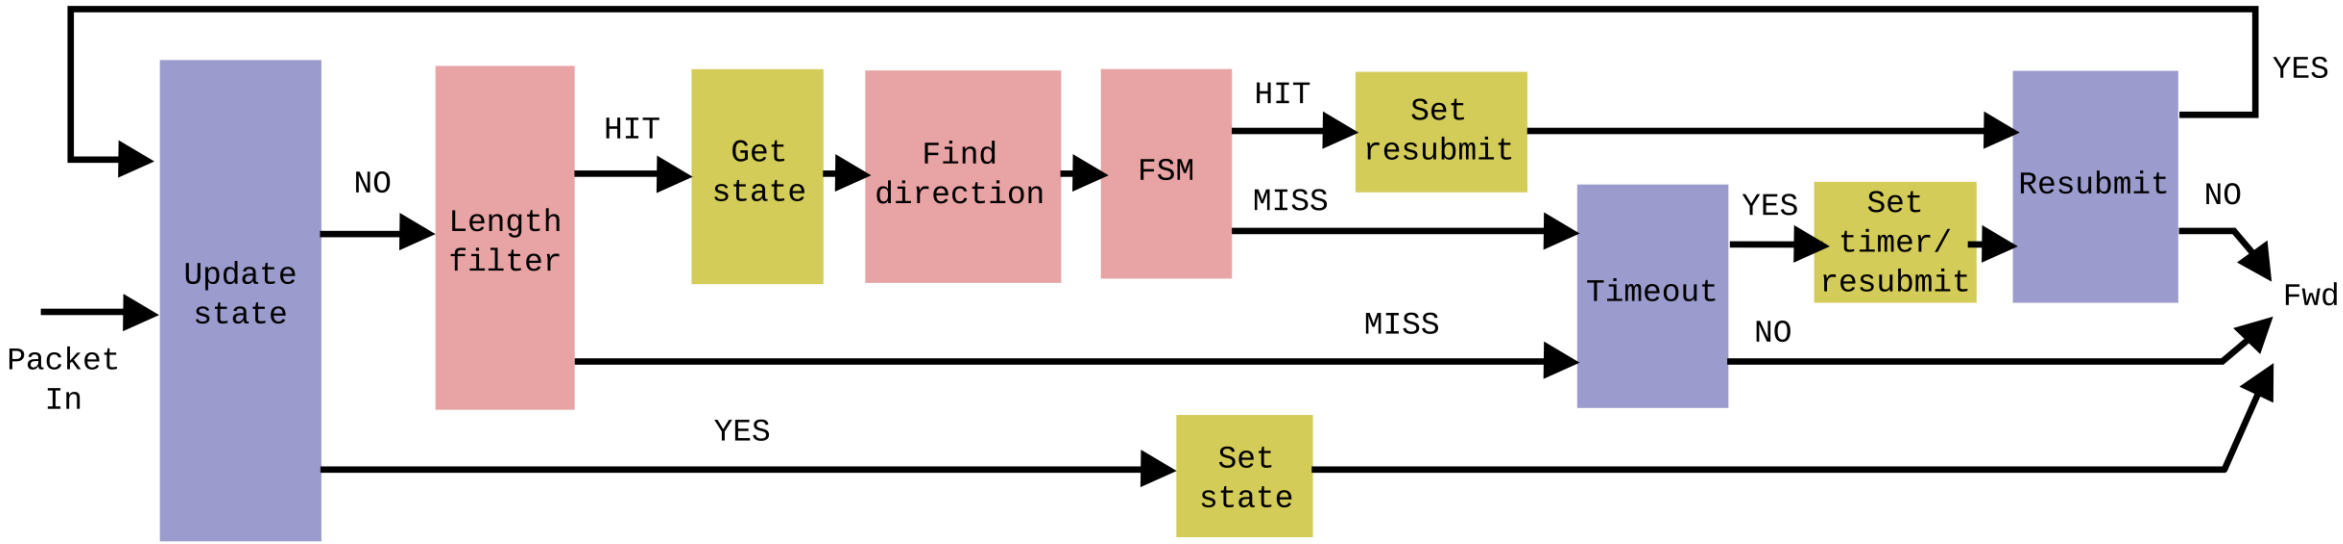

- Resubmit used instead of recirculation to update information without costs of engaging traffic manager

Evaluation

PoirIoT implemented on Wedge 100BF-32X with Tofino ASIC

Code Available (https://github.com/PINetDalhousie/poiriot)

IoT traffic from publicly available dataset

- 14 devices
- ~30 Signatures

More results/details in Paper

Device Detection

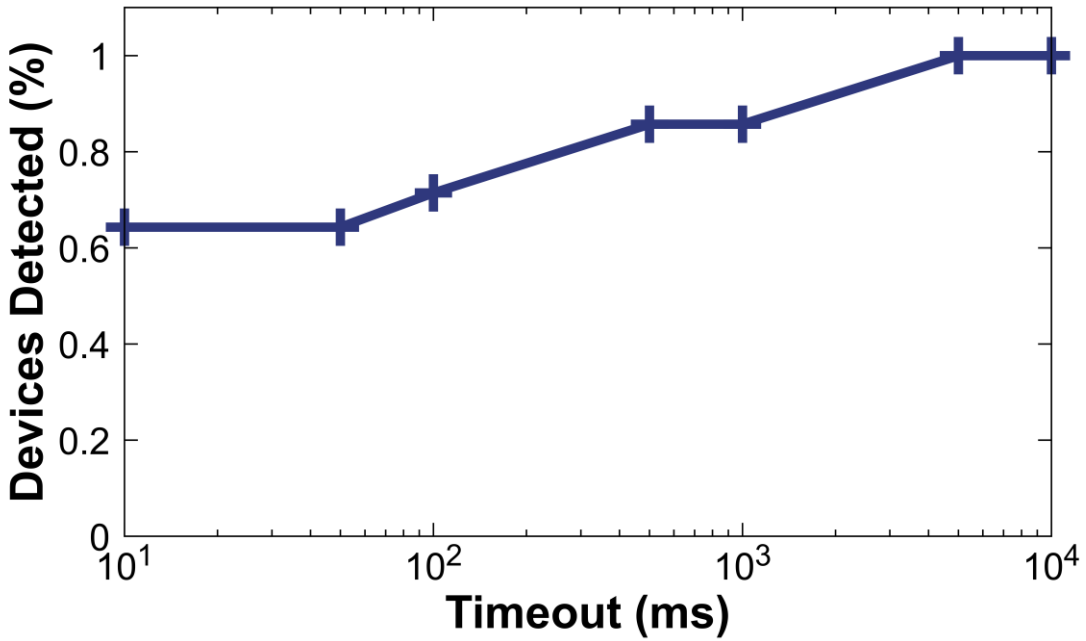

- Detect all 14 devices
- Timeouts needs to be sufficiently long to allow for RTT and longer signatures

Device Detection

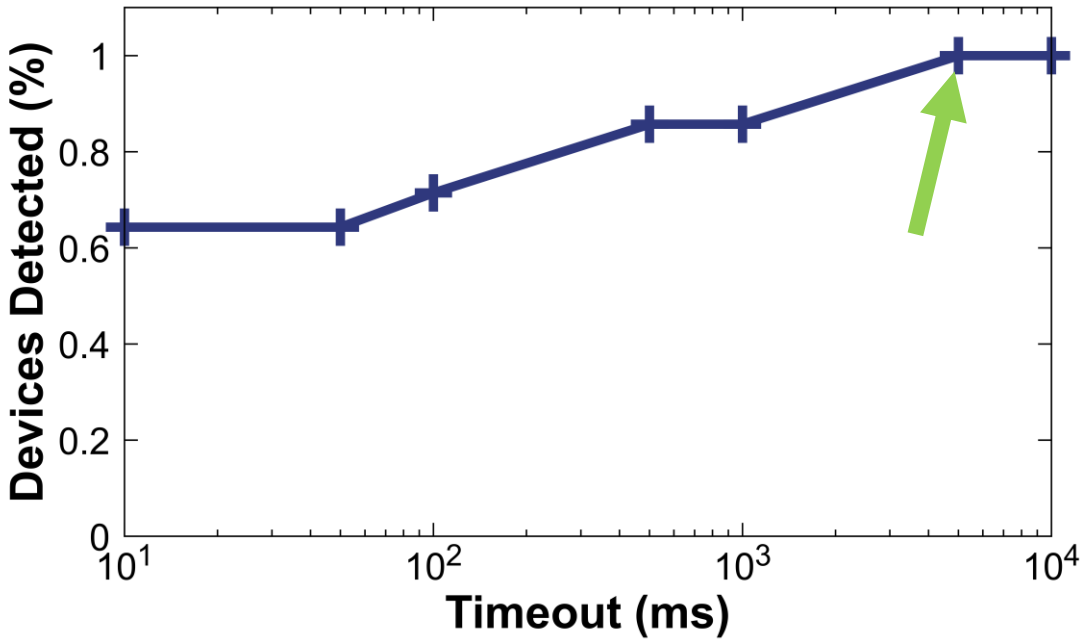

- Detect all 14 devices
- Timeouts needs to be sufficiently long to allow for RTT and longer signatures

Device Detection

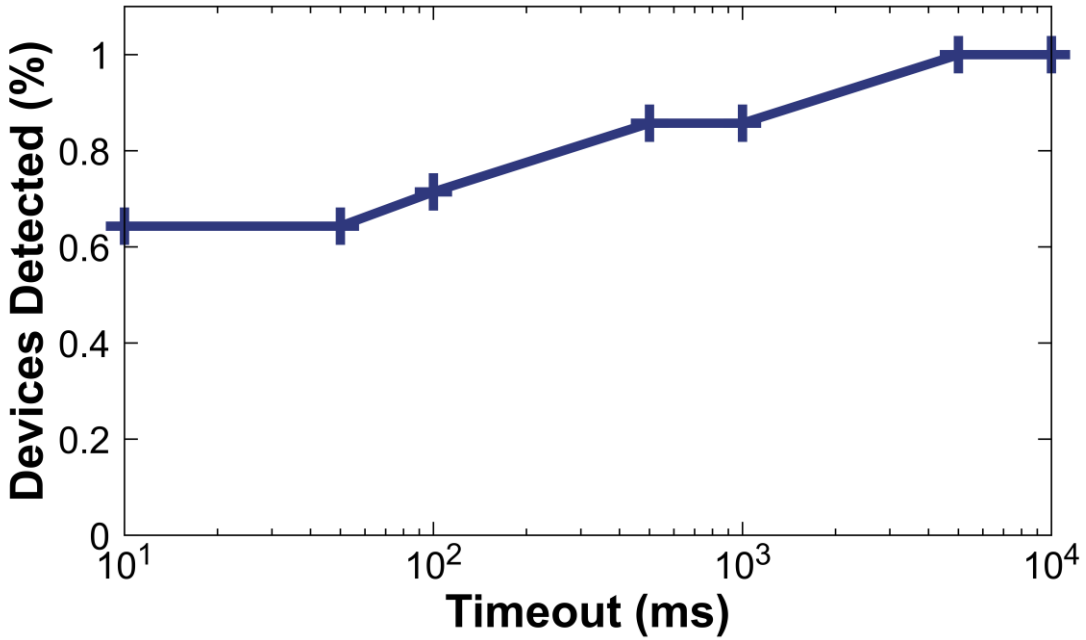

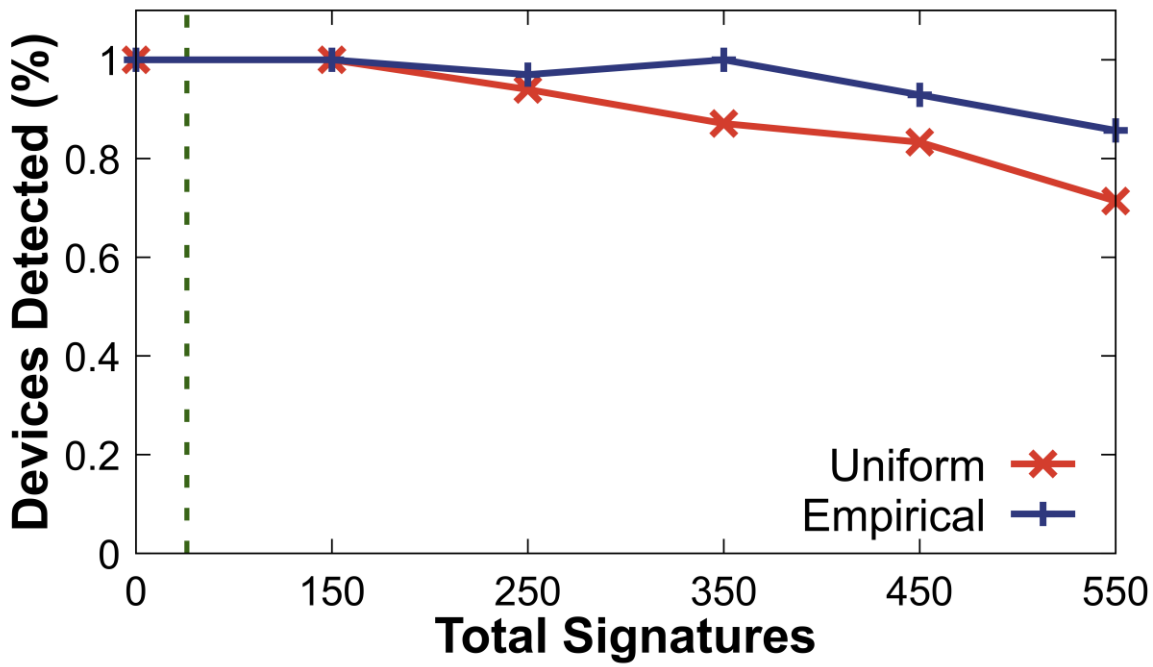

- Two types of synthetic signatures: uniformly sampled or derived from existing signature set
- Detection accuracy remains high (80%+) with hundreds of additional signatures

Device Detection

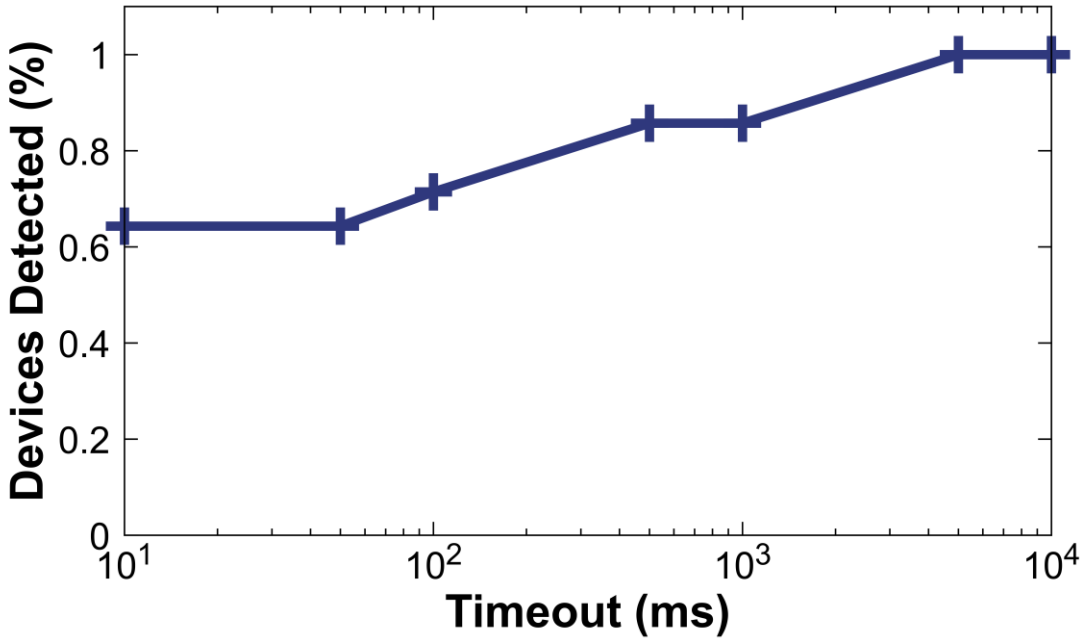

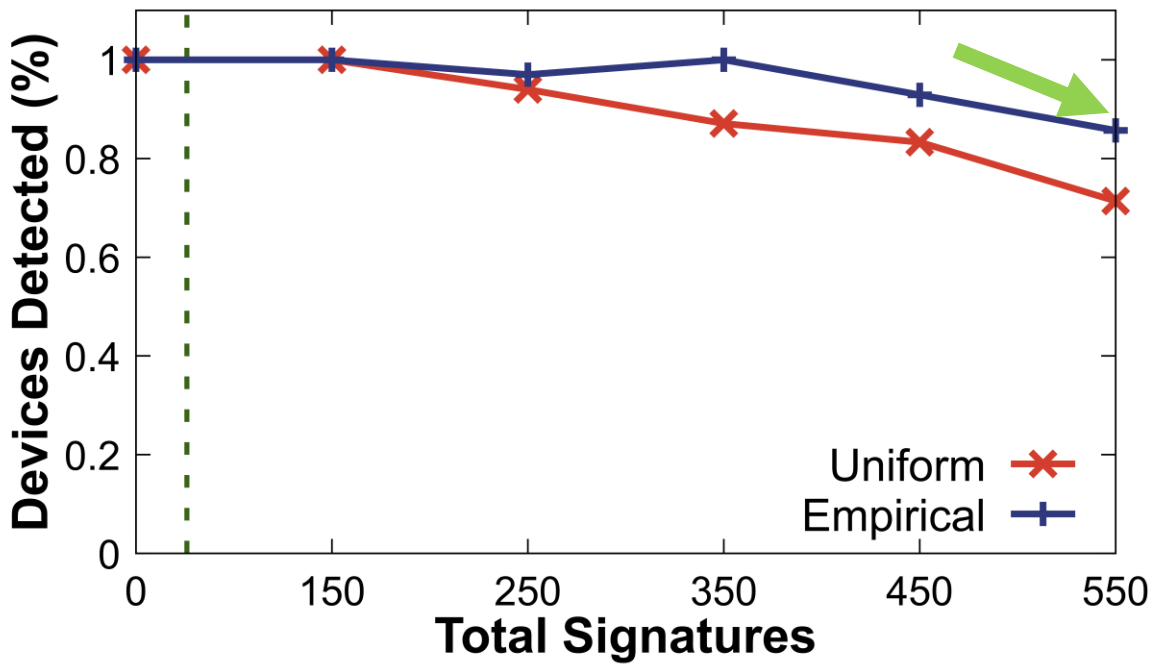

- Two types of synthetic signatures: uniformly sampled or derived from existing signature set
- Detection accuracy remains high (80%+) with hundreds of additional signatures

Scalability

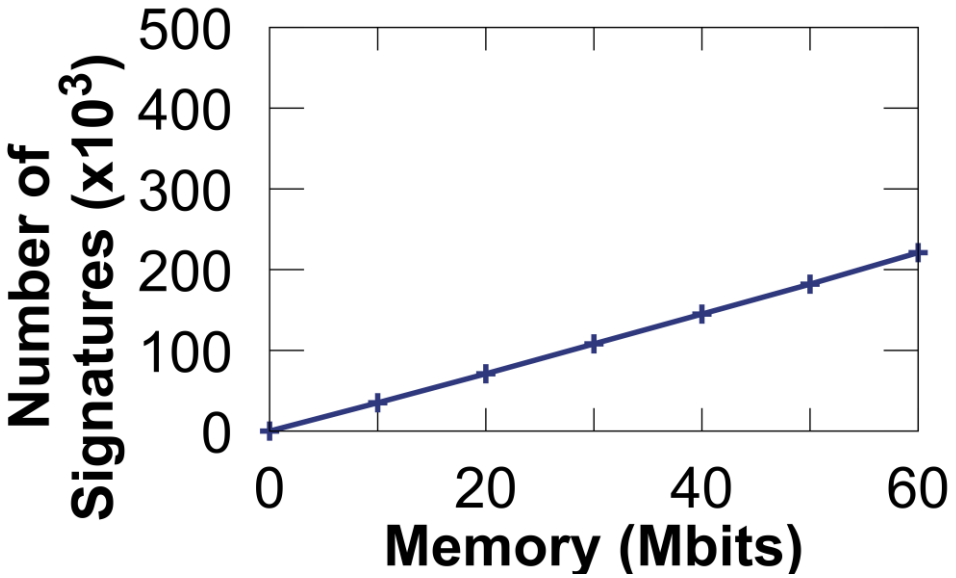

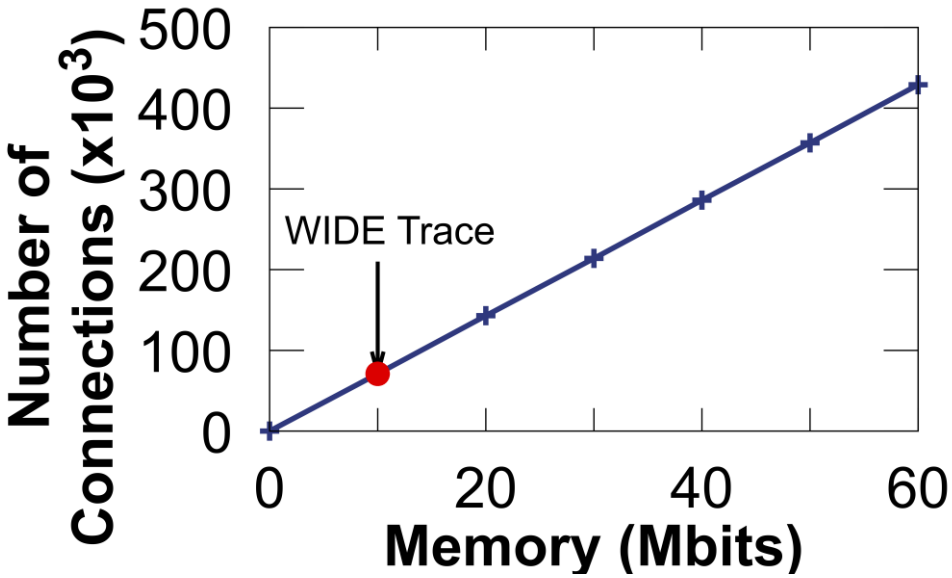

- Switch memory supports 200,000 signatures or 400,000 connections
- WIDE Trace shows flows the have a packet length matching one of our initial lengths

Scalability

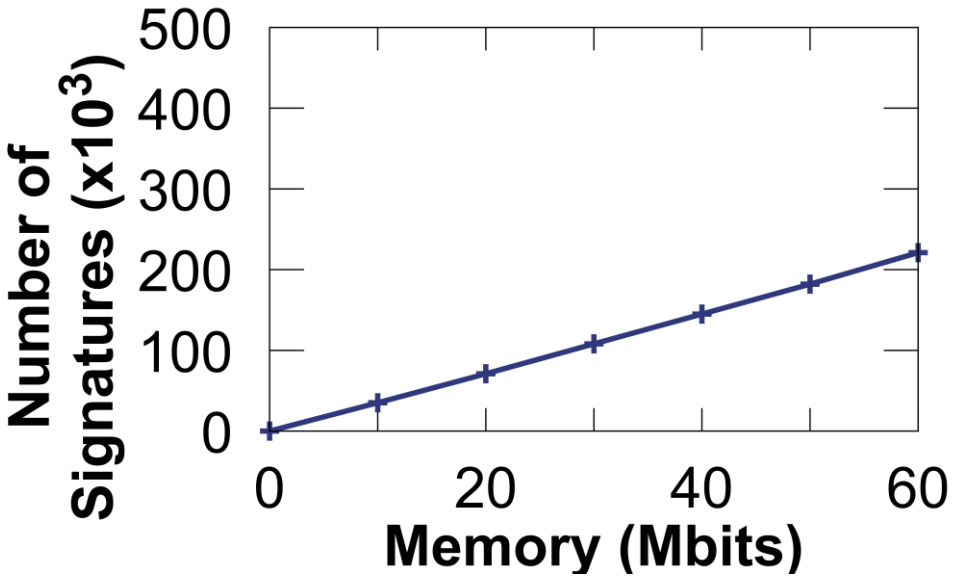

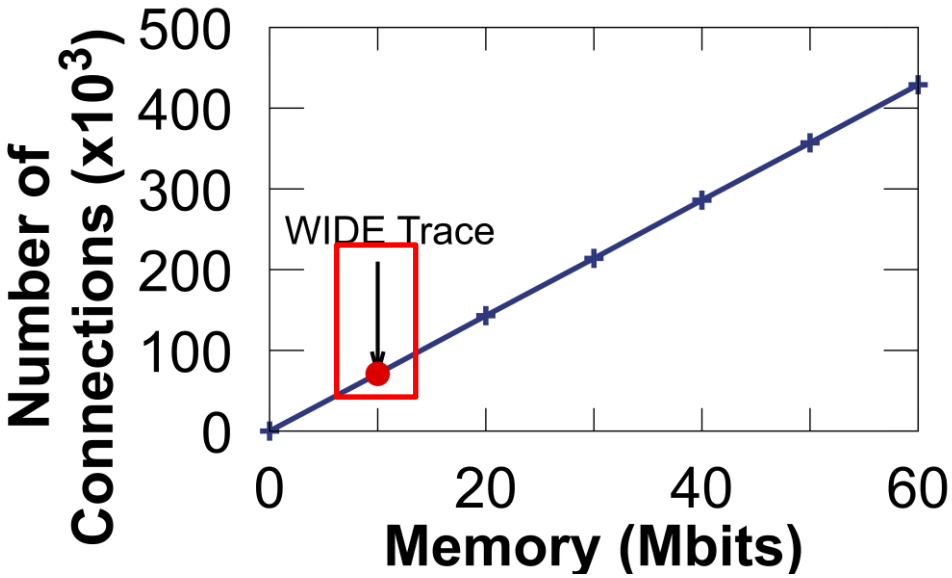

- Switch memory supports 200,000 signatures or 400,000 connections
- WIDE Trace shows flows the have a packet length matching one of our initial lengths

Resource Consumption

Resource	Usage
Match Crossbar	3.1%
SRAM	5.1%
TCAM	0%
VLIW Instruction	3.4%
Hash Bits	4.7%

- Application takes minimal amounts of switch resources (No TCAM at all)
- SRAM most consumed resource (stateful information + table entries)

Resource Consumption

Resource	Usage
Match Crossbar	3.1%
SRAM	5.1%
TCAM	0%
VLIW Instruction	3.4%
Hash Bits	4.7%

- Application takes minimal amounts of switch resources (No TCAM at all)
- SRAM most consumed resource (stateful information + table entries)

Resource Consumption

Resource	Usage
Match Crossbar	3.1%
SRAM	5.1%
TCAM	0%
VLIW Instruction	3.4%
Hash Bits	4.7%

- Application takes minimal amounts of switch resources (No TCAM at all)
- SRAM most consumed resource (stateful information + table entries)

Summary

- State of the art fingerprinting solutions face challenges (volume and granularity) at scale
- PoirIoT brings fingerprinting to data plane offering:
  - High Speed
  - Packet level granularity
- System consists of two components controller + switch data plane
- Solution detects 100% of devices in a publicly available data set while using minimal switch resources

Questions?

Carson Kuzniar
carson.kuzniar@dal.ca

References

1. A. Acar, H. Fereidooni, T. Abera, A. K. Sikder, M. Miettinen, H. Aksu, M. Conti, A.-R. Sadeghi, and S. Uluagac, “Peek-a-boo: I see your smart home activities, even encrypted!” in Proceedings of the 13th ACM Conference on Security and Privacy in Wireless and Mobile Networks, ser. WiSec ’20, 2020, p. 207–218.
2. Y. Afek, A. Bremler-Barr, D. Hay, R. Goldschmidt, L. Shafir, G. Avraham, and A. Shalev, “Nfv-based iot security for home networks using mud,” in NOMS 2020 - 2020 IEEE/IFIP Network Operations and Management Symposium, 2020, pp. 1–9
3. R. Trimananda, J. Varmarken, A. Markopoulou, and B. Demsky, “Packet-Level Signatures for Smart Home Devices,” Proceedings of the 2020 Network and Distributed System Security (NDSS) Symposium, February 2020.
4. C. Duan, S. Zhang, J. Yang, Z. Wang, Y. Yang, and J. Li, “Pinball: Universal and robust signature extraction for smart home devices,” in 2021 IFIP/IEEE International Symposium on Integrated Network Management (IM), 2021, pp. 1–9.

Server Vs Switch

- Server drops packets at high rates and misses events
- Switch operates at line rate so no degradation as rates increase

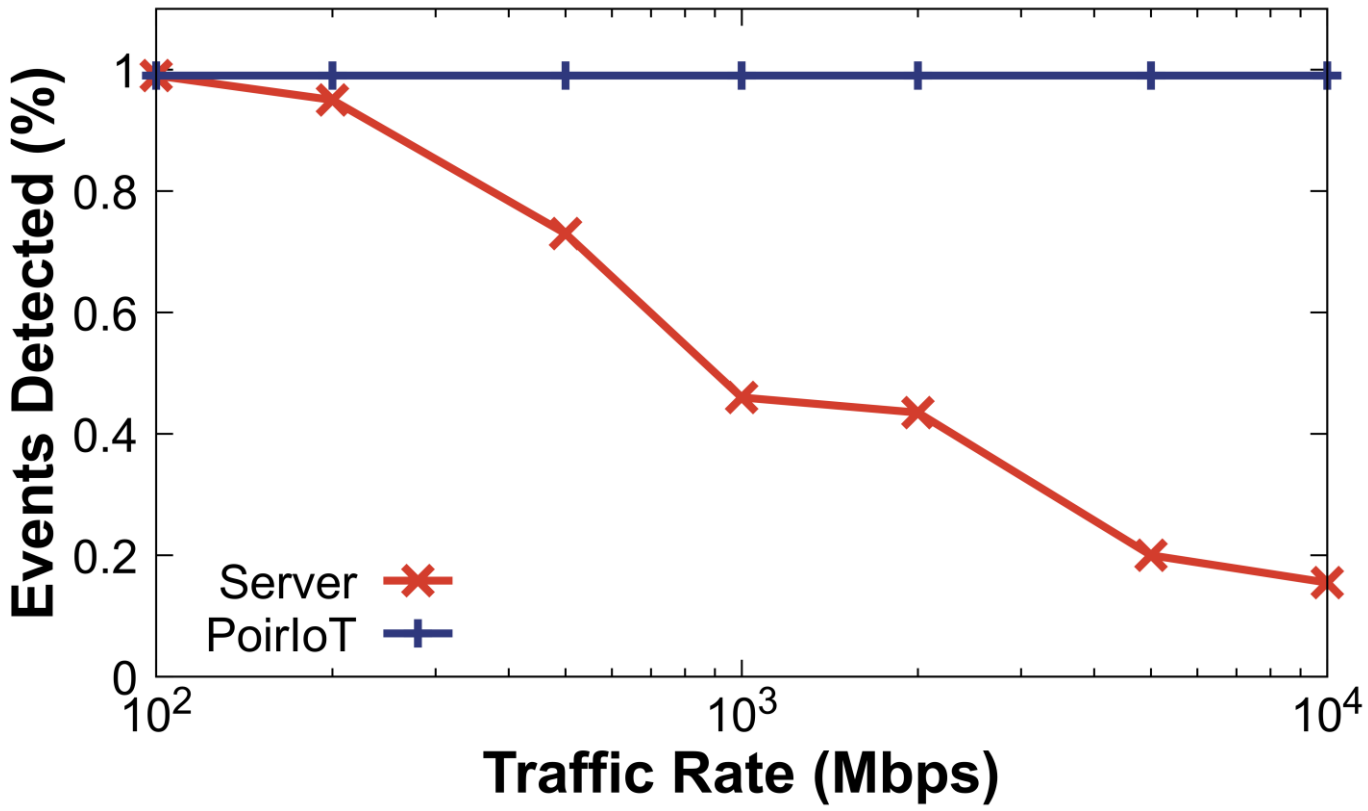

Event Detection

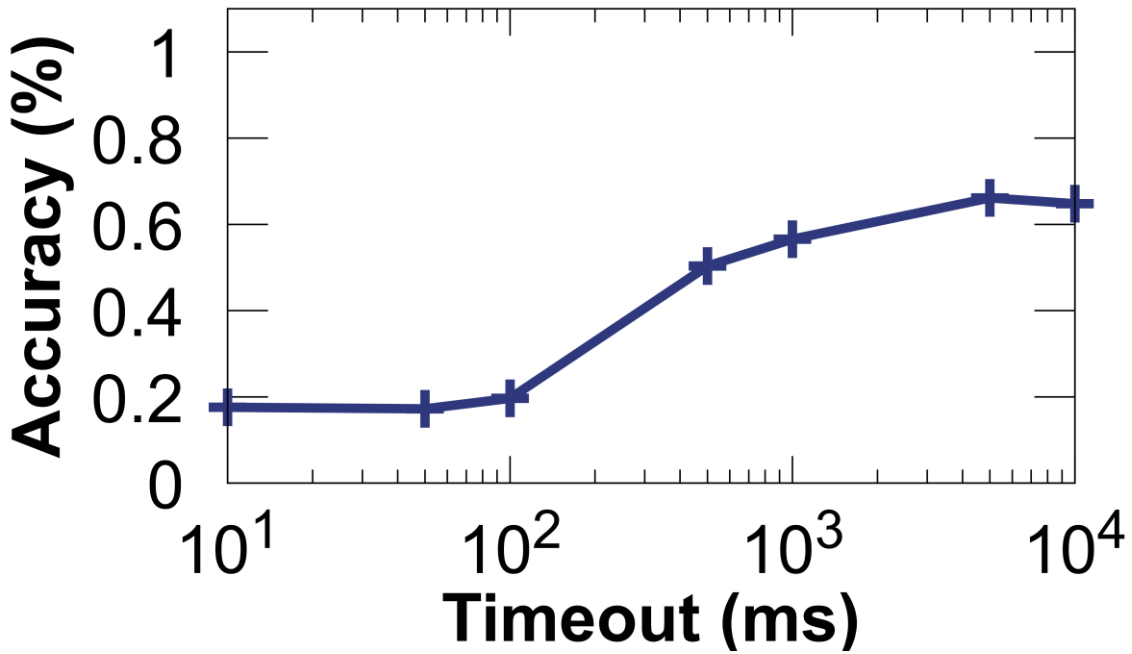

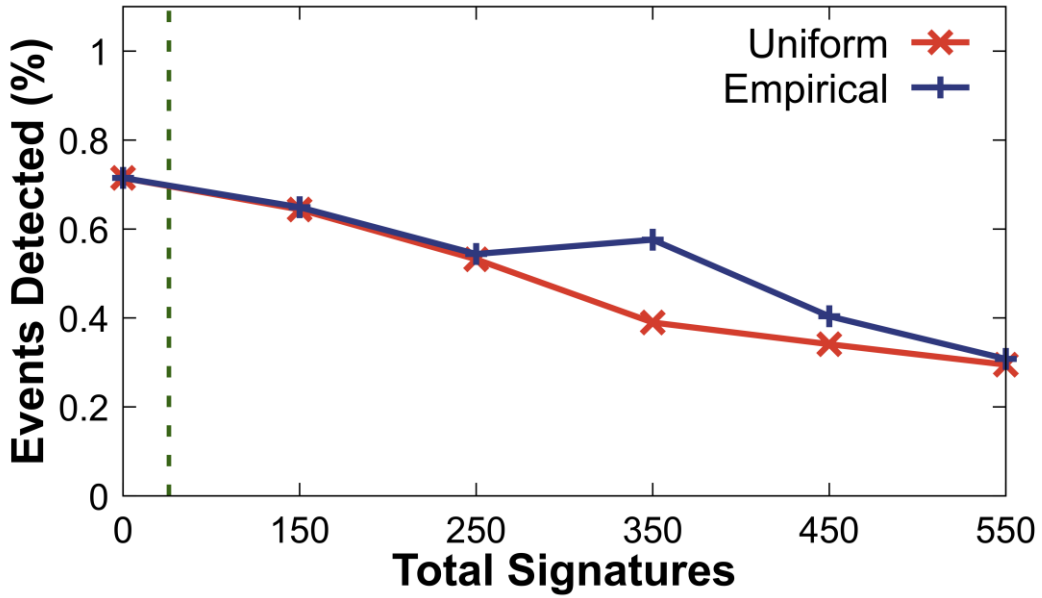

- Event accuracy starts to decrease at longer timeouts because of a lack of resets
- Event detection more granular so suffers more as signatures are added