Reinforcement Learning-Controlled Mitigation of Volumetric DDoS Attacks

Hauke Heseding

Institute of Telematics, Research Group Prof. Zitterbart

Volumetric DDoS Attacks

- Attackers send high-volume attack traffic
  - Attack traffic saturates bottleneck links
    - Elephant flows, amplification attacks (DNS, NTP, SSDP), …
  - Legitimate traffic is suppressed and targets availability is impeded

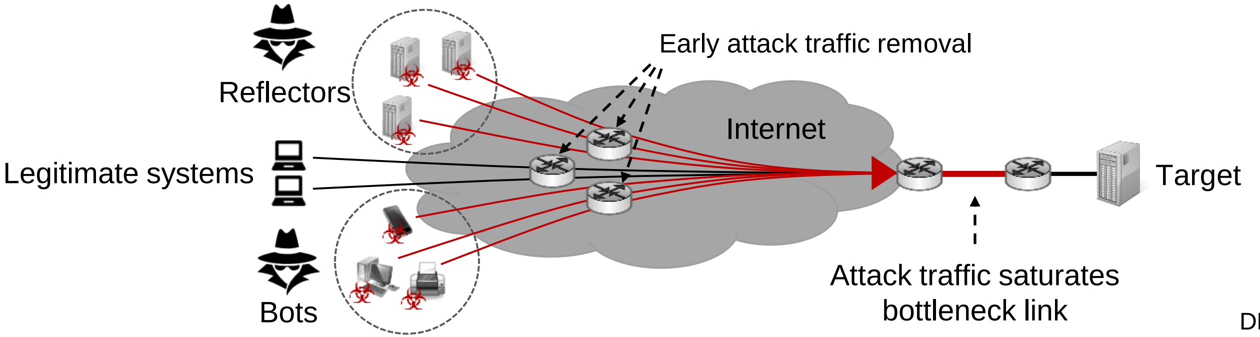

DNS Domain Name System
NTP Network Time Protocol
SSDP Simple Service Discovery Protocol

TCAM-Based Ingress Filtering

- Reduce infrastructure load
  - Identify suspicious IP source subnets
  - Establish upstream filter rules in TCAM
  - Cost and power consumption limit TCAM capacity

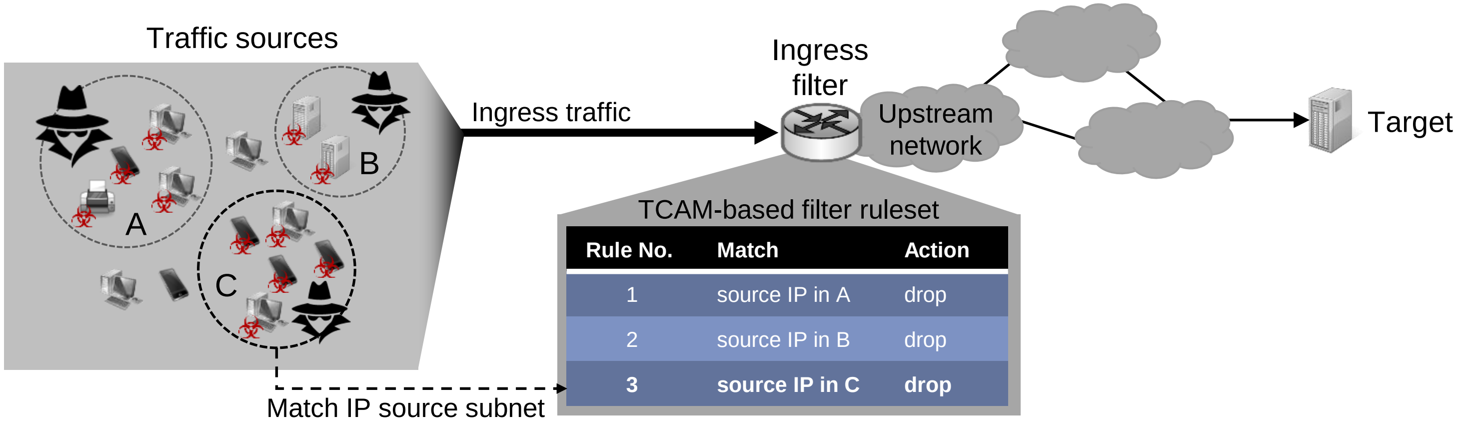

TCAM-based filter ruleset

Rule No.	Match	Action
1	source IP in A	drop
2	source IP in B	drop
3	**source IP in C**	**drop**

TCAM: Ternary content-addressable memory

Finding Effective Filter Rules

- Hierarchical Heavy Hitters (HHH) ⟶ detect suspicious IP subnets
- Reinforcement learning ⟶ adjust HHH thresholds

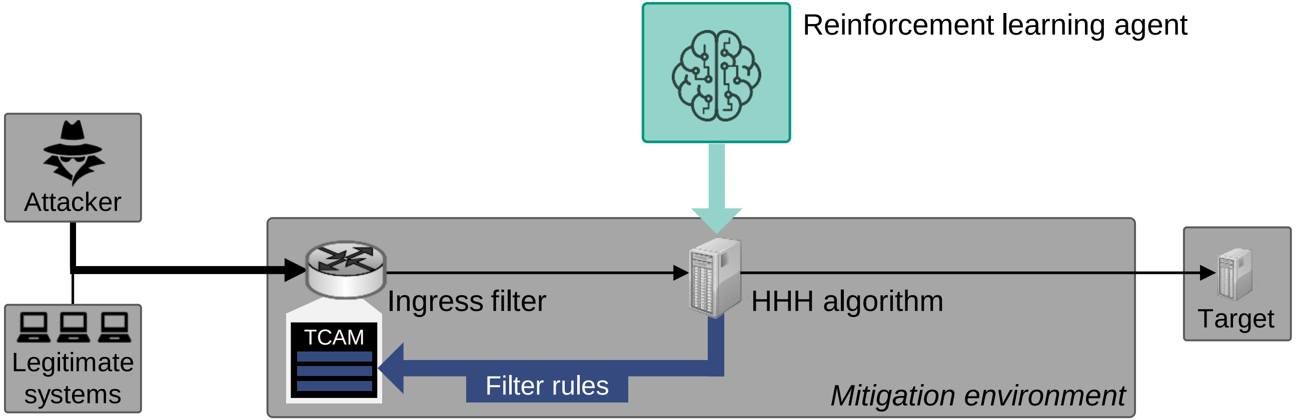

Hierarchical Heavy Hitters (HHH)

- Find IP subnets sending at least fraction ϕ of total traffic
  - Aggregate traffic volume by IP subnet
  - Select filter rules from identified HHHs

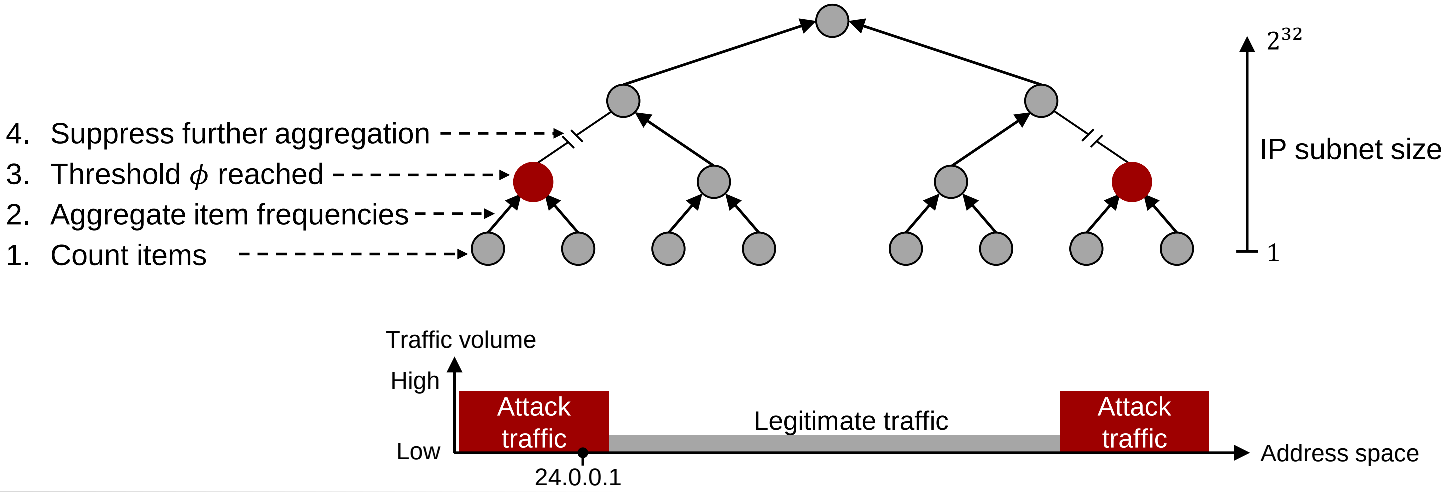

Filter Rule Selection

- Avoid over-aggregation
  - Hierarchy threshold $H^{\max}$ limits aggregation
- Avoid excessive TCAM utilization
  - Adjust minimum frequency threshold ϕ

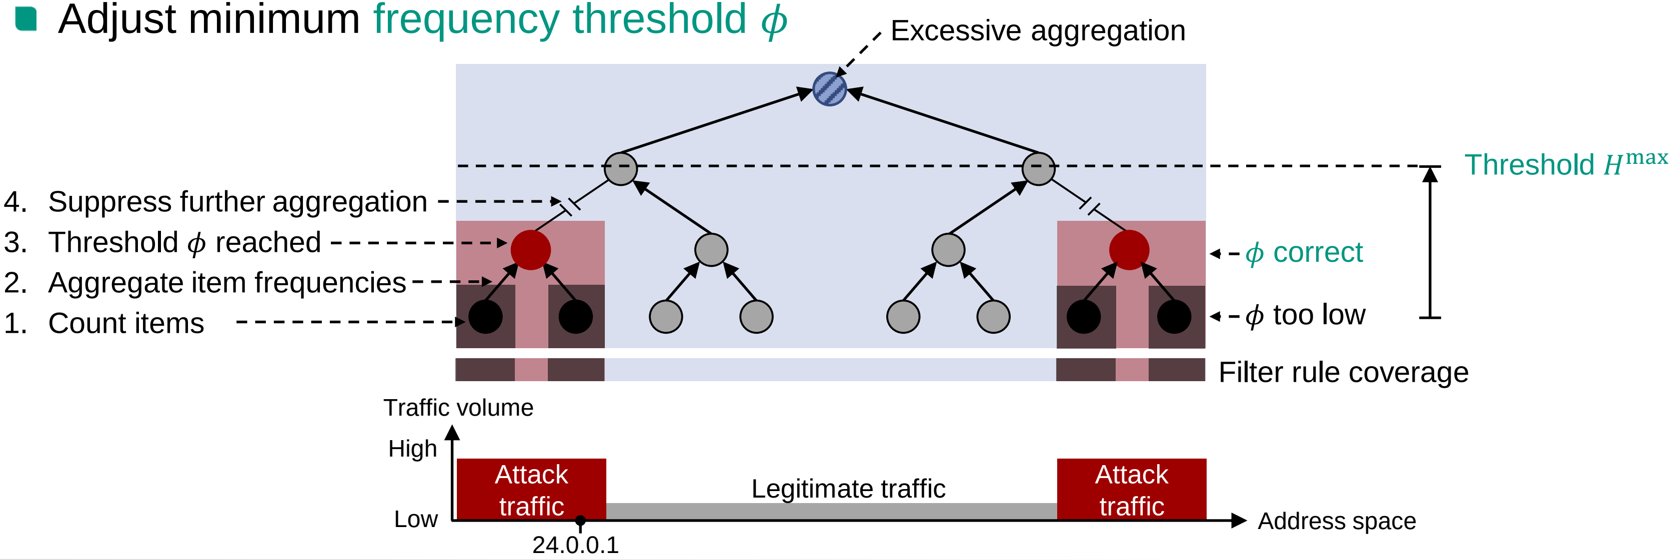

Finding Effective Filter Rules

- Hierarchical Heavy Hitters (HHH) ⟶ detect suspicious IP subnets
- Reinforcement learning ⟶ adjust HHH thresholds

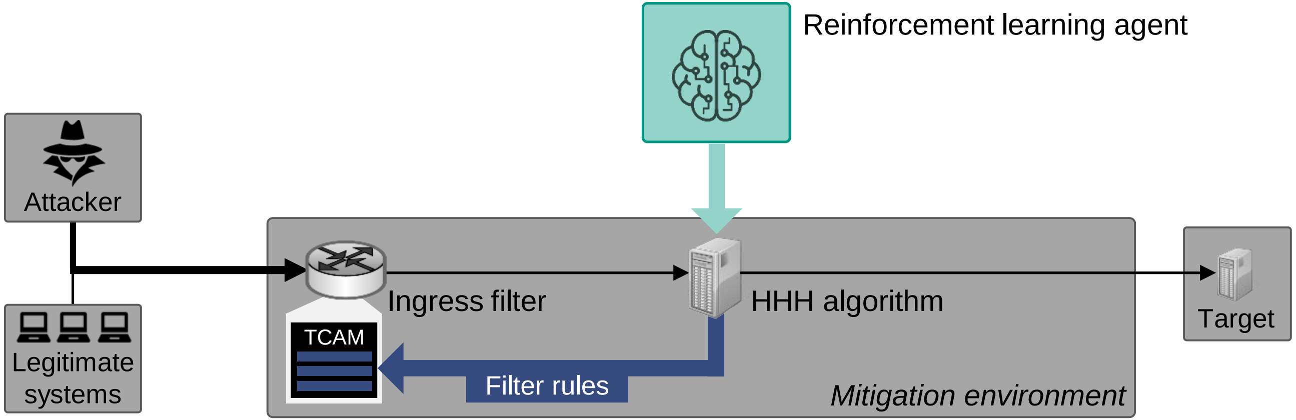

How to Choose Effective Thresholds?

Karlsruhe Institute of Technology

- Deep Reinforcement Learning with Deep Q-Networks (DQN)
  - Agent observes traffic distribution and filter effectiveness
  - Agent adapts thresholds when traffic patterns evolve over time
  - Agent learns over time from interaction with mitigation environment

Filter Rule Adaptation

- Continuously executed control loop for threshold adaptation
  1. Query HHH with selected parameters
  2. Propagate HHH-derived filter rules upstream to TCAM
  3. Agent observes TCAM utilization and filter rule effectiveness
  4. Agent adapts thresholds to match traffic pattern

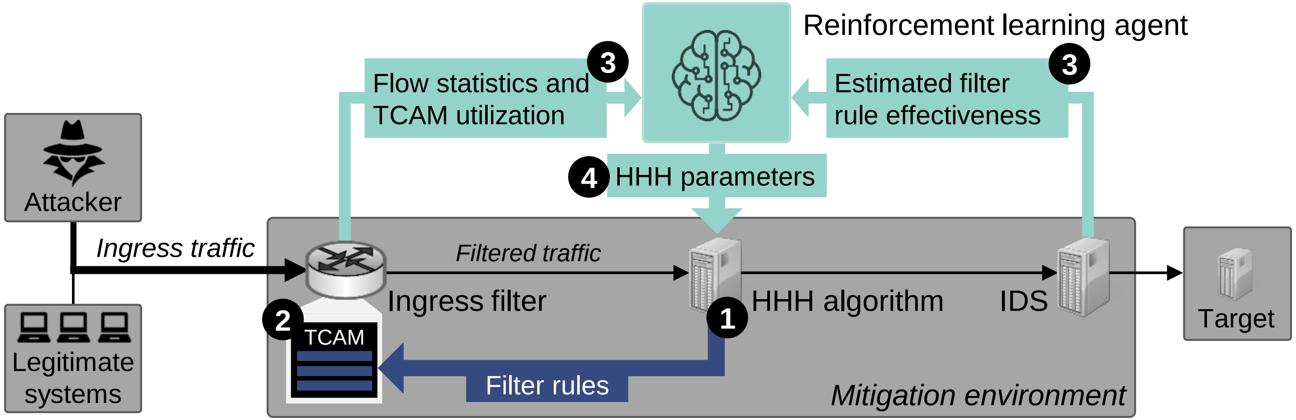

Observations and Actions

- Observation space
  - Number and of distribution of detected HHHs and filter rules
  - Estimated precision, sensitivity, false positive rate

- Discrete action space: $A_1, A_2, \ldots$
  - Represents possible parameter combinations

Coarse-grained rules

	$\phi_1 = 0.01$	$\phi_2 = 0.02$	…	$\phi_{24} = 0.24$	$\phi_{25} = 0.25$
$H^{\max} = 24$	A_1	A_2	…	A_{24}	A_{25}
…	…	…	…	…	…
$H^{\max} = 16$	A_{201}	A_{202}	…	A_{224}	A_{225}

Fine-grained rules

Reward Function Modelling

- Reward function $\mathfrak{r}(p, s, f, r) = \mathfrak{r}_p(p) \cdot \mathfrak{r}_s(s) \cdot \mathfrak{r}_f(f) \cdot \mathfrak{r}_r(r)$
  - Polynomial factors
    - Precision p, sensitivity s, false positive rate f, filter rule count r
  - Different emphasis on mitigation goals

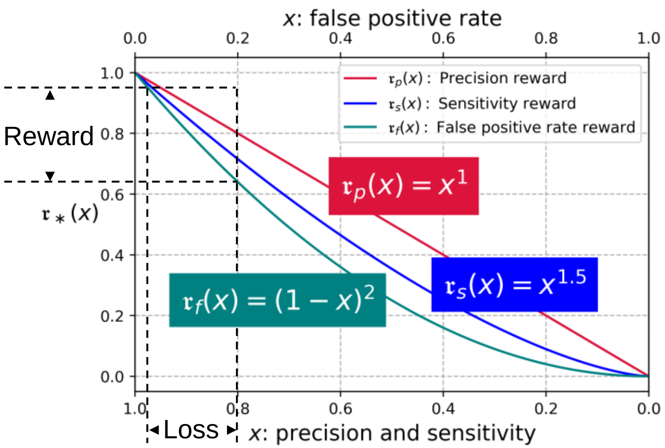

Simulated Traffic Scenario

- Randomized traffic source activity over time
  - Four phases with different attack traffic patterns

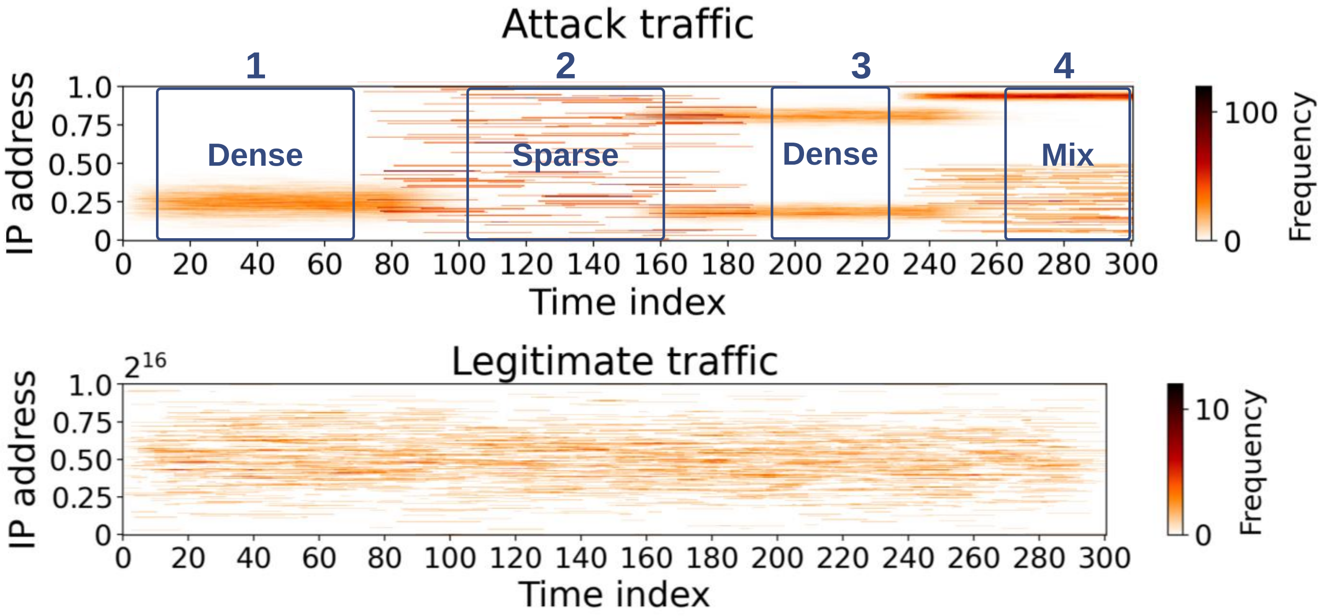

A Snapshot Filter Rule Selection

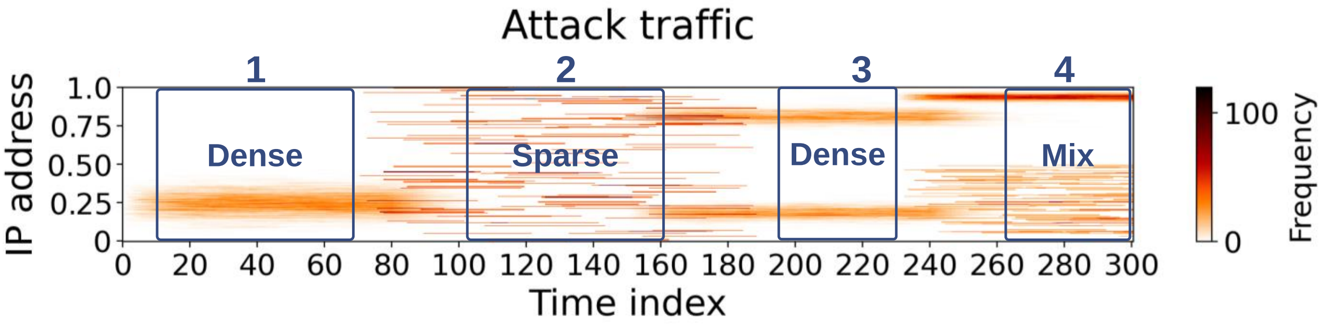

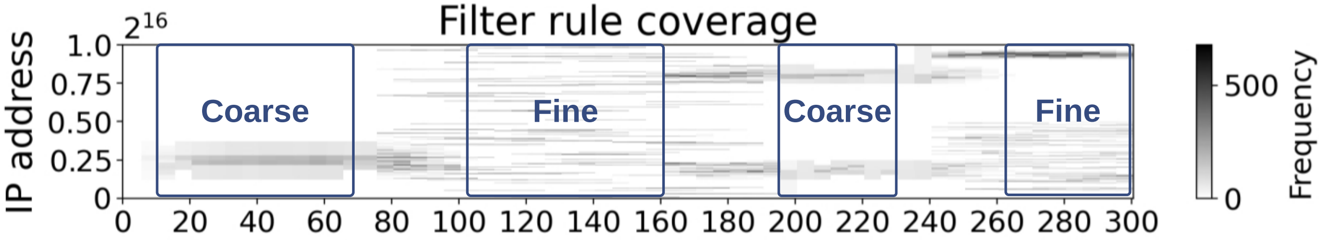

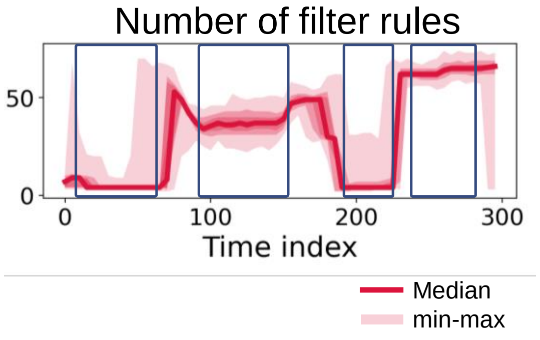

Selected Results

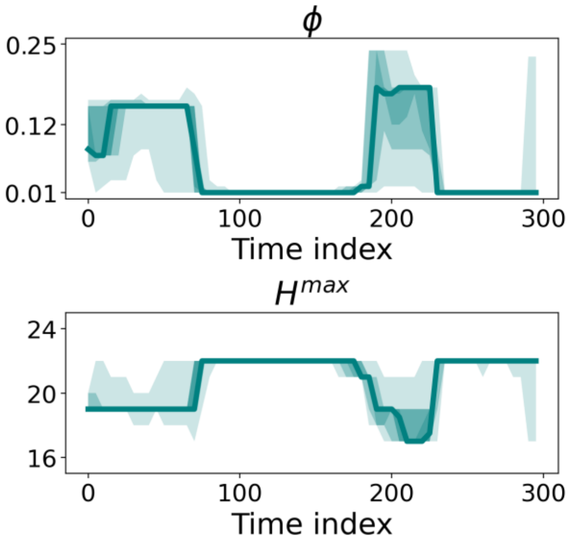

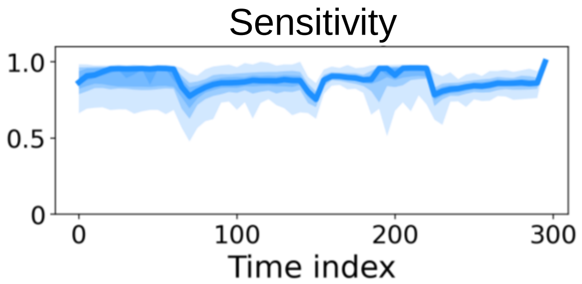

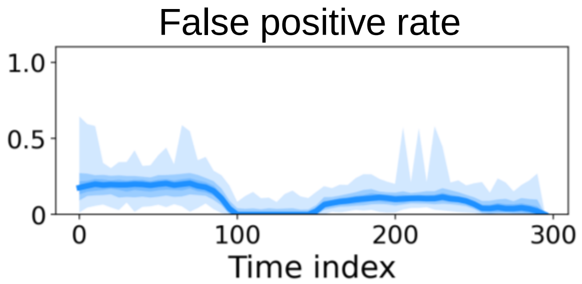

Agent adapts to traffic patterns

Significant attack traffic reduction

Low false positive rate maintained

Conclusion

- TCAM-based ingress filtering
  - Hierarchical heavy hitters for attack traffic source identification
  - Upstream propagation of filter rules for early traffic filtering
  - Agent learns and adapts thresholds to adapt filter rules

- In comparison
  - Avoids extensive state keeping of microflow-based traffic segmentation[3,4,5,6]
  - Respects traffic composition typically disregarded by router throttling[1,2]

- Simulative evaluation
  - Significant attack traffic reduction
  - Maintains low false positive rates

References

[1] Mahajan, Ratul, et al. "Controlling high bandwidth aggregates in the network." *ACM SIGCOMM Computer Communication Review 32.3* (2002): 62-73.

[2] Malialis, Kleanthis, and Daniel Kudenko. "Distributed response to network intrusions using multiagent reinforcement learning. *Engineering Applications of Artificial Intelligence* 41 (2015): 270-284.

[3] Simpson, Kyle A., Simon Rogers, and Dimitrios P. Pezaros. "Per-host DDoS mitigation by direct-control reinforcement learning.“ *IEEE Transactions on Network and Service Management* 17.1 (2019): 103-117.

[4] Zhang, Menghao, et al. "Poseidon: Mitigating volumetric ddos attacks with programmable switches." the *27th Network and Distribute System Security Symposium (NDSS 2020)*. 2020.

[5] Phan, Trung V., et al. "DeepGuard: Efficient anomaly detection in SDN with fine-grained traffic flow monitoring." *IEEE Transactions on Network and Service Management* 17.3 (2020): 1349-1362.

[6] Liu, Zaoxing, et al. "Jaqen: A High-Performance Switch-Native Approach for Detecting and Mitigating Volumetric DDoS Attacks with Programmable Switches." *30th USENIX Security Symposium (USENIX Security 21)*. 2021