Network Measurements

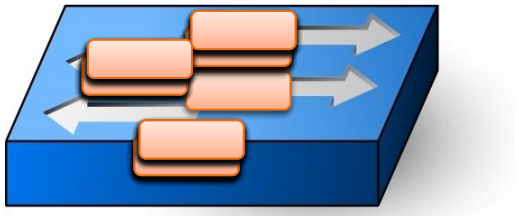

Network Measurements

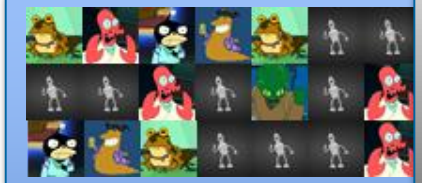

Elephant Flow

Load Balancing
Traffic Engineer
C

Sliding Windows Statistics

Link Utilization
Trend Detection

tinct Elements

ion

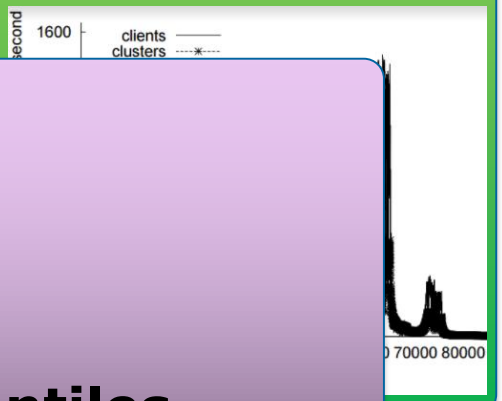

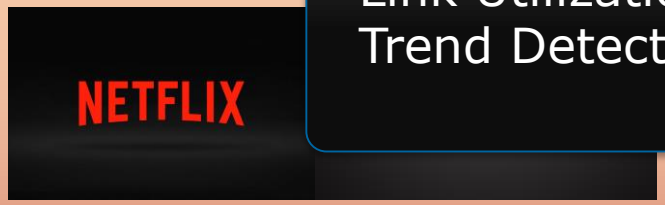

Estimating the fraction of rare flows

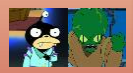

Customer Satisfaction
DDoS Detection

Computing Quantiles

Data Log Analysis
Network Health Monitoring

Heavy Hitters

- How many packets has  sent?
- Which flows are larger than T?
- Traditionally – must fit in the SRAM

Year	2012	2014	2016
SRAM (MB)	10-20	30-60	50-100

(SilkRoad, SIGCOMM 2017)

Can't allocate a counter for each flow!

Distributed Denial of Service

CNN tech BUSINESS CULTURE GADGETS FUTURE STARTUPS

Massive cyberattack turned ordinary devices into weapons

by Samuel Burke @CNNTech

October 22, 2016: 10:37 AM ET

Recommend 4.2K

Major DDoS attacks see huge increase, says Akamai

Alex Scroxton
Networking Editor
14 Feb 2017 12:40

Akamai's State of the Internet/Security Report for the fourth quarter of 2016 finds that distributed denial of service attacks larger than 100Gbps are rapidly increasing as more IoT devices are compromised

Hierarchical Heavy Hitters (HHH)

Hierarchical Heavy Hitters identify traffic clusters.

They are at the core of numerous DDoS mitigation systems…

DDoS attack (Aug. 2014)

DREAM: dynamic resource allocation for software-defined Counting.
ACM SIGCOMM 2014

LADS: Large-scale Automated DDoS Detection System.
USENIX ATC 2006

Automatically Inferring Patterns of Resource Consumption in Network Traffic.
ACM SIGCOMM 2003

DDoS Mitigation

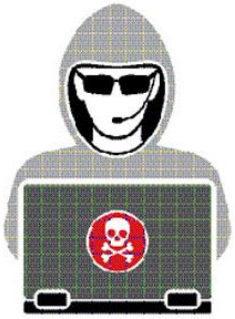

181.7.20.1
181.7.20.2
…
181.7.21.1
181.7.21.2
…

Can we block only the attacking devices?

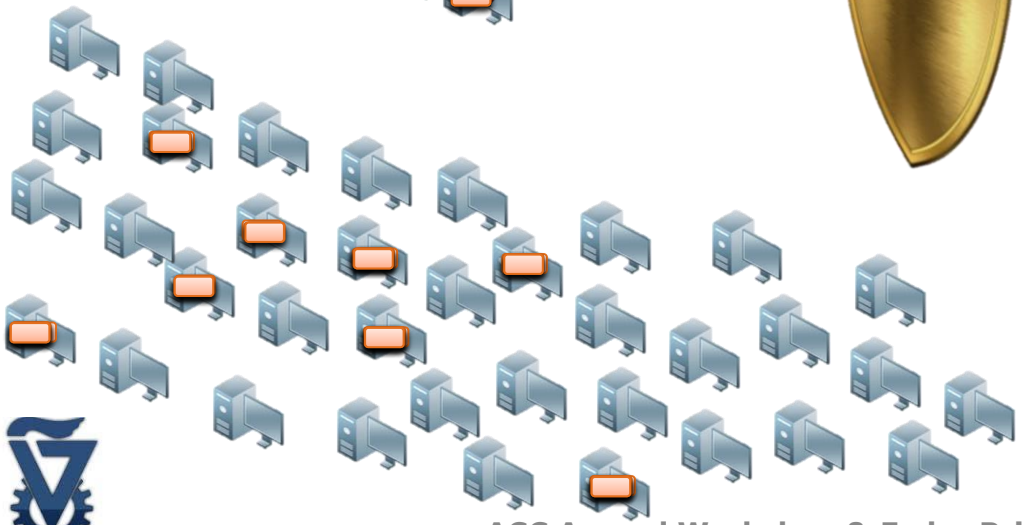

Hierarchical Heavy Hitters

Hierarchical Heavy Hitters identifies frequent:

- Flows (*heavy hitters)*
- Source networks.
- Source-Destination pairs.

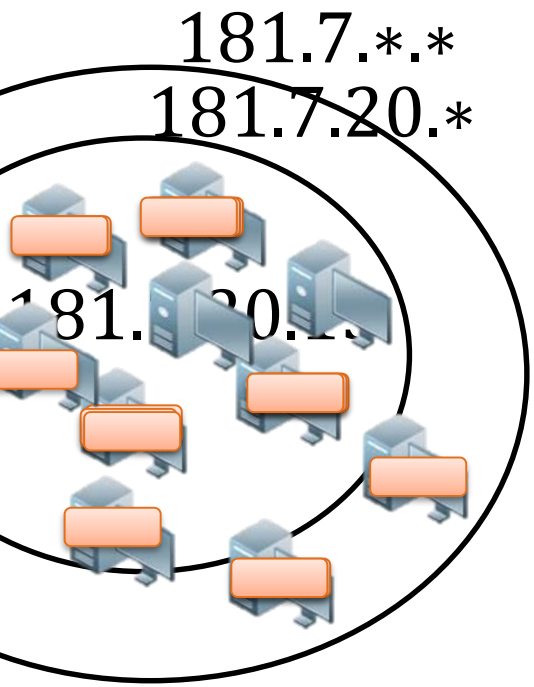

State of the art

"Count each **prefix** independently."

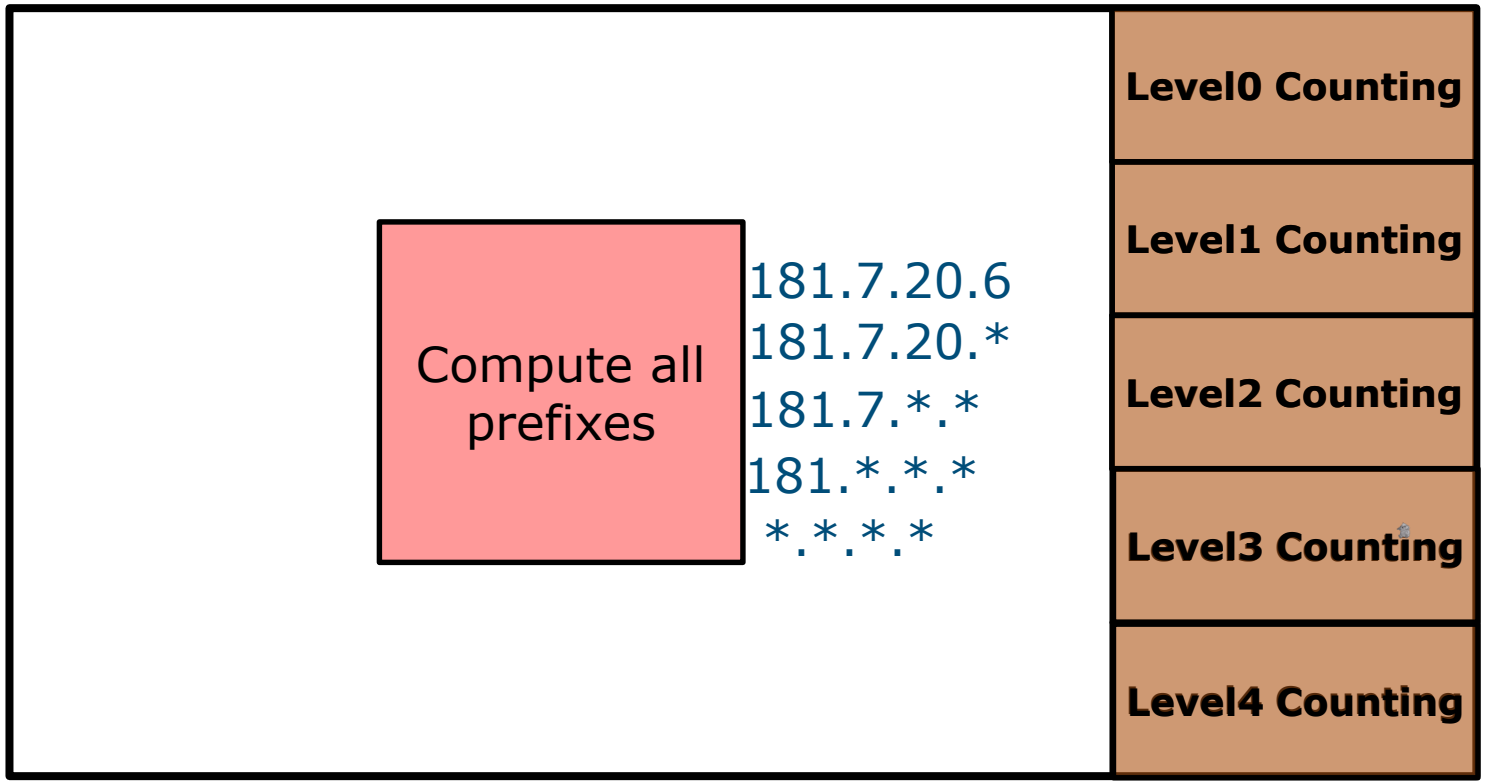

Mitzenmacher et al., Hierarchical Heavy Hitters with the Space Saving Algorithm, ALENEX 2012

Randomized HHH (Our work)

"Select a prefix at random and count it"

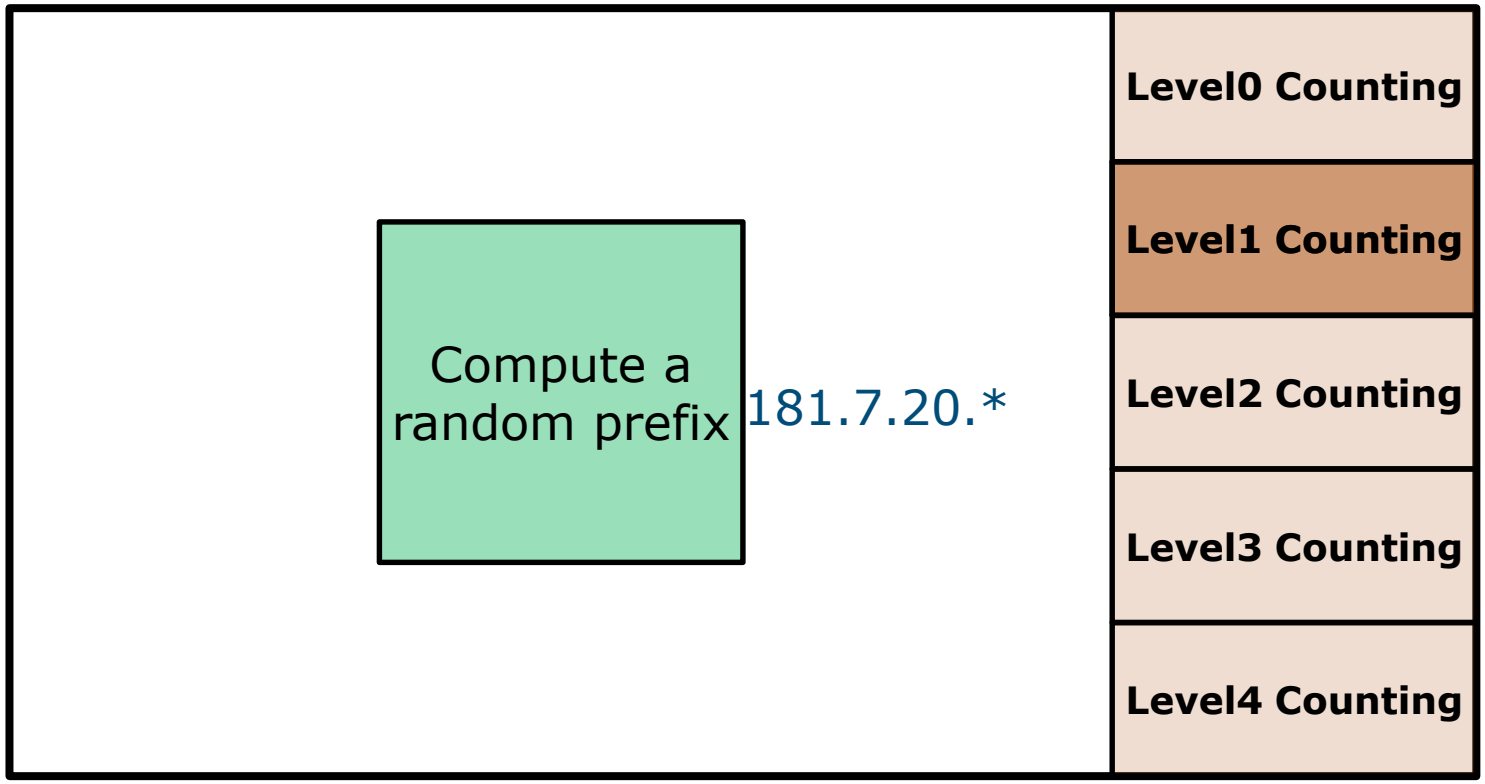

Additional Speedup

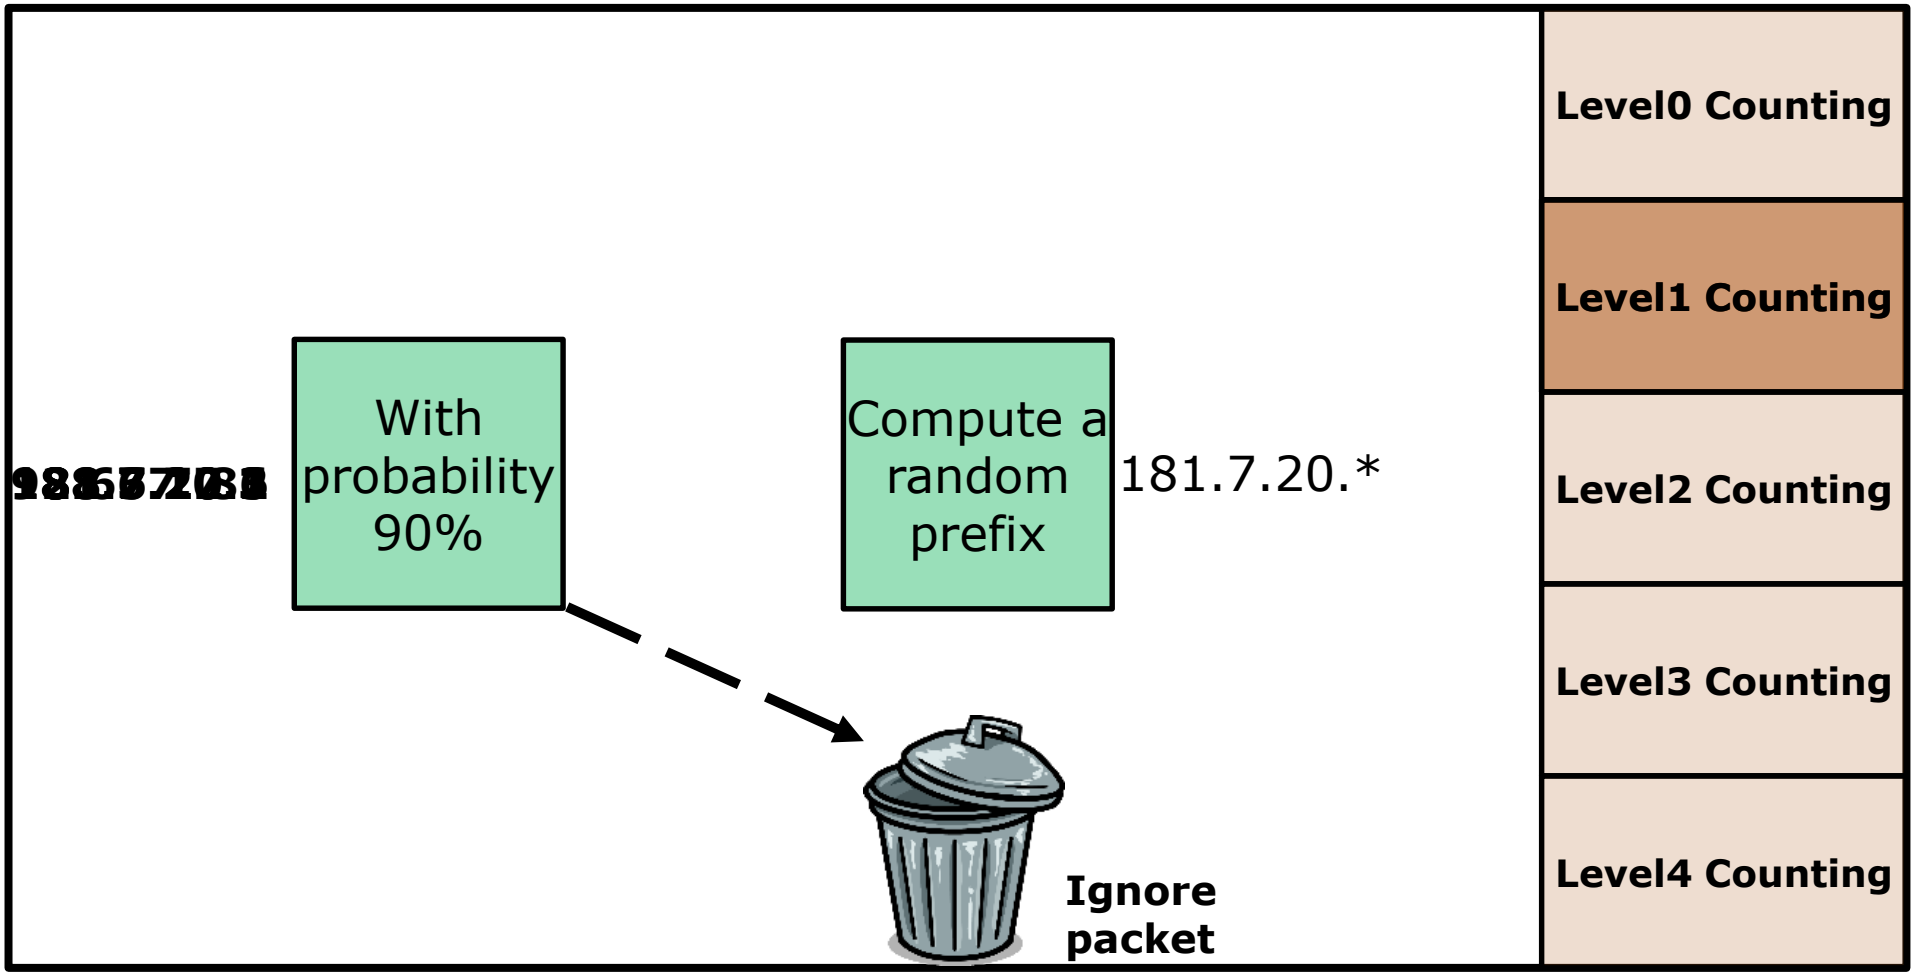

We did the math

Accuracy and convergence guarantees .

After enough packets there are:

1. No false negatives.
2. No counting errors.
3. Only a few false positives.

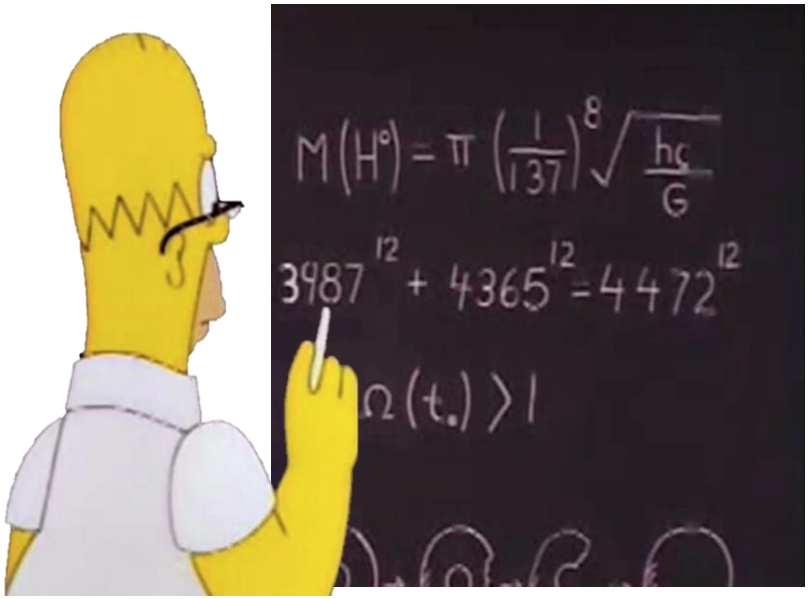

How much traffic is needed for convergence?

"Accuracy improves with the number of packets"

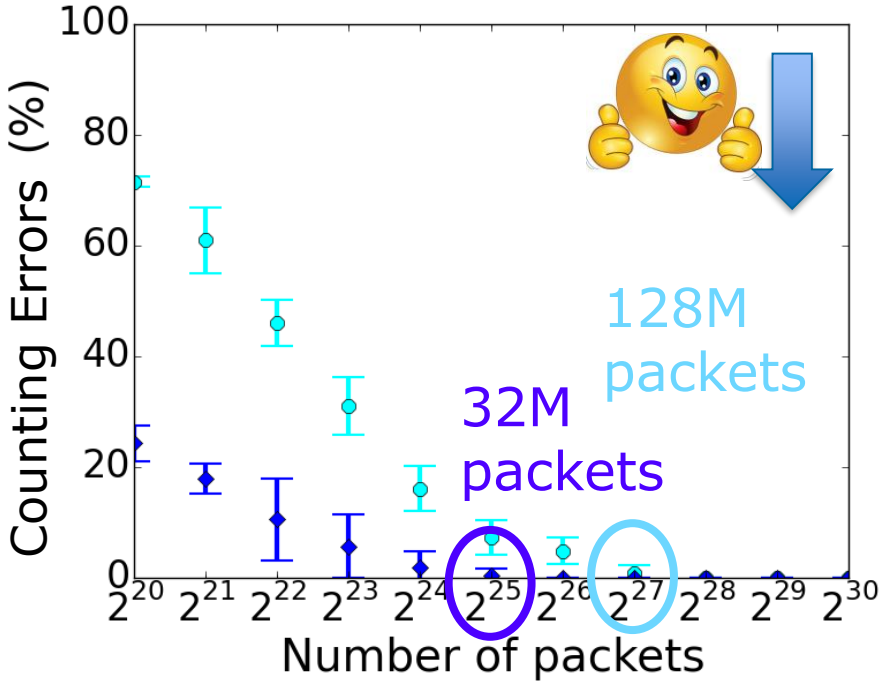

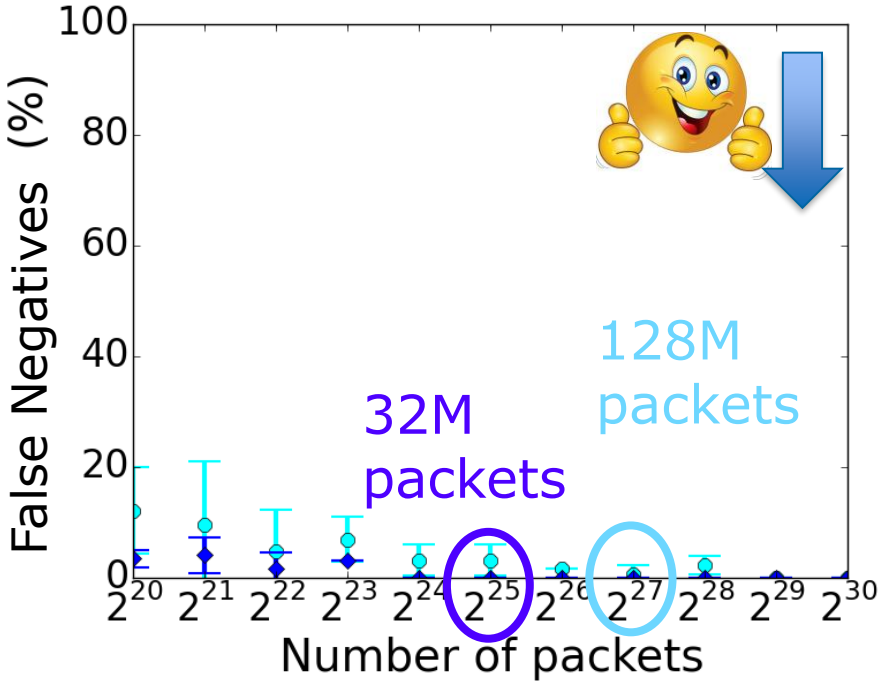

Comparison with other HHH algorithms

"Accuracy improves with the number of packets"

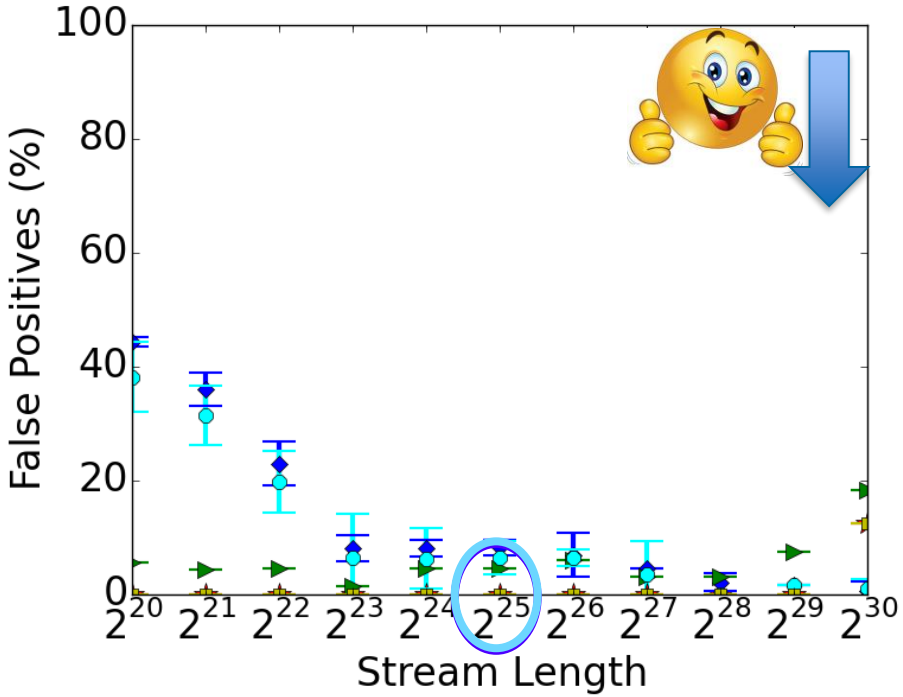

One prefix per packet

One prefix per 10 packets

Partial Ancestry

Full Ancestry

Mitzenmacher et al.

Cormode et al., Finding hierarchical heavy hitters in streaming data, TKDD 2008

Virtual Switching

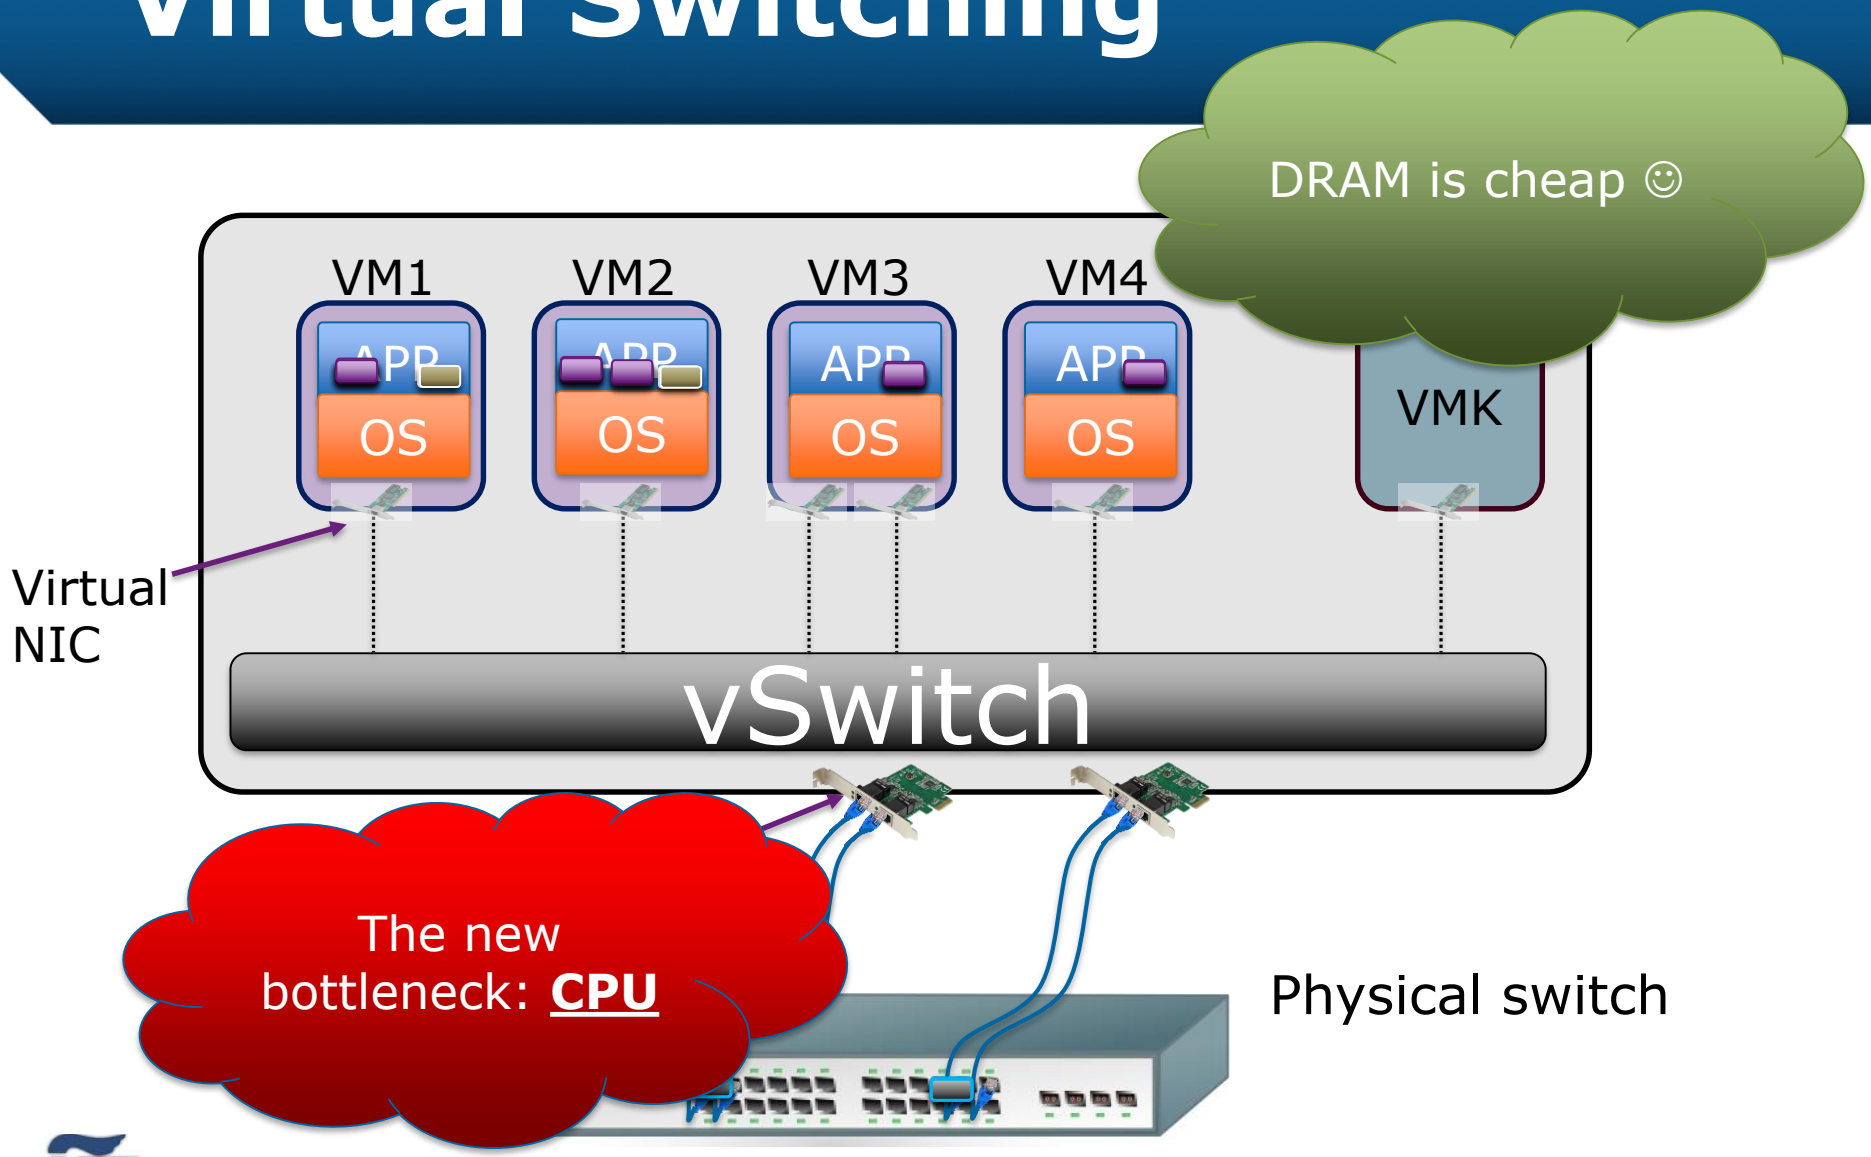

Open vSwitch Implementation

- Server A: Traffic Generator
  - We send min-sized packets with headers from Internet traces.
- Server B: DPDK enabled Open vSwitch
  - Performs HHH Counting in data plane

Traffic Generator

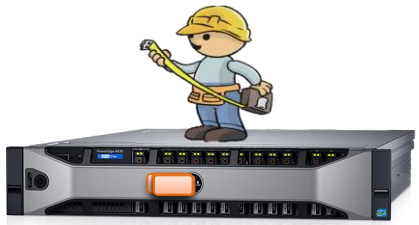

Open vSwitch

Comparing Implementation Overhead

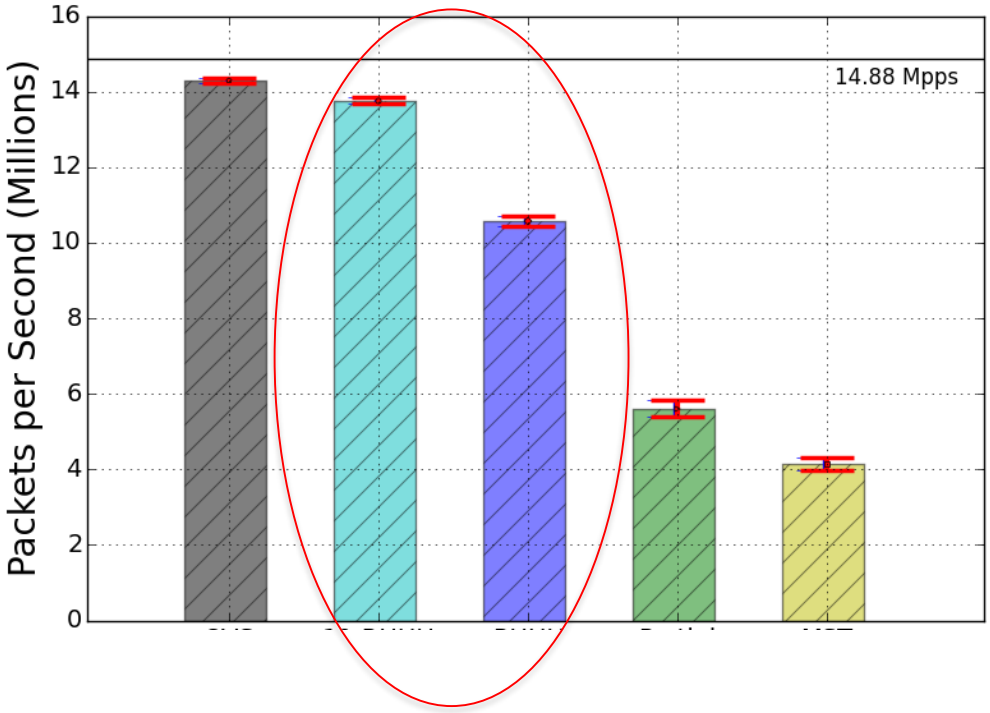

Highlights:

- Only –4% overheads for HHH in the OVS data plane!
- +250% throughput improvement compared to previous work.

One prefix per packet | One prefix per 10 packets | OVS

Partial Ancestry | Mitzenmacher et al.

Takeaways

- Real time hierarchical heavy hitters measurement in networking devices.

- Provable accuracy guarantees.

- Open source code: https://github.com/ranbenbasat/RHHH

Limitations and current projects

- Support for weights
- Support for sliding windows
  - No convergence time!
- Allowing time-based queries
  - "What are the HHH for Jan 20th 2018, 4PM-5PM?"

Any Questions