

Heavy Hitter Detection (TUDelft)

Identifying, via the dataplane, the “heavy hitter” flows or flows with large traffic volumes is important for several applications e.g., flow-size-aware routing, traffic engineering, network management applications, link congestion avoidance and detection, to plan network capacity, and to detect network anomalies and attacks (DoS detection, etc.). The problem of finding these flows in a network is a type of “frequent items” problem, i.e. given N flows, find M of those that occur most frequently or which occur more than $x\%$ of the time.

Heavy hitter detection implementations usually depend on some type of packet sampling or are based on specialized data structures called sketches that hash and count all packets in switch hardware. Packet sampling approaches have some well-known limitations, like the trade-off between scalability, overhead, and accuracy.

We would like to investigate whether these limitations can be mitigated within a P4 context. Our objective is to use P4 programmable switches (we plan to use SURFnet's Barefoot switches) to test the performance of various (possibly novel) algorithms for heavy hitter detection that were initially not developed with P4 in mind.

Deliverable: A report on heavy hitter detection in the dataplane, due Dec. 31, 2018