Cybersecurity (Security+) and P4 Programmable Switches

Heavy Hitters, Syn Flood

Elie Kfoury, Jorge Crichigno
University of South Carolina
http://ce.sc.edu/cyberinfra

Western Academy Support and Training Center (WASTC)
University of South Carolina (USC)
Energy Sciences Network (ESnet)

June 23rd, 2023

Outline

- Heavy hitters
- Heavy hitters detection
- Count-min sketch
- SYN flood

Heavy Hitters

Heavy Hitters

- Heavy hitters are a small number of flows that constitute most of the network traffic over a certain amount of time
- It is important to promptly detect heavy hitters in order to react to them
- Example: volumetric DDoS attack (like SYN flood) sending a large number of packets to the same destination (victim)

Heavy Hitter Detection

- Heavy hitter detection consists of tracking flows' packet counts
- Using a simple register array to store the counts is not scalable
  - Millions of flows traverse through the network
  - The memory on switches is not large enough to store counts for all flows

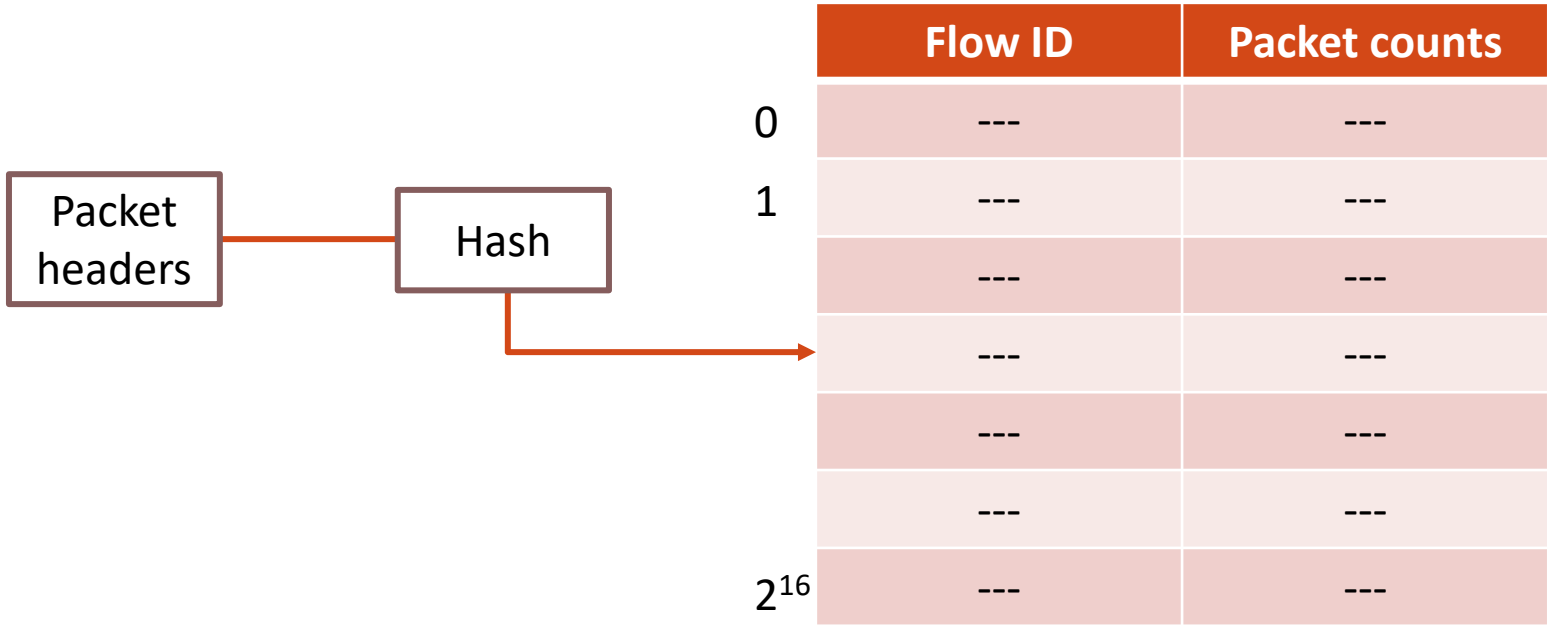

Heavy Hitter Detection

- Heavy hitter detection consists of tracking flows' packet counts
- Using a simple register array to store the counts is not scalable
  - Millions of flows traverse through the network
  - The memory on switches is not large enough to store counts for all flows

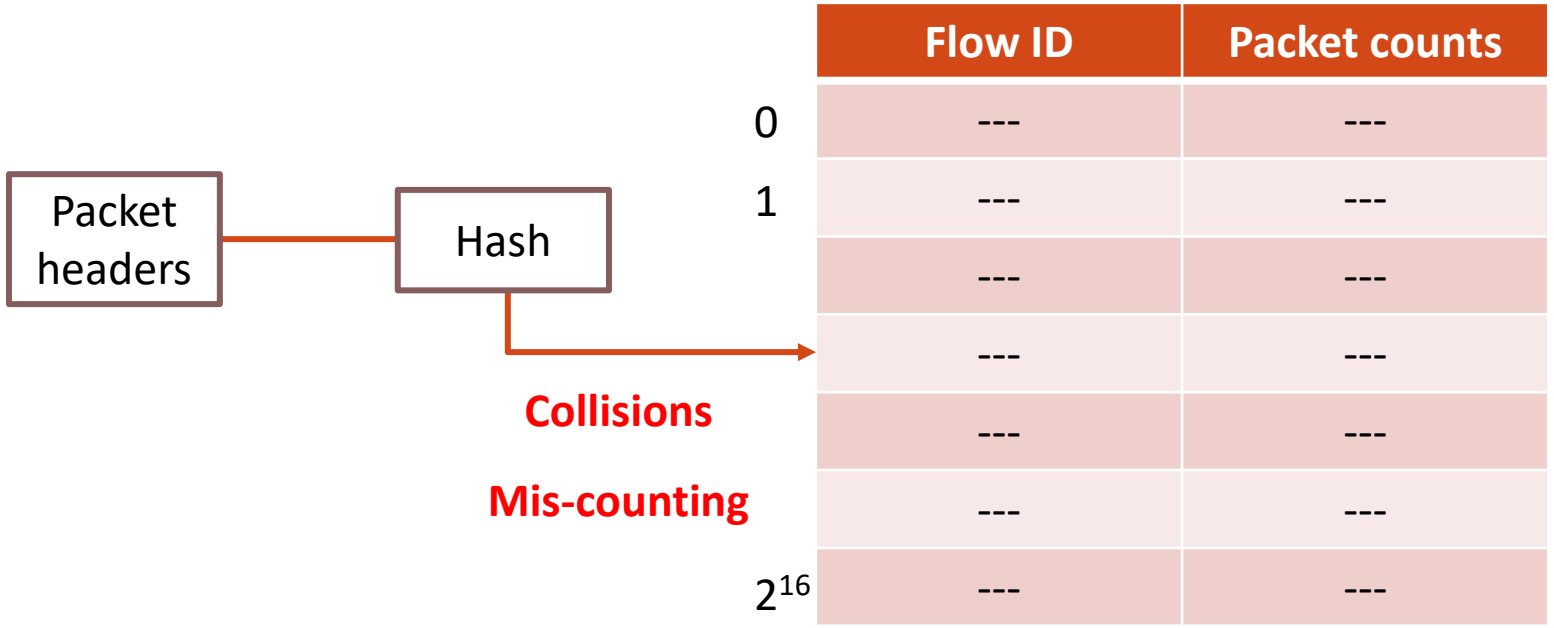

Count-min Sketch (CMS)

- Count-Min Sketch (CMS) is a probabilistic data structure
- CMS provides an efficient solution for estimating item frequencies
- CMS offers a tunable trade-off between accuracy and memory usage by adjusting its parameters

Count-min Sketch (CMS)

- The CMS consists of ***d*** register arrays that contain ***w*** cells each
- CMS uses ***d*** pairwise-independent hash functions h_i, ..., h_d that are applied to the 5-tuple fields in the packet headers
- The results are indices where counts are stored and incremented
- Calculating the minimum between these counts gives an approximation of the packet counts per flow

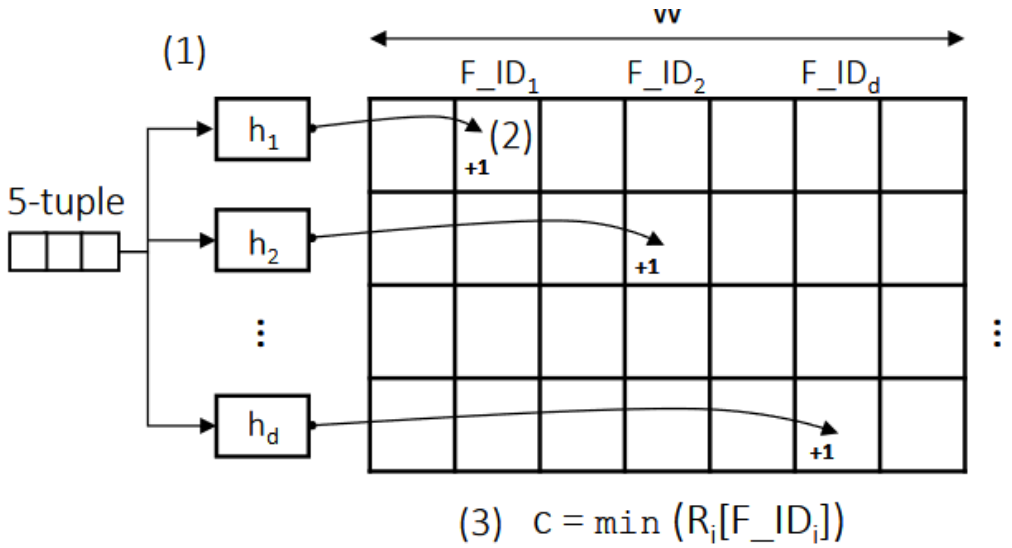

SYN Flood

SYN Flood Attack

- Massive amount of TCP SYN requests with spoofed IP addresses are sent to the server
- These connections consume the server's resources, making it unresponsive to legitimate traffic
- Server start "half-open" connections
- Connections build up until queue is full and all additional requests are blocked

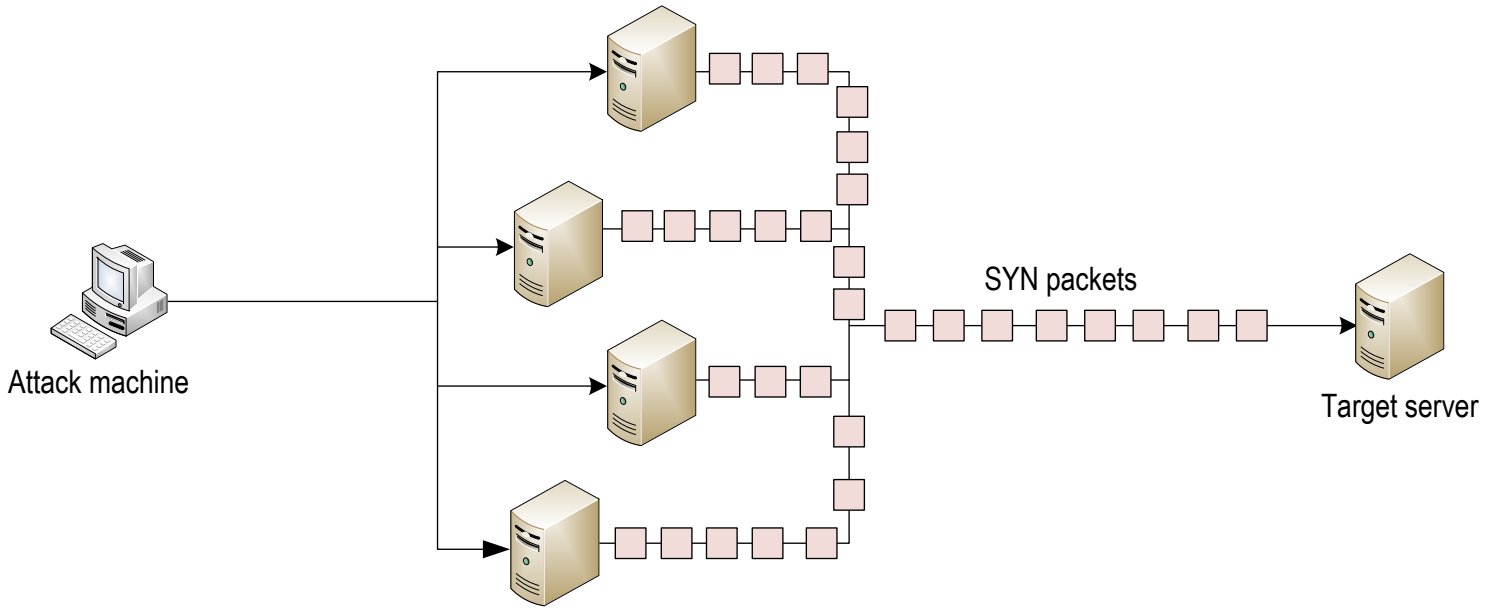

SYN Flood Detection in P4

- Count the number of SYN packets per second in the data plane
- When the count exceeds a predefined threshold, the switch starts dropping SYN packets
- The dropping percentage is configured by the administrator from the control plane