USC University of Southern California

# AMON-SENSS

## Scalable DDoS Detection for ISPs

Jelena Mirkovic (USC/ISI), Rajat Tandon (USC)

# DDoS Attack Detection Challenge

- Most attacks create large volume at the target
  - Some attacks do not
  - Some targets can handle large volume
- Most attacks are very short or intermittent
  - We do not want frequent false positives but want to detect and handle large attacks
- Most attacks launched by numerous sources
  - So are port scans
  - Some attacks launched by one source or a few of them

# DDoS Attack Signature Challenge

- Signature derivation is hard
  - Usually requires modeling how normal traffic looks like for a given destination, using many features
  - It does not scale to keep statistics about every potential attack target at an ISP (many records, many targets)
- CPU cost of processing each packet/flow
- Memory cost of storing statistics
- Many of the records are stored needlessly
  - The destination does not come under attack
  - Most of the stats stored not relevant for the signature

# AMON

- We were inspired by AMON [1] by Merit Networks
  - Keeps statistics for detection in **a matrix of bins**, aggregating traffic between many source-destination pairs
    - Volume and/or number of packets
  - Use Boyer-Moore algorithm to detect heavy-hitter sources and destinations for each bin

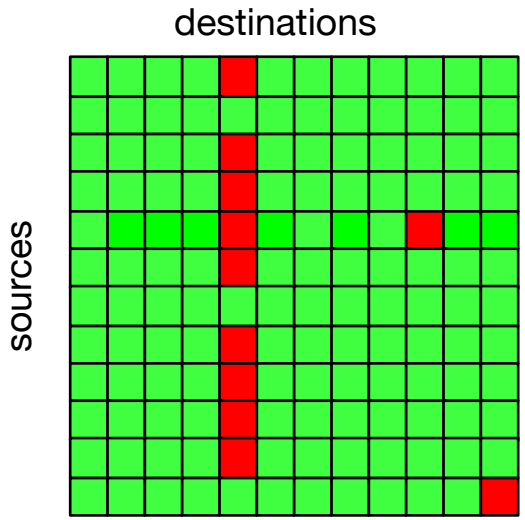


[1] https://arxiv.org/abs/1509.00268

# AMON-SENSS

- Keeps statistics for detection in **an array of bins**, aggregating traffic to many addresses
  - Volume


  - Asymmetry score (number and type of asymmetric pkts)
    - For a flow: asym_score = asym_factor * num_pkts

| proto | flags | src port | dst port | asym_factor |
|---|---|---|---|---|
| TCP | PSH | any | any | 0 |
| TCP | no PSH | service | user | -1 |
| TCP | no PSH | user | service | 1 |
| UDP | n/a | service | user | -1 |
| UDP | n/a | user | service | 1 |

# Volume and Asymmetry

- Both volume and asymmetry must be *abnormal* to detect a possible attack
  - Abnormal here means not within their historic ranges
    - mean ± 5*stdev
  - High volume but asymmetry within expected ranges may mean large data transfers, which destination can handle
  - High asymmetry but volume within expected ranges may mean scanning activity
- We can also require that abnormality lasts for some sustained period
  - To avoid large scans triggering detection
- To detect an attack's stop:
  - Both volume and asymmetry must remain within their historical ranges for a sustained period of time

# Illustration

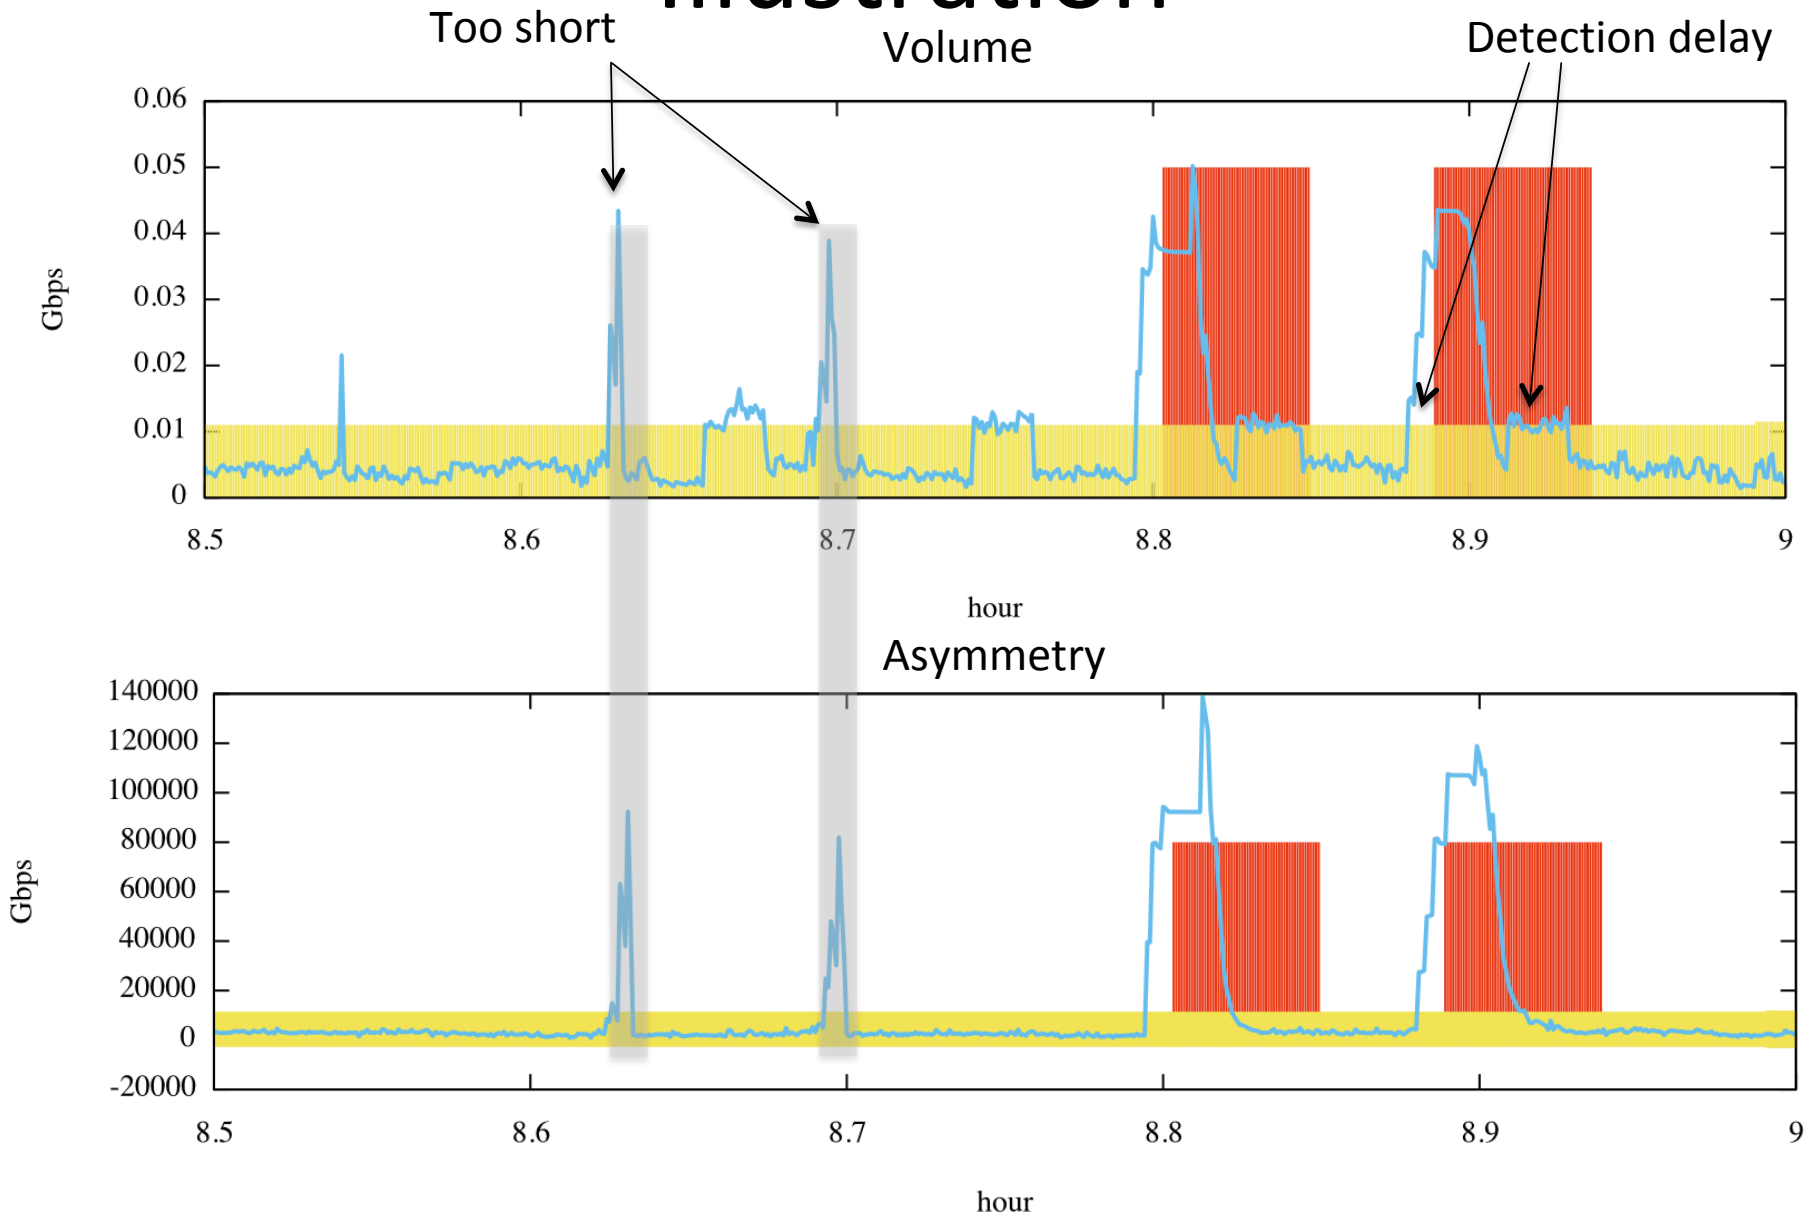

# Signature Generation

- Proactively sample flows whose asymmetry matches asymmetry of the bin
  - Whenever both volume and asymmetry are abnormal
- Proactively generate signatures over samples
  - Masking src IP, src port, dst port
  - Keeping proto and dst IP
  - For each combination keep only the most representative signature − covering most samples
- Find a signature that covers enough samples
  - And explains most of the asymmetry seen
  - Prefer more specific signatures but only when they are not much worse at explaining the asymmetry

# Illustration

| Signature | Asymmetry explained |
|---|---|
| *:* ➔164.76.176.0:* udp | 97% |
| *:* ➔164.76.176.0:43967 udp | <1% |
| *:53 ➔164.76.176.0:* udp | 95% |
| *:53 ➔164.76.176.0:43967 udp | <1% |
| 58.177.216.0:* ➔164.76.176.0:* udp | <1% |
| 58.177.216.0:* ➔164.76.176.0:43967 udp | <1% |
| 58.177.216.0:53 ➔164.76.176.0:* udp | <1% |
| 58.177.216.0:53 ➔164.76.176.0:43967 udp | <1% |

A more specific signature performs a bit worse than less specific one
But has lower chance of false positives

# Signature Testing

- Once signature is generated we test it to see if it works well
  - Collect matching flows in direct and reverse direction
  - Once enough flows are collected, evaluate how many are good (non TCP symmetric or TCP PSH) and how many are bad (asymmetric)
  - If good/(good+bad) < threshold proclaim this is a good signature and install it
- Always collect bin statistics prior to dropping
  - Also collect info how much traffic is dropped

# Illustration

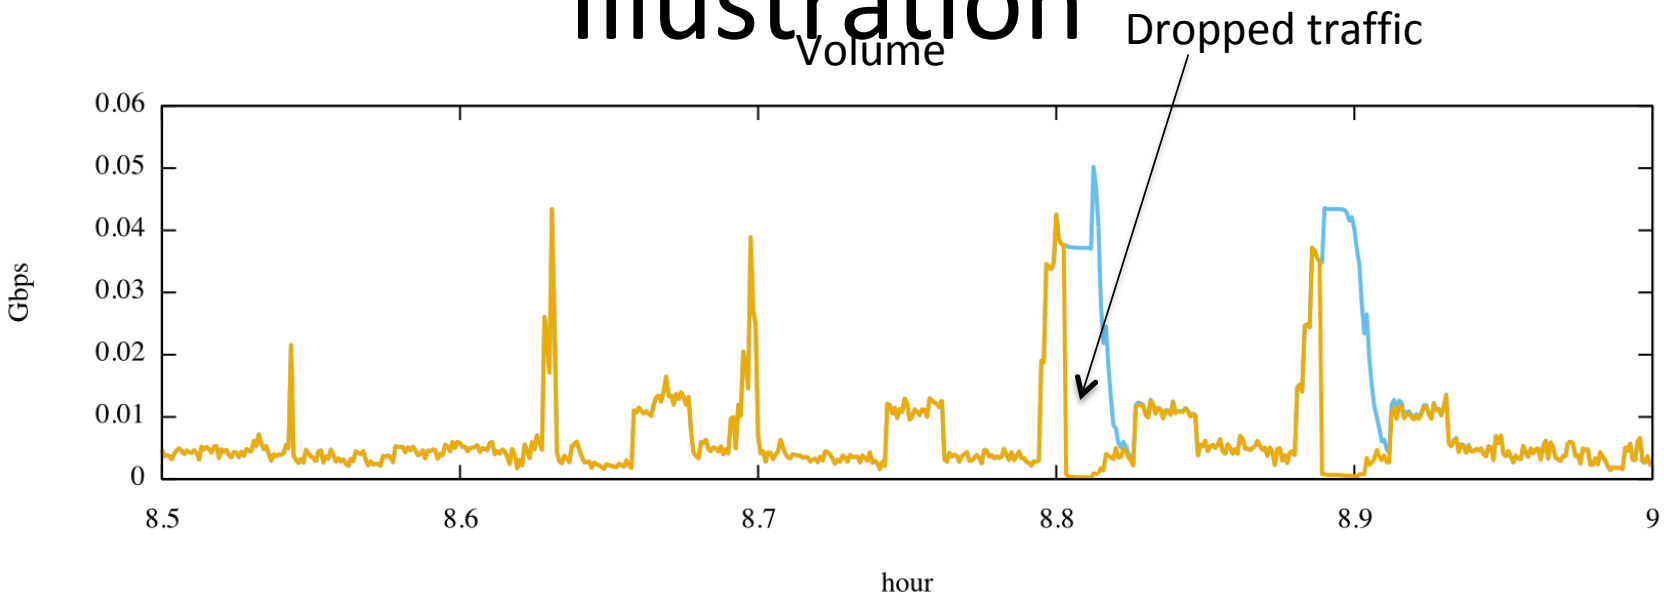


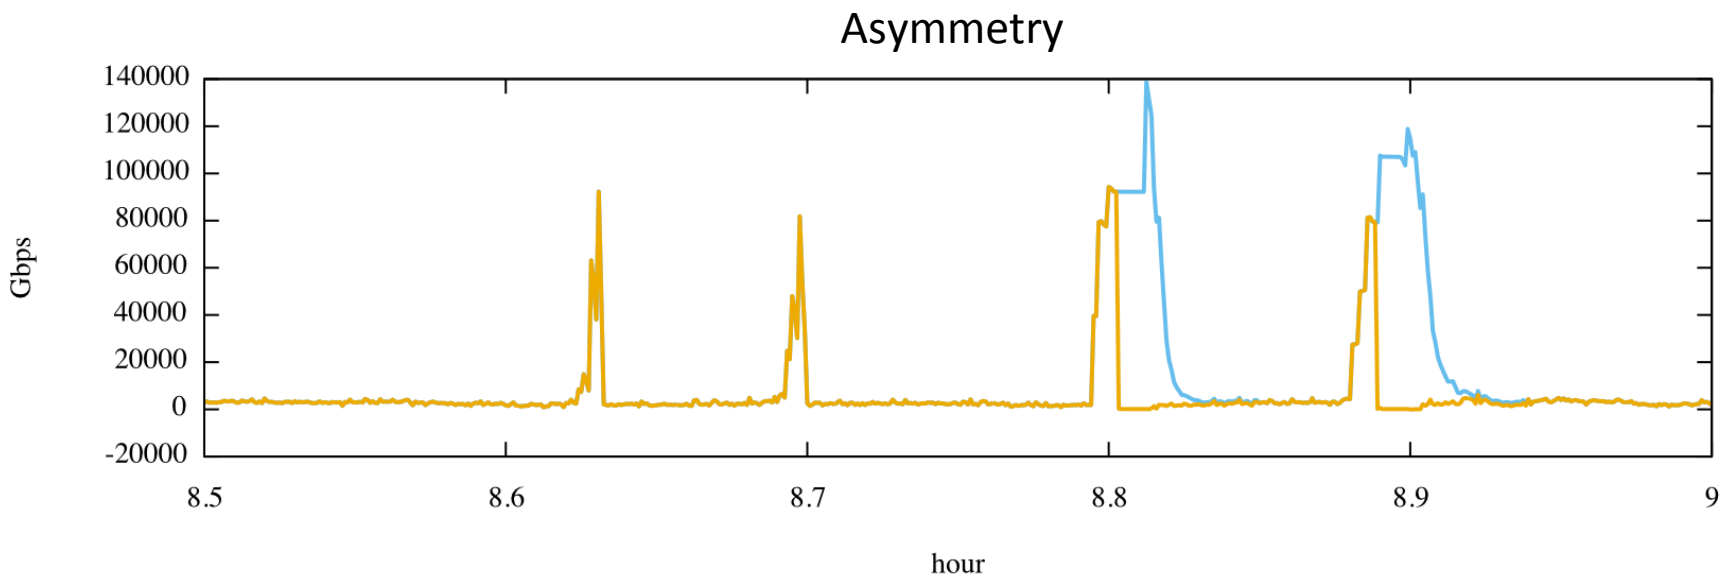

# Testing

- Tested on five Merit Network attack traces from IMPACT
  - Detected all the attacks noted in the metadata
  - Detected many more attacks

| trace | duration | #attacks | Largest size | Longest duration |
|---|---|---|---|---|
| chargen | 1 day | 61 | 0.9 Gbps syn flood | 19h syn flood |
| dns_ampl | 1 day | 43 | 4.5 Gbps DNS reflection | 0.5h NTP reflection |
| ntp-ddos | 2 weeks | 2,448 | 2.4 Gbps NTP reflection | 6 days syn flood |
| radb_ddos | 2 days | 71 | 1.8 Gbps DNS reflection | 0.5 h UDP flood |
| ssdp | 2 hours | 1 | 0.03 Gbps, 5 min SSDP reflection flood | |

# Observations About Attacks

- Number of attacks per day consistent with known literature
  - Recent IMC paper [1] finds 20M attacks in 2 years ~ 100 per day in a network of MeritNet's size
- Duration of attacks is also consistent with [1]
  - Most attacks are short and on-off, which makes detection and mitigation hard
- Long-lasting attacks are usually low volume

[1] Millions of Targets Under Attack: a Macroscopic Characterization of the DoS Ecosystem
M. Jonker, A. King, J. Krupp, C. Rossow, A. Sperotto, A. Dainotti
ACM SIGCOMM Internet Measurement Conference IMC 2017

# AMON-SENSS Performance

- Processes 6h of traffic in 1h
- Very small memory footprint
  - Large CPU footprint, mostly for Netflow reading, can be controlled per process