

Design and Development of Volumetric DDoS Detection Strategies in P4-enabled Programmable Switches

Damu Ding

Joint work with: Domenico Siracusa, Marco Savi and Federico Pederzolli

University of Oxford

Coseners
7th July, 2022



Network monitoring functionalities in Software-Defined Networks

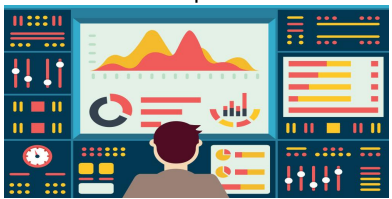

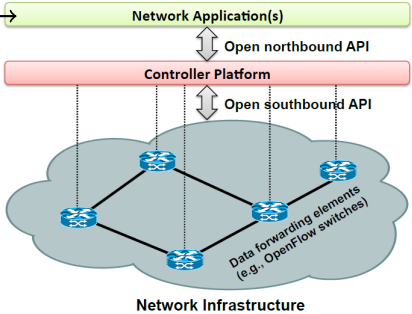

Figure source: Kreutz, Diego, et al. "Software-defined networking: A comprehensive survey." Proceedings of the IEEE 103.1 (2015): 14-76. and https://n0where.net/real-time-network-monitoring-cyberprobe

Network monitoring functionalities in Software-Defined Networks

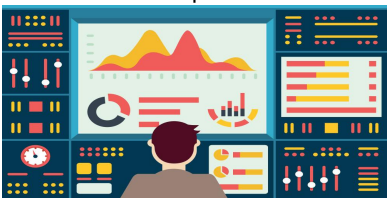

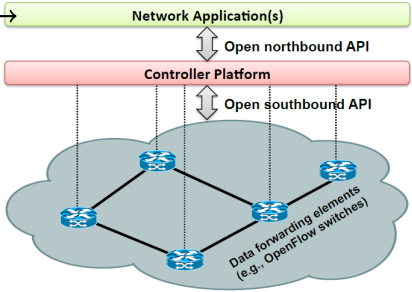

1. Significant communication overhead
2. The latency caused by interaction
3. Cannot perform monitoring at line-rate speed (Up to 100 Gbps)

Figure source: Kreutz, Diego, et al. "Software-defined networking: A comprehensive survey." Proceedings of the IEEE 103.1 (2015): 14-76. and https://n0where.net/real-time-network-monitoring-cyberprobe

Network monitoring functionalities in Software-Defined Networks

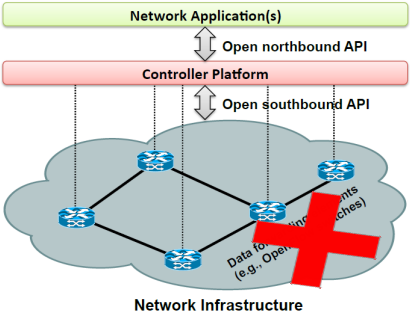

Data plane programmable switches

Figure source: Kreutz, Diego, et al. "Software-defined networking: A comprehensive survey." Proceedings of the IEEE 103.1 (2015): 14-76. and https://n0where.net/real-time-network-monitoring-cyberprobe

Network monitoring functionalities in Software-Defined Networks

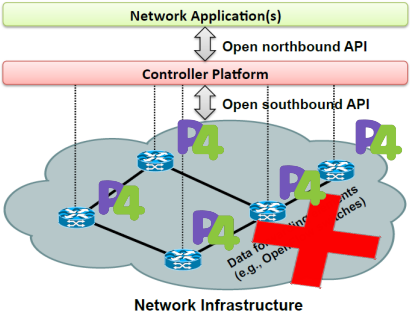

Data plane programmable switches

Figure source: Kreutz, Diego, et al. "Software-defined networking: A comprehensive survey." Proceedings of the IEEE 103.1 (2015): 14-76. and https://n0where.net/real-time-network-monitoring-cyberprobe

Network monitoring functionalities in Software-Defined Networks

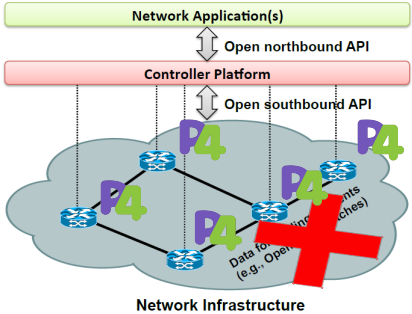

Data plane programmable switches

Figure source: Kreutz, Diego, et al. "Software-defined networking: A comprehensive survey." Proceedings of the IEEE 103.1 (2015): 14-76. and https://n0where.net/real-time-network-monitoring-cyberprobe

Network monitoring functionalities in Software-Defined Networks

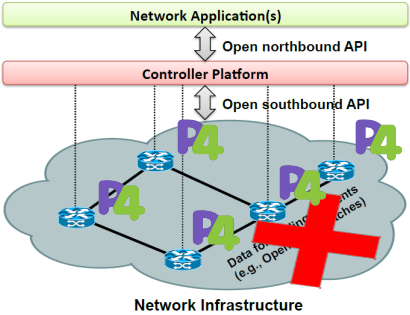

Heavy-hitter detection
Flow cardinality estimation
Network traffic entropy estimation
Traffic volume estimation
Volumetric DDoS detection

Data plane programmable switches

Figure source: Kreutz, Diego, et al. "Software-defined networking: A comprehensive survey." Proceedings of the IEEE 103.1 (2015): 14-76. and https://n0where.net/real-time-network-monitoring-cyberprobe

Network monitoring functionalities in Software-Defined Networks

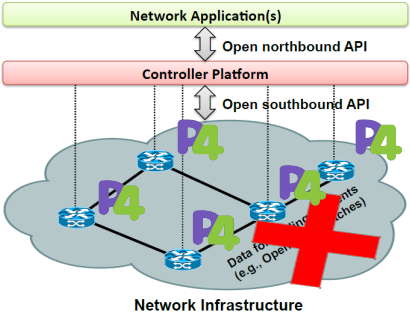

Heavy-hitter detection

Flow cardinality estimation

Network traffic entropy estimation

Traffic volume estimation

Volumetric DDoS detection

Data plane programmable switches

Figure source: Kreutz, Diego, et al. "Software-defined networking: A comprehensive survey." Proceedings of the IEEE 103.1 (2015): 14-76. and https://n0where.net/real-time-network-monitoring-cyberprobe

Damu Ding damu.ding@eng.ox.ac.uk

Challenges

32x 100Gbps QSFP ports

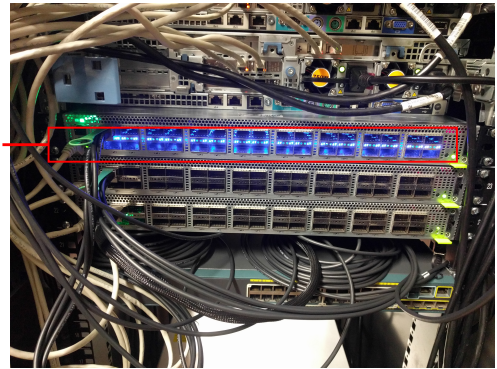

Pros:

1. Higher monitoring throughput

Cons:

1. Limited hardware resources
2. Computational constraints

Figure: Edgecore Tofino switch

Volumteric DDoS detection strategies in literature cannot be directly of-floaded to programmable switch data plane

Motivation

Goal

Design and develop new strategies for volumteric DDoS detection in P4-enabled programmable data planes considering the switch constraints

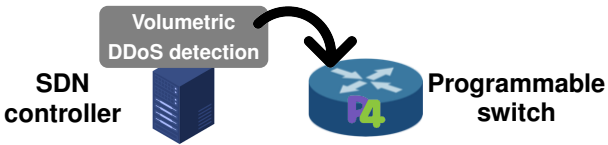

Motivation

Goal

Design and develop new strategies for volumteric DDoS detection in P4-enabled programmable data planes considering the switch constraints

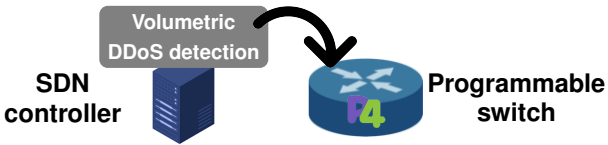

Focus on
ISP networks

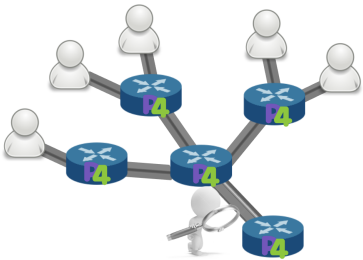

- Minimize out-of-band actions
- Good network performance
- High detection accuracy

Normalized network traffic entropy-based DDoS detection

Normalized network traffic entropy

Normalized network traffic entropy H_{norm} indicates network traffic distribution

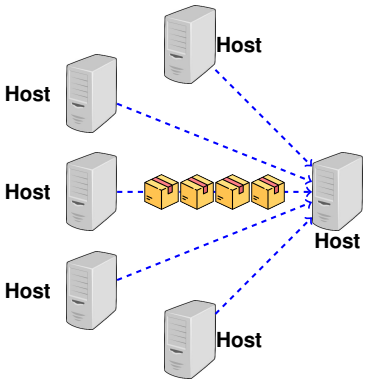

$H_{norm} = 1$

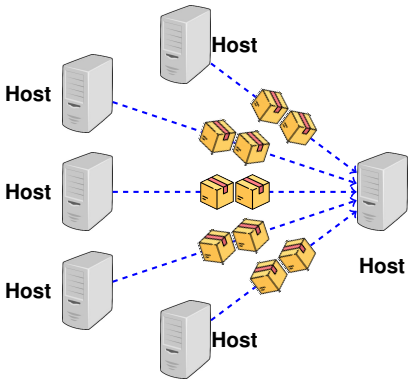

$H_{norm} = 0$

Normalized network traffic entropy in programmable switches

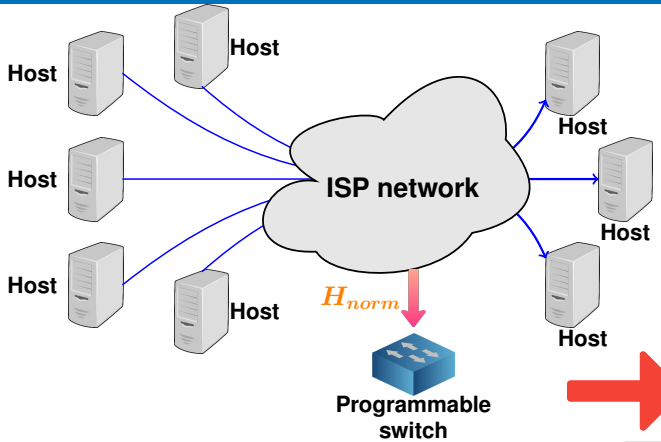

I can know the **overall condition** of high-throughput network

Network operator

Normalized Shannon entropy

$$H_{norm} = \frac{-\sum_{i=1}^{n_{tot}} \frac{f_i}{S_{tot}} \log_2 \frac{f_i}{S_{tot}}}{\log_2 n_{tot}}$$

- f_i: Packet count of flow i
- S_{tot}: Overall number of packets ✓
- n_{tot}: Overall number of flows

Sketch

Sketch is a fast and memory-efficient data structure to store flow statistics

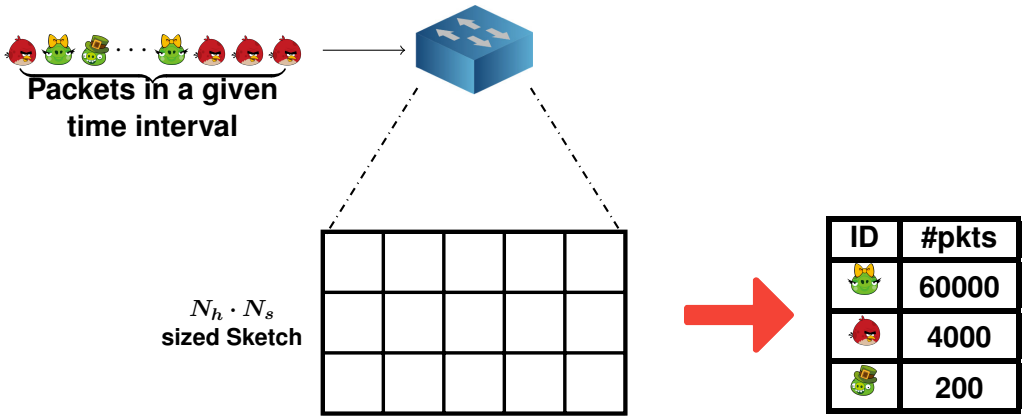

N_h : Number of hash functions, N_s : Output size of hash functions

Normalized network traffic entropy-based DDoS detection

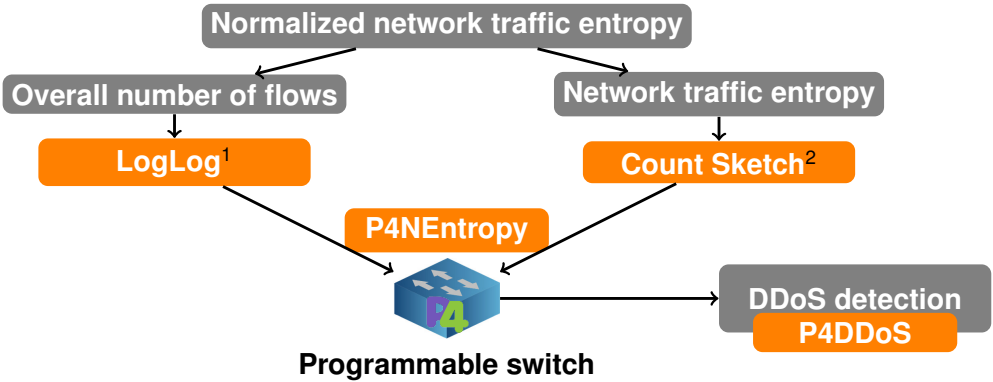

[1] Durand, Marianne et al. "Loglog counting of large cardinalities." European Symposium on Algorithms. Springer, Berlin, Heidelberg, 2003.

[2] M. Charikar et al, "Finding frequent items in data streams," in Springer International Colloquium on Automata,Languages, and Programming (ICALP), 2002.

Property of volumetric DDoS attacks

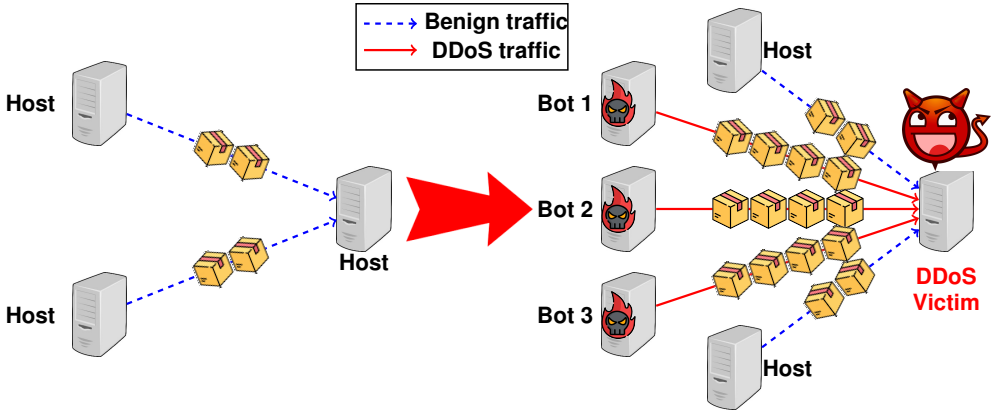

Adaptive threshold for DDoS detection

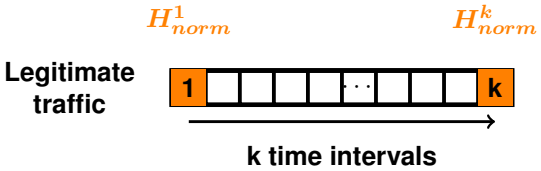

Exponentially weighted moving average (EWMA) of normalized entropy

$$EWMA^1_{norm} = H^1_{norm}$$
$$EWMA^2_{norm} = \alpha H^1_{norm} + (1-\alpha) H^2_{norm}$$
$$\cdots$$
$$EWMA^k_{norm} = \alpha H^{k-1}_{norm} + (1-\alpha) H^k_{norm}$$

DDoS threshold

$$\lambda^k_{norm} = EWMA^k_{norm} - \epsilon$$

Damu Ding, Marco Savi, and Domenico Siracusa. *Tracking Normalized Network Traffic Entropy to Detect DDoS Attacks in P4* IEEE Transactions on Dependable and Secure Computing (TDSC).

Per-flow cardinality-based DDoS detection

Property of volumetric DDoS attacks

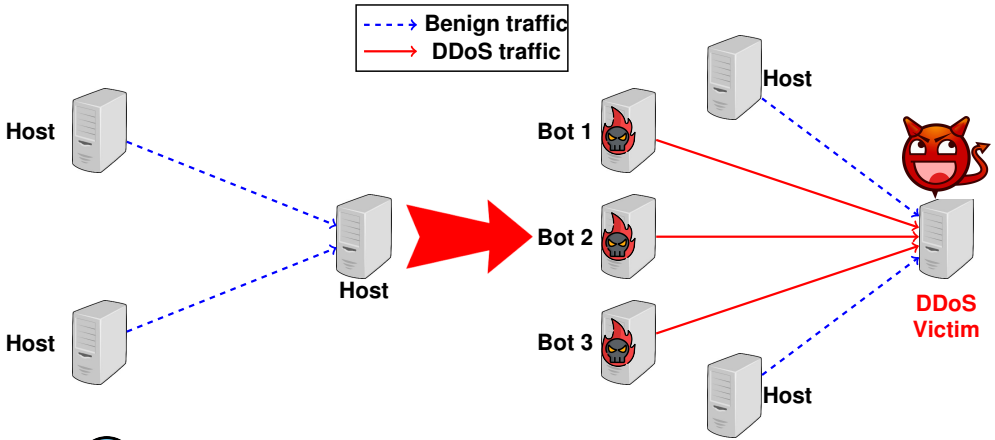

Threat model and deployment scenario

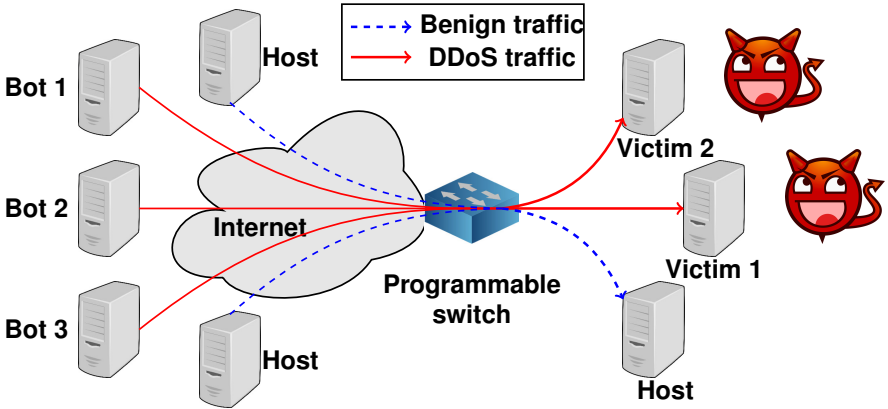

A memory-efficient data structure to count the number of flows targeting different destinations in the programmable switch is necessary

BACON Sketch

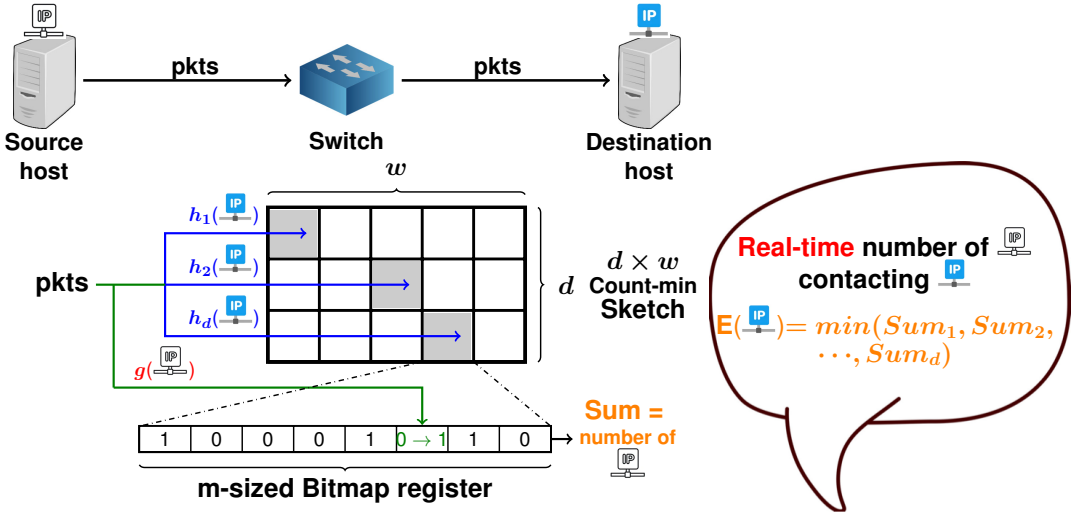

In-network DDoS victim identification (INDDoS)

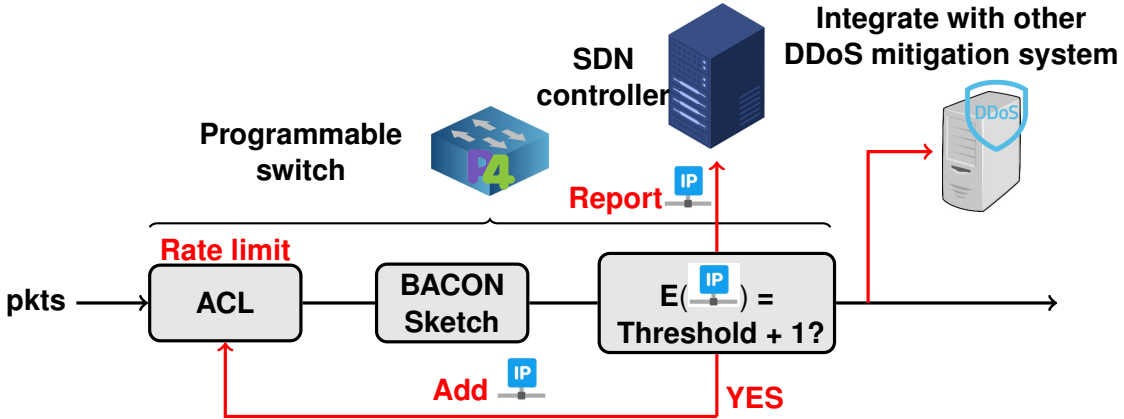

Damu Ding, Marco Savi, Federico Pederzolli, Mauro Campanella, and Domenico Siracusa. *In-Network Volumetric DDoS Victim Identification Using Programmable Commodity Switches* IEEE Transactions on Network and Service Management (TNSM).

Results

Algorithm	False-positive rate D_{fp}	True-positive rate D_{tp} / Detection accuracy D_{acc}				
		Booter 6	Booter 7	Booter 1	Booter 4	Mixed
P4DDoS	8%	100% / 96%	82% / 87%	96% / 94%	98% / 95%	100% / 96%
SOTA_DDoS [3]	10%	100% / 95%	74% / 82%	100% / 95%	94% / 92%	100% / 95%

Algorithm	Recall	Precision	F1 score
INDDoS	**0.96**	**0.99**	**0.97**
Spread Sketch[4]	**0.92**	**0.94**	**0.93**

DDoS attack flow trace	Recall	Precision	F1 score
Booter 6	1.0 (1/1)	1.0 (1/1)	1.0 (1/1)
Booter 7	1.0 (1/1)	1.0 (1/1)	1.0 (1/1)
Booter 1	1.0 (1/1)	1.0 (1/1)	1.0 (1/1)
Booter 4	1.0 (1/1)	1.0 (1/1)	1.0 (1/1)
Mixed	1.0 (4/4)	1.0 (4/4)	1.0 (4/4)

Both **P4DDoS** and **INDDoS** can be entirely executed in P4 programmable switches

[3] Lapolli, Angelo Cardoso, Jonatas Adilson Marques, and Luciano Paschoal Gaspary. "Offloading real-time ddos attack detection to programmable data planes." 2019 IFIP/IEEE Symposium on Integrated Network and Service Management (IM). IEEE, 2019.

[4] Tang, Lu, Qun Huang, and Patrick PC Lee. "Spreadsketch: Toward invertible and network-wide detection of superspreaders." IEEE INFOCOM 2020-IEEE Conference on Computer Communications. IEEE, 2020.

Conclusion

- Offload volumetric DDoS detection to programmable switches leveraging sketches
  - Memory efficient
  - Accurate estimation
  - Fast
- Two different volumetric DDoS detection strategies in programmable data planes
  - Normalized entropy-based
  - Per-flow cardinality-based
- Proved DDoS detection performance using programmable switches
  - High DDoS detection accuracy
  - Low packet processing time for detection
  - Valid for high-throughput networks

Thank you!

damu.ding@eng.ox.ac.uk