UNIVERSITY OF CAMBRIDGE

Hidden Hierarchical Heavy Hitters (H4)

Salvator Galea, Gianni Antichi, Roberto Bifulco, Giuseppe Bianchi, Andrew W. Moore

Department of Computer Science and Technology

NEC \Orchestrating a brighter world

Network Management

Monitoring Network Traffic

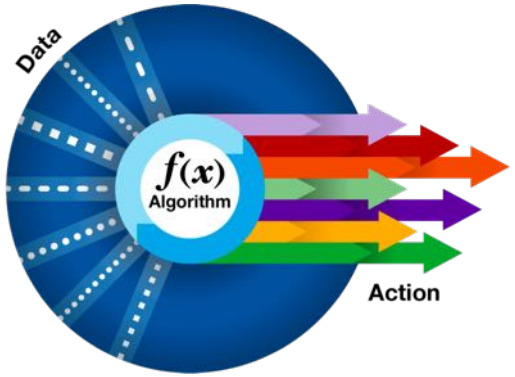

Management Techniques

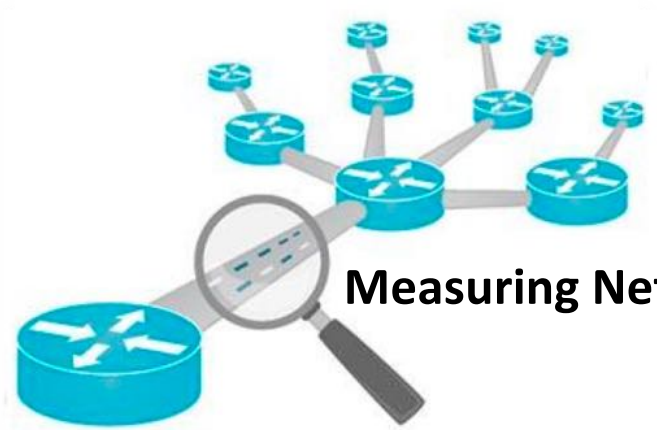

Measuring Network Traffic

Measurements as a support for Management

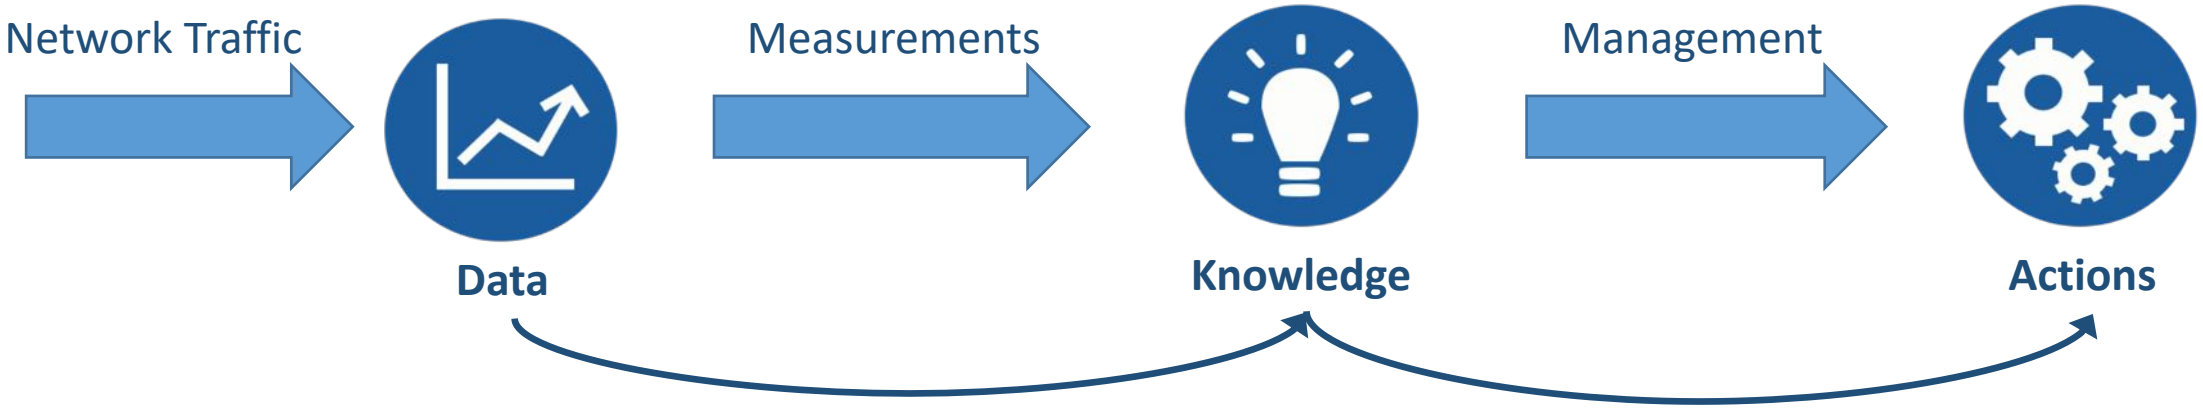

Network Measurements

- Network state
- Traffic demands
- Performance

Network Management

- Traffic Engineering techniques
- Quality of Service
- Security

Traffic characterization

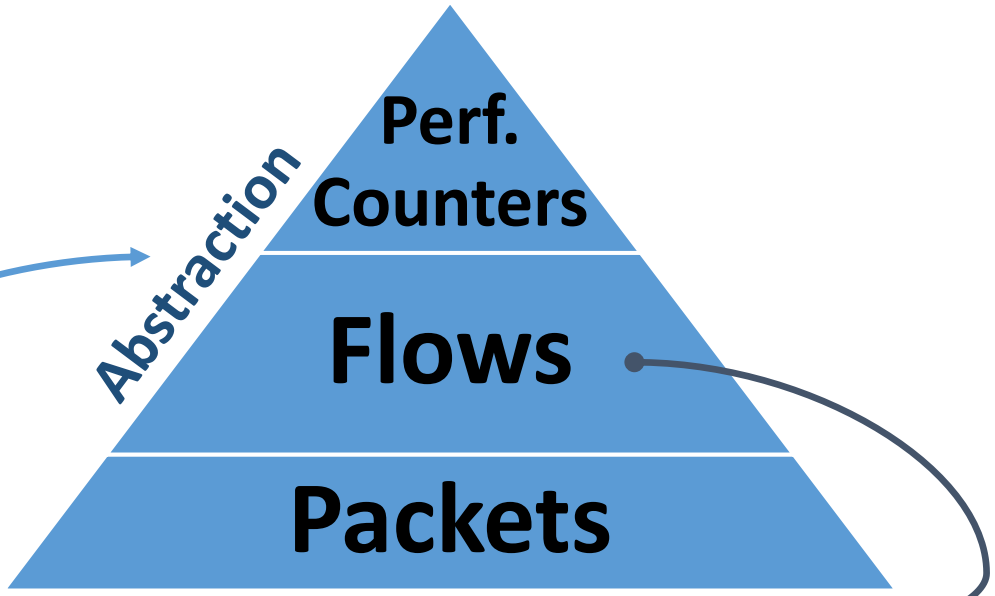

Types of Traffic:

- Bursty (Heavy Flows)
- Latency-sensitive
- Intrusion traffic-pattern (Heavy Flows)

Traffic engineering
Accounting

Security Measure

Collecting data

Data Stream

Sampling

Aggregation

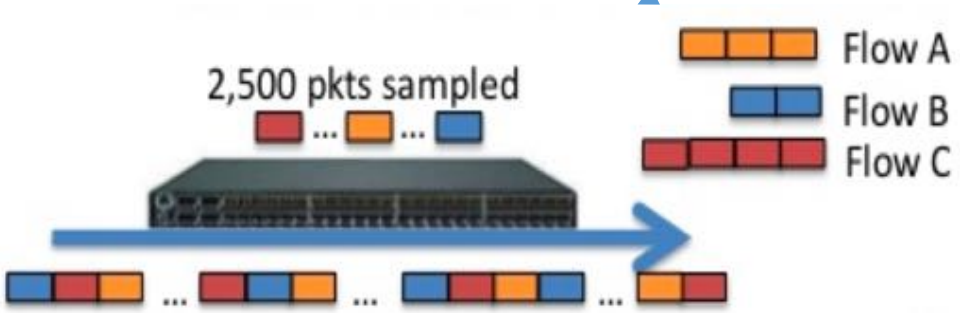

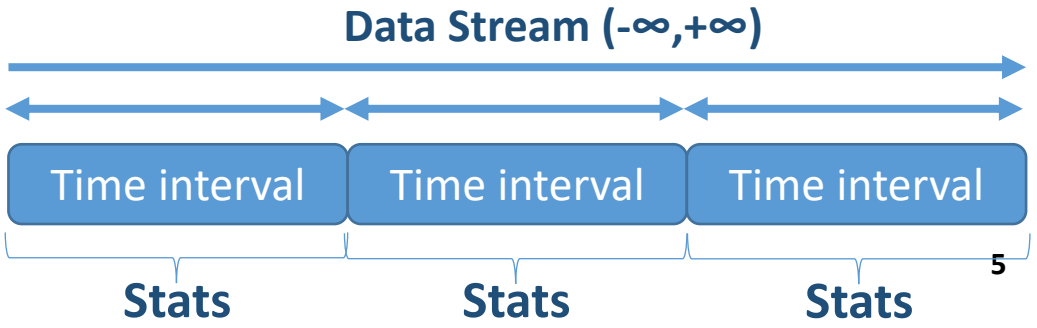

Aggregation + Streaming Algorithms

Aggregation method applied by Streaming algorithm

- Easy to implement
- Prevent counters overflow by flushing
- Stats. In Probabilistic data structures

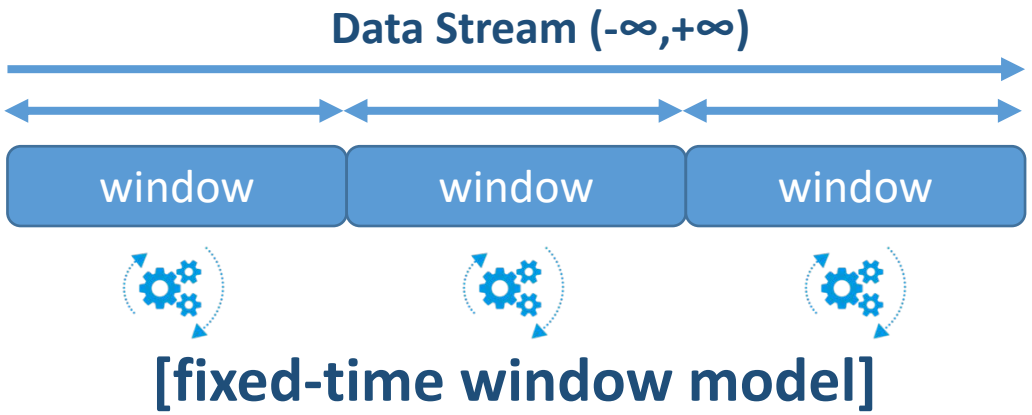

[fixed-time window model]

So far so good ehm, so what's the **pitfall**?

- Discard statistics at the end of each window
- Coupling between statistics and window size

Window Size

Is the window size a problem?

What's the "right" window size?

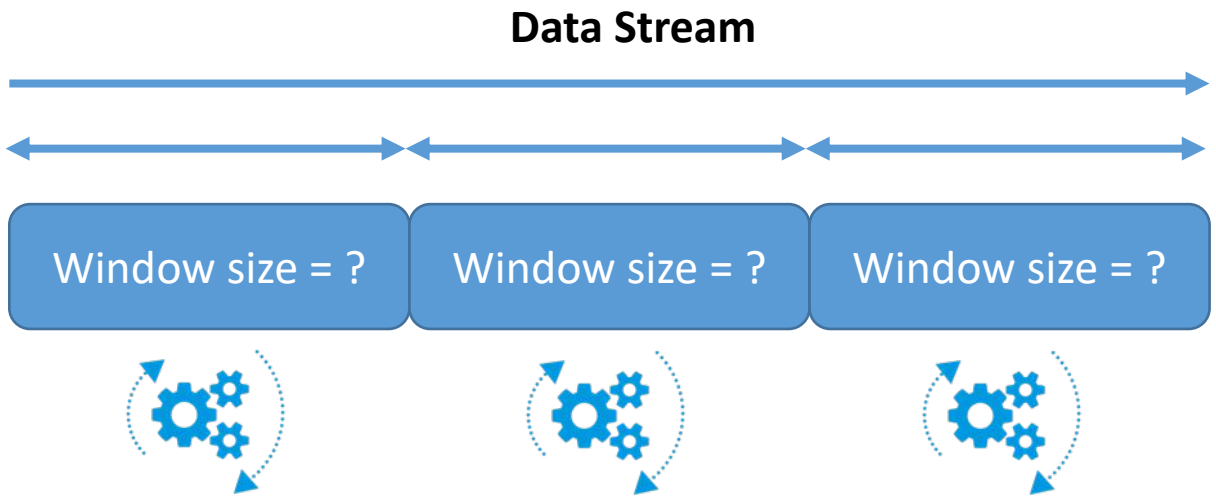

Question: Can we quantify the differences in the aggregates based on the windows characteristics and window model?

Offline experiments

Traces : CAIDA2016 DirA
Prefixes : Source IPs
Baseline Win. : 10sec window
Threshold : % of the total traffic in the window
Comparison metric : Jaccard similarity coefficient
Detection : Heavy Hitters, Hierarchical Heavy Hitter, Leaf Heavy Hitters

	Threshold	Detection	Trace	Windows Comparison (sec)
Test 1	5%	HH,HHH	10 min	[10] to [9.9, 9.8, 9.7]
Test 2	5%	HH,HHH	20 min	[10] to [9.99, 9.96, 9.93, 9.90]
Test 3	1%, Top-50	LHH	60 min	[10] to [10]+offset[1, 2, 3, 4]
Test 4	1%, 5%, 10%	HHH	60 min	[20], [10], [5]

Experiment 1 (HH + HHH), 100msec

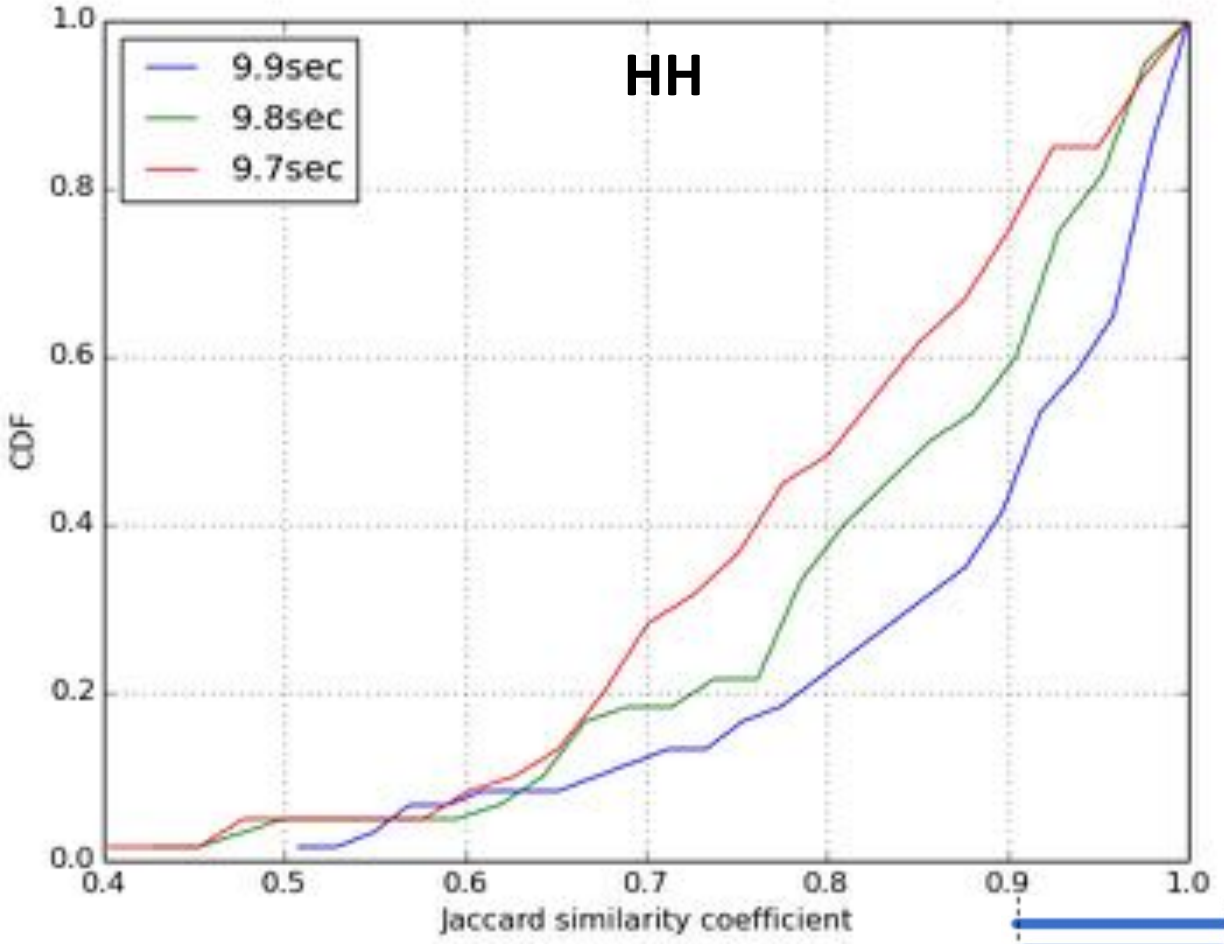

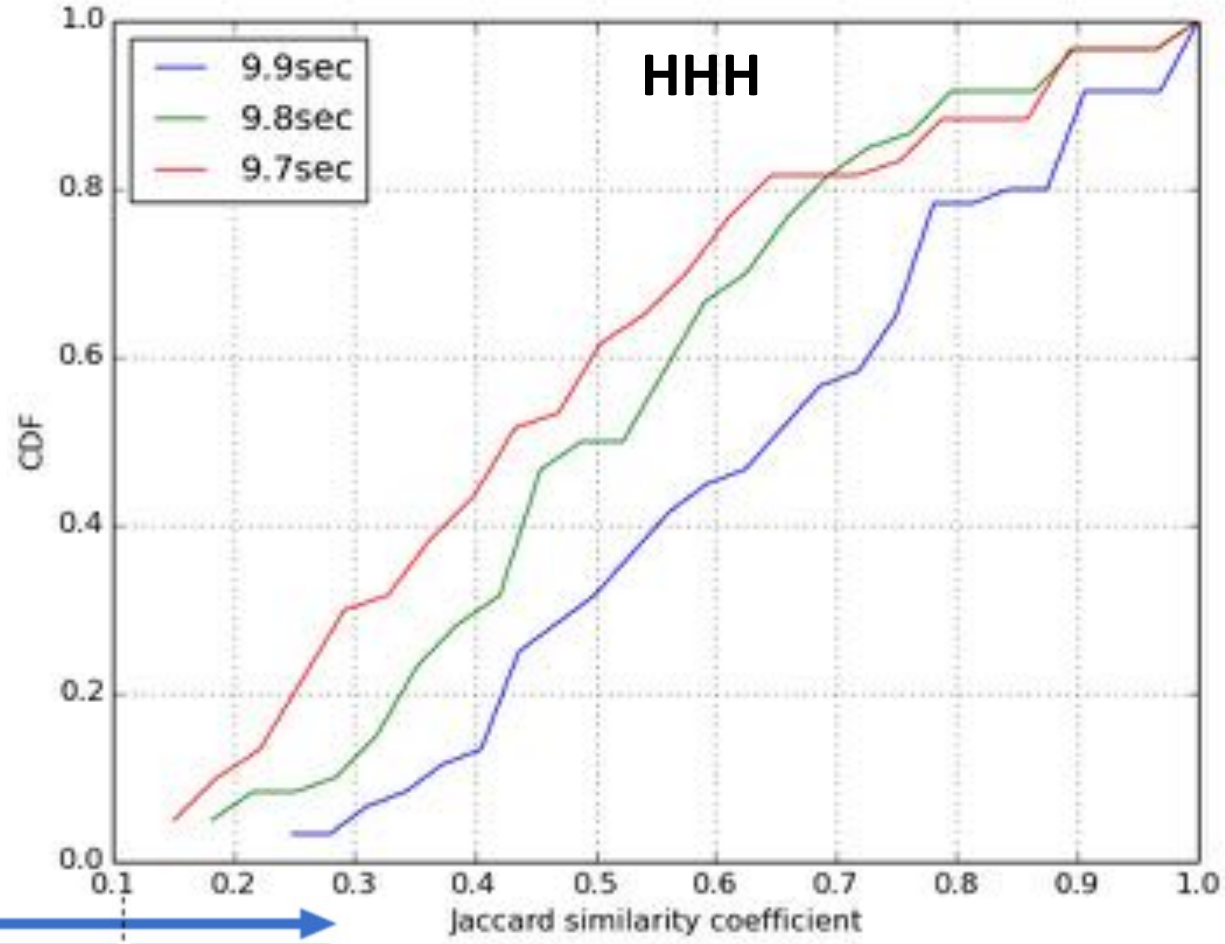

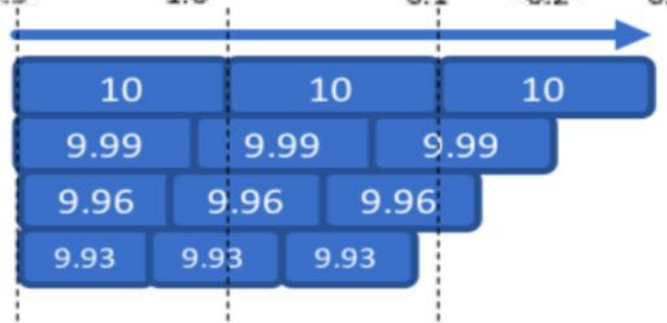

Thr	Detection	Trace(min)	Windows (sec)
5%	HH,HHH	10	[10]-[9.9, 9.8, 9.7]

Experiment 2 (HH + HHH), 10msec

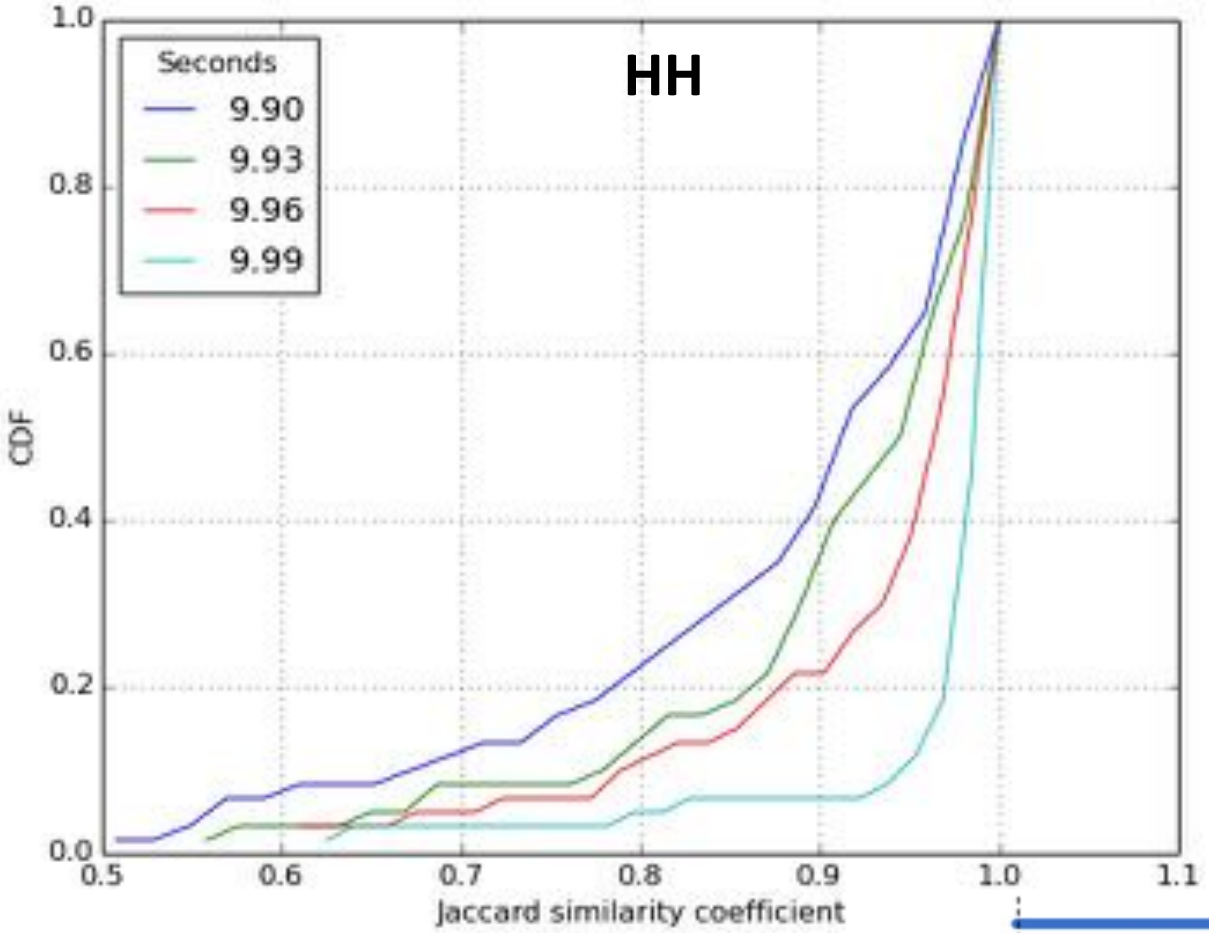

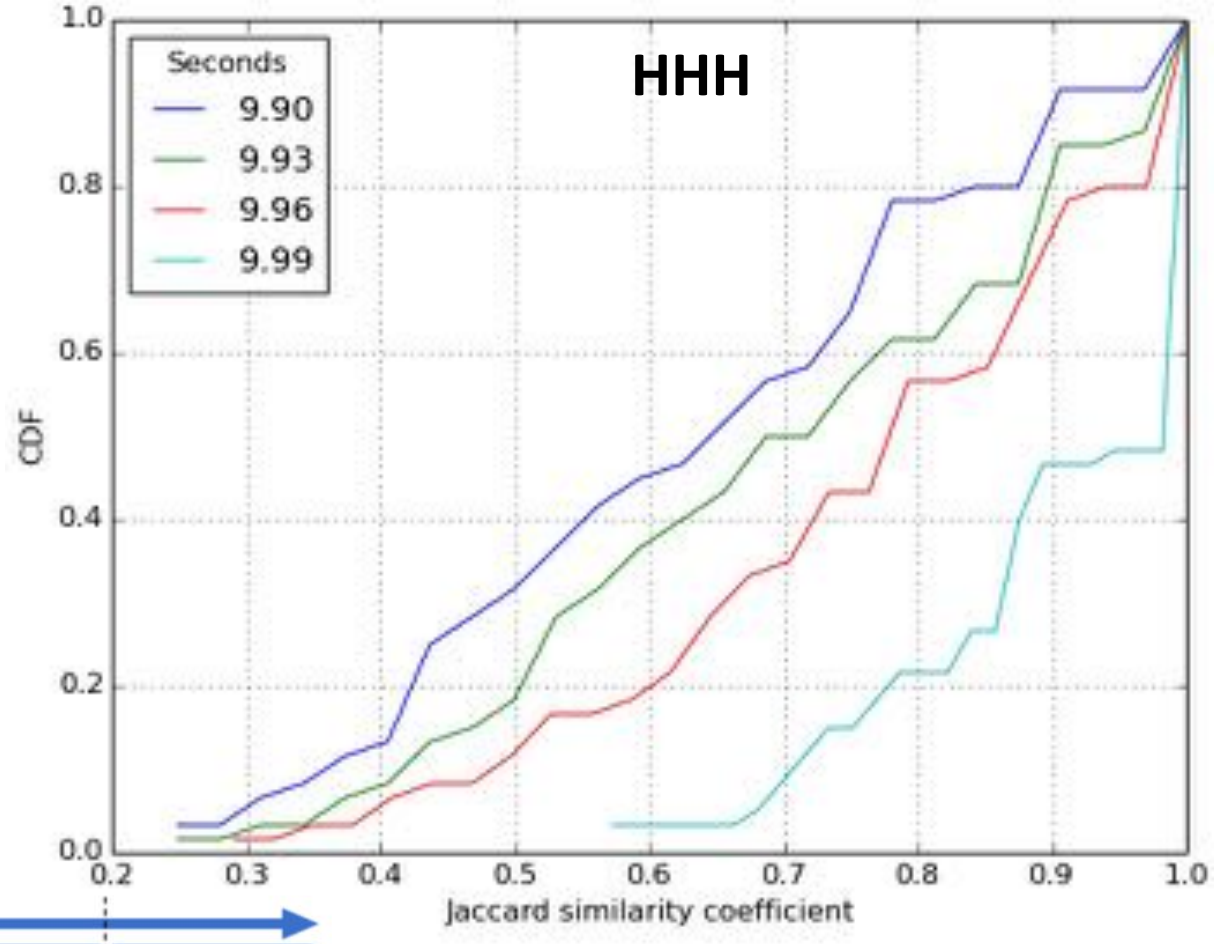

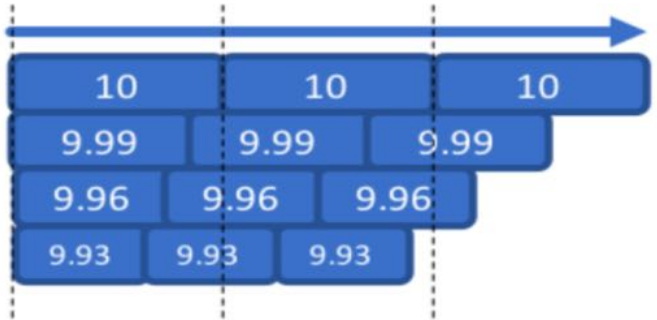

Thr	Detection	Trace(min)	Windows (sec)
5%	HH,HHH	20	[10]-[9.90, 9.93, 9.96, 9.99]

Experiment 3 (LeafHH)

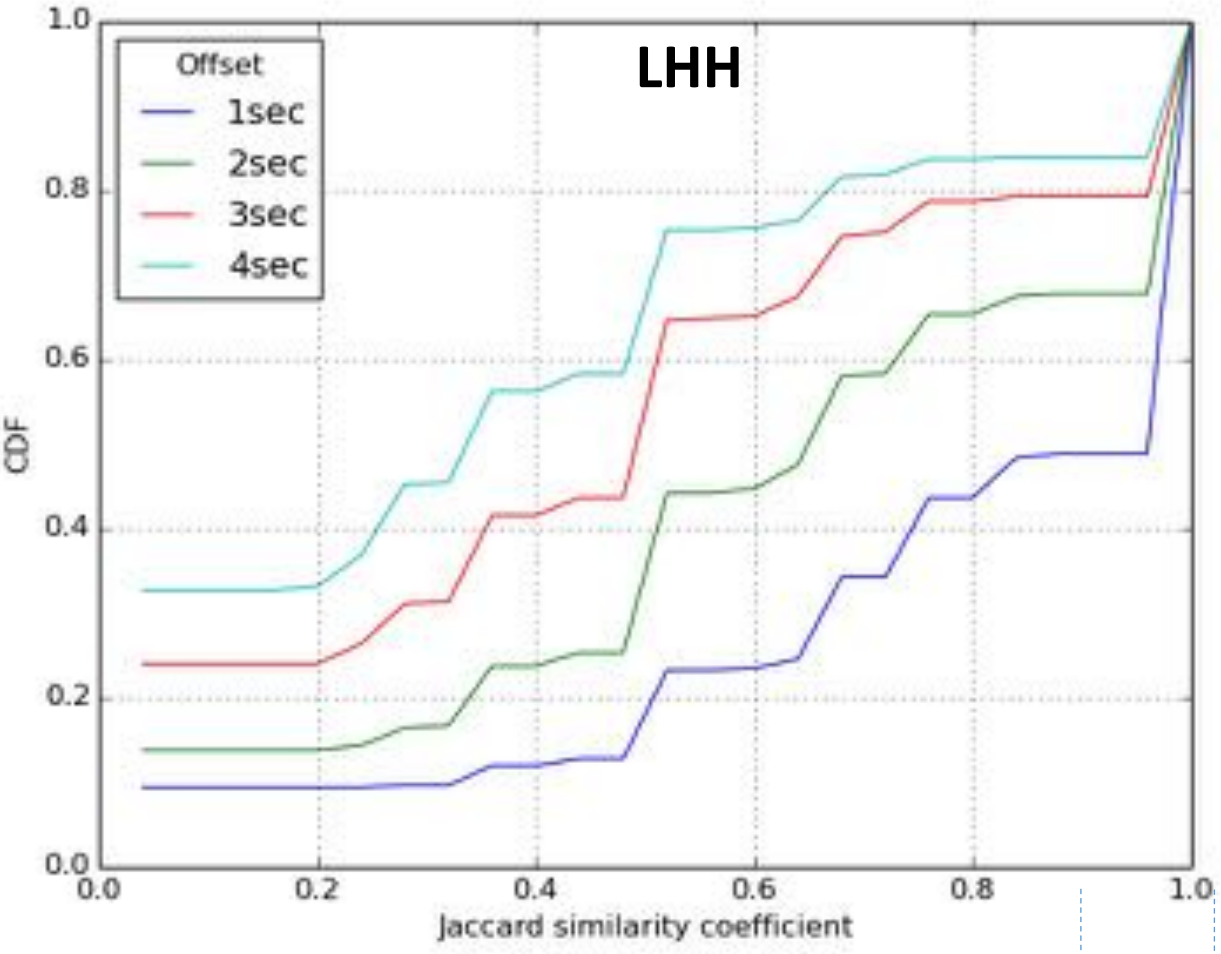

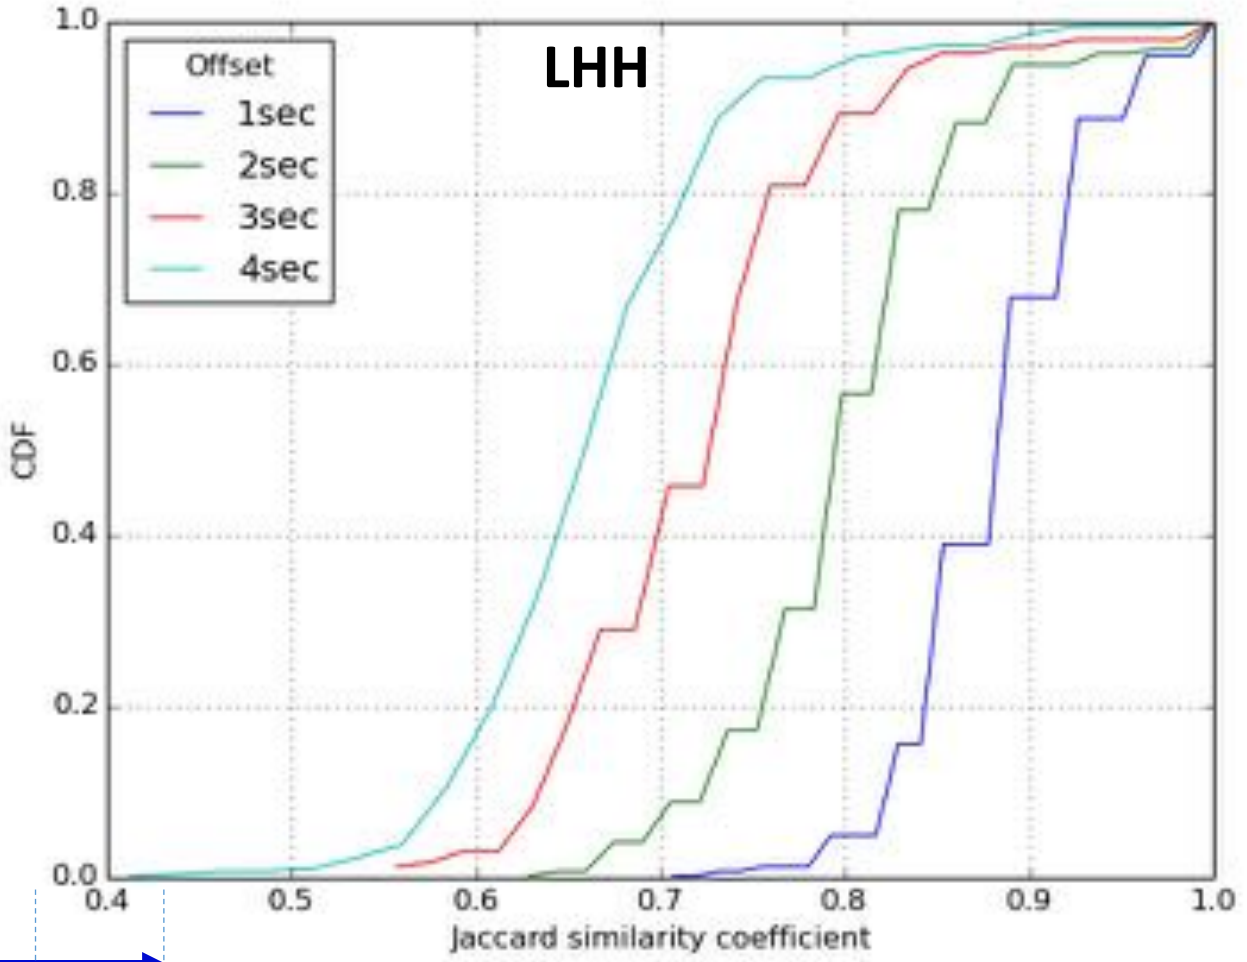

Thr	Detection	Trace	Windows (sec)
1%, Top-50	LHH	60 min	[10]-[9, 8, 7, 6]

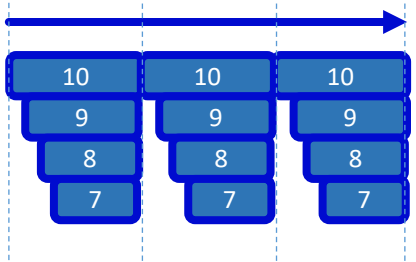

Experiment 4 (HHH, fixed vs sliding window)

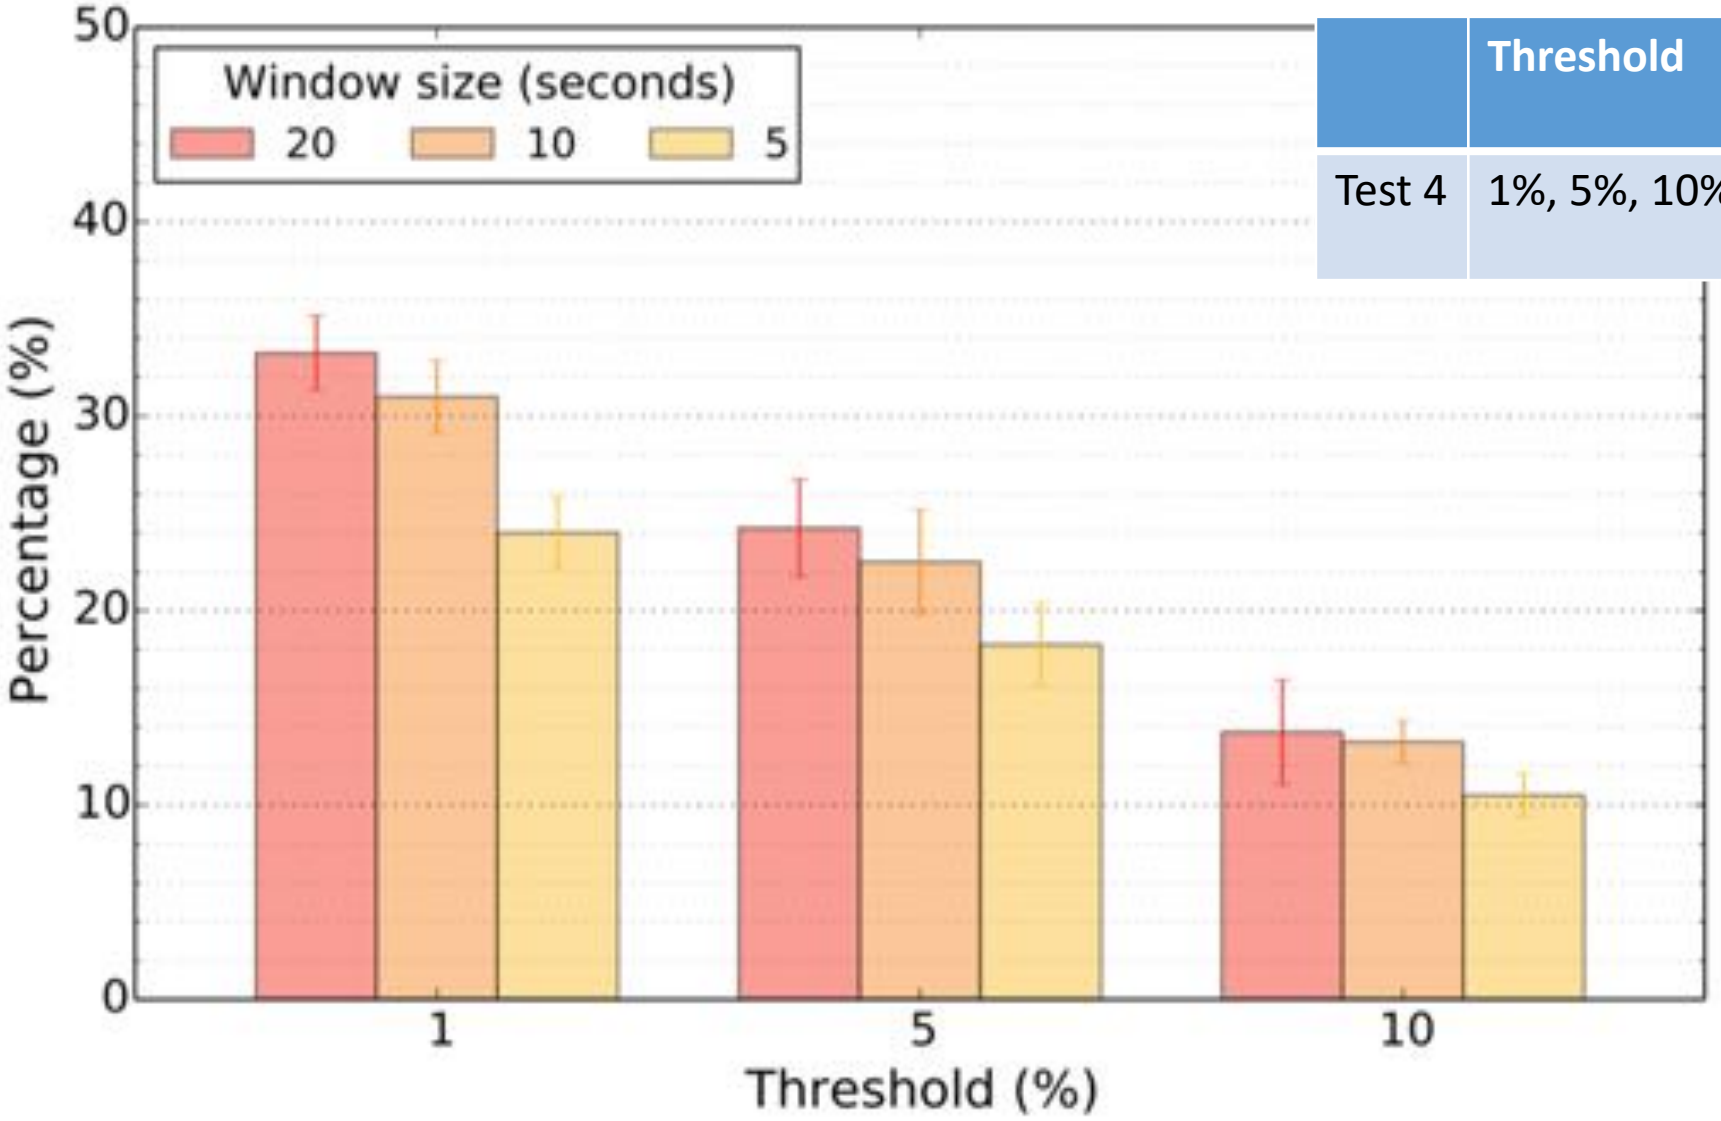

	Threshold	Detection	Trace	Windows (sec)
Test 4	1%, 5%, 10%	HHH	60 min	[20], [10], [5]

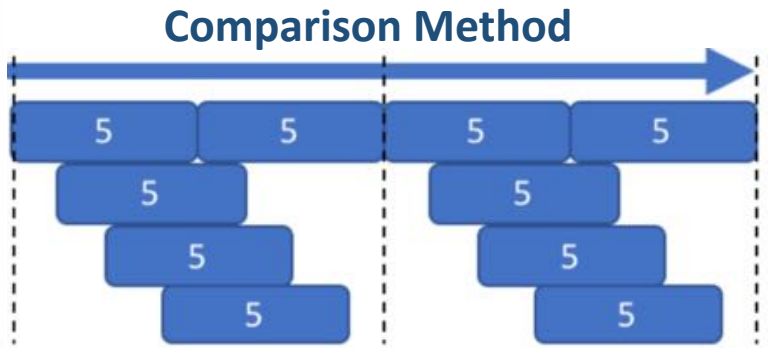

Food for thought

-- Fixed time or sliding window? Your call…

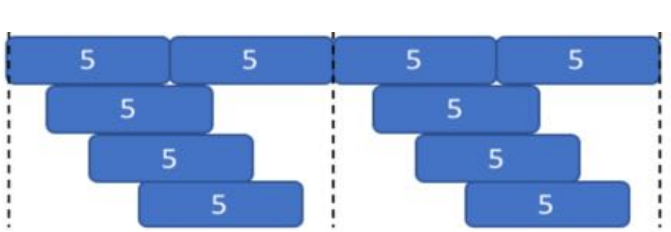

-- The size matters, even small variations make a difference..

-- Can you afford the missing information (H4)?

-- Can the “Hidden” information (34%) have ramifications on network management decisions?

-- Can it be the case for new windowless-based algorithms?