Karlsruhe Institute of Technology

Institute of Telematics
Research Group Prof. Zitterbart

Reinforcement Learning-Controlled Mitigation of Volumetric DDoS Attacks

Hauke Heseding

Volumetric DDoS Attacks

- Attackers send high-volume attack traffic
  - Attack traffic congests bottleneck links
- Reduce infrastructure load by early traffic filtering

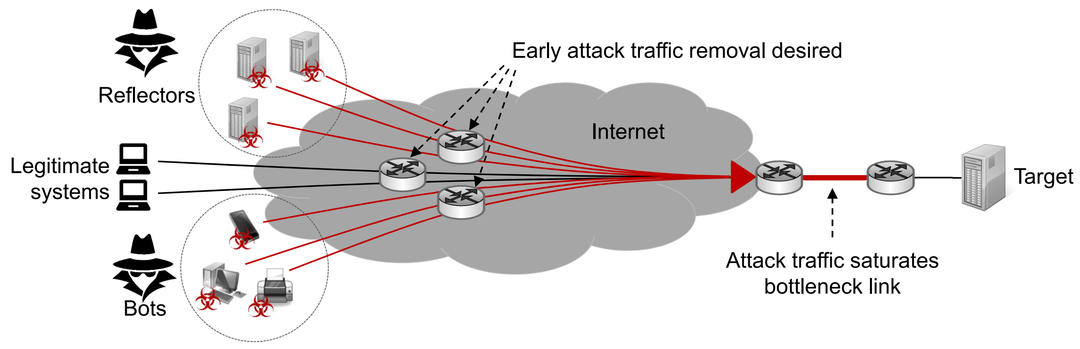

TCAM-Based Ingress Filtering

- Establish ingress filter in upstream network
  - Ternary content-addressable memory (TCAM)
  - Evaluate filter rules in a single clock cycle
- Adapt to changing traffic patterns

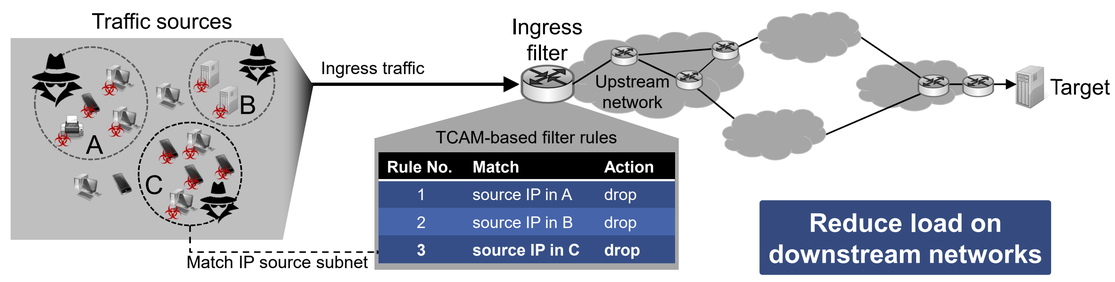

Hierarchical Heavy Hitters (HHH)

- Identify attack traffic sources
  - Find subnets sending fraction ϕ of total traffic
  - Preserve limited TCAM capacity ⟶ choose ϕ
  - Avoid excessive aggregation ⟶ choose $H^{\max}$

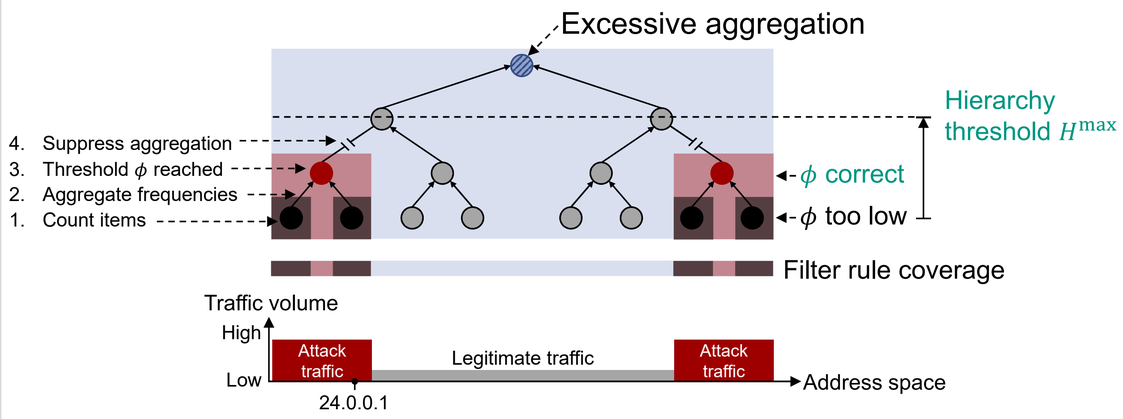

Adaptive Filter Rule Selection

- Reinforcement learning with Deep Q-Networks
  - Agent learns from traffic patterns
  - Agent selects ϕ, $H^{\max}$ for HHH algorithm
  - Agent adapts filter rule granularity to traffic pattern

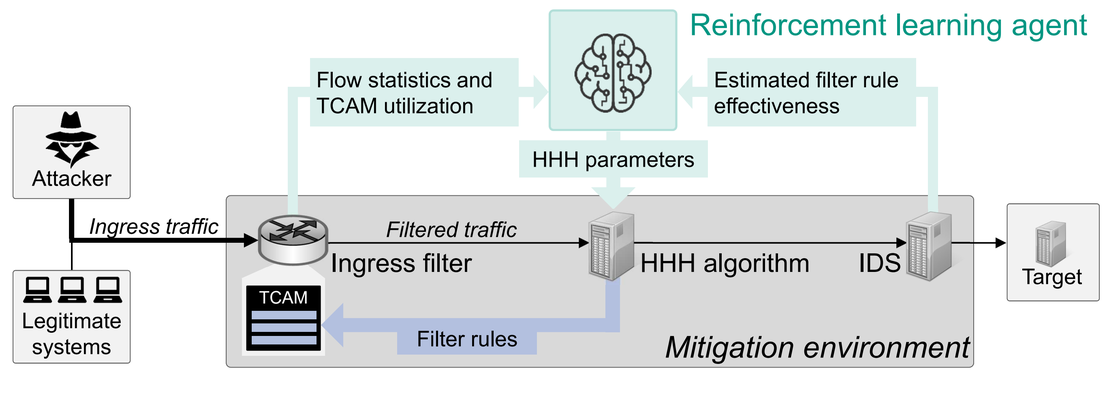

Effective Trade-Offs

- Reward function prioritizes mitigation goals
  - Precision p, sensitivity s, false positive rate f, filter rule count r
  - $\mathfrak{r} = \mathfrak{r}_p \cdot \mathfrak{r}_s \cdot \mathfrak{r}_f \cdot \mathfrak{r}_r$
  - Prioritization by tuning polynomial factors $\mathfrak{r}_*$
  - Agent learns to realize effective trade-offs

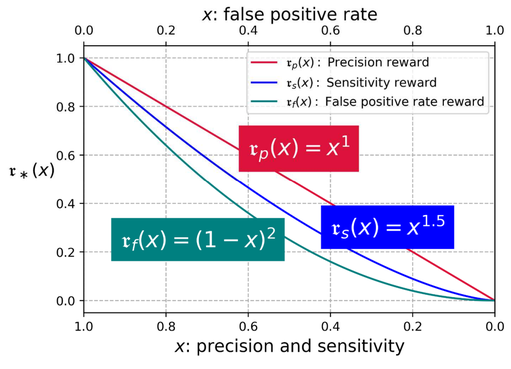

Simulated Dynamic Traffic Scenario

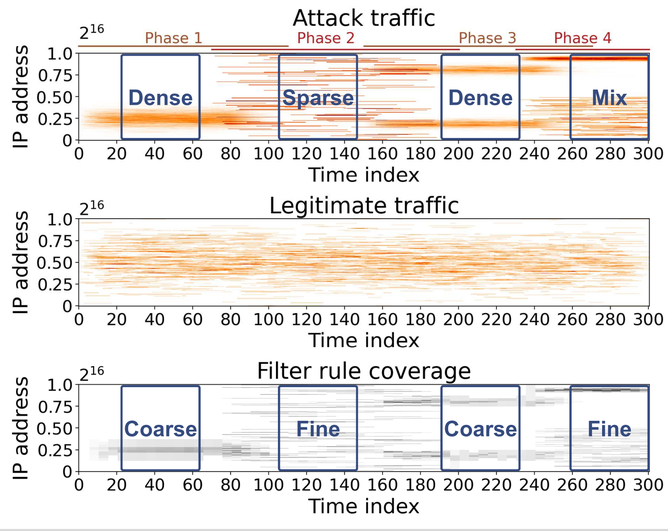

Changing attack pattern

Filter rules adapted to attack traffic

Hauke Heseding
Institute of Telematics
hauke.heseding@kit.edu

KASTEL
Institute of Information
Security and Dependability