Measurement

COS 597E: Software Defined Networking

Jennifer Rexford
Princeton University
MW 11:00am-12:20pm

Measurement vs. Metrics

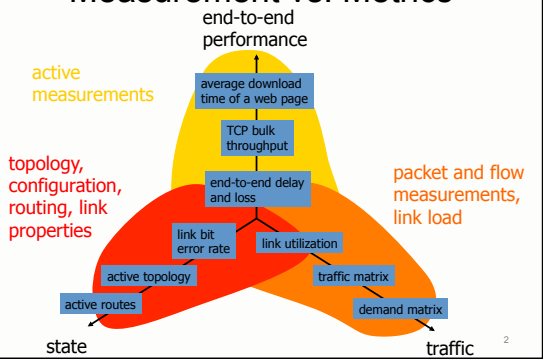

2

Reducing Measurement Overhead

Filtering, Aggregation, and Sampling

3

Reducing Measurement Overhead

- Measurement overhead
  - In some areas, you could measure everything
  - Information processing not the bottleneck
  - Networking: thinning is crucial!
- Reducing measurement traffic:
  - Filtering
  - Aggregation
  - Sampling
  - ...and combinations thereof

4

Filtering

- Measure selectively
  - Only record statistics for a subset of traffic
- Examples
  - Matching a destination prefix
  - For a certain service class
  - Violating an access control list
  - TCP SYN or RST packets (attacks, abandoned http download)

5

Aggregation

- Coarse-grained statistics
  - Combine related traffic together
- Example: srcip and dstip

src	dest	# pkts	# bytes
a.b.c.d	m.n.o.p	374	85498
e.f.g.h	q.r.s.t	7	280
i.j.k.l	u.v.w.x	48	3465
....	....	....	

6

Sampling

- Select a "random" subset of traffic
  - Representative of all traffic
- Examples
  - Random
  - Round-robin
  - Hash-based

16 packets: estimate ¾ blue and ¼ red

7

Comparison

	Filtering	Aggregation	Sampling
Precision	exact	exact	approximate
Generality	constrained a-priori	constrained a-priori	general
Local Processing	filter criterion for every object	table update for every object	only sampling decision
Local memory	none	one bin per value of interest	none
Compression	depends on data	depends on data	controlled

8

Traffic Monitoring Techniques

Links, Flows, and Packets

9

Traffic Monitoring Techniques

- Link monitoring
  - Group all packets on the same link
  - Average load, loss, corruption, …
- Flow monitoring
  - Group similar packets into flows
  - Same header fields and close in time
- Packet monitoring
  - Capture first n bytes of a packet

10

IP Flows

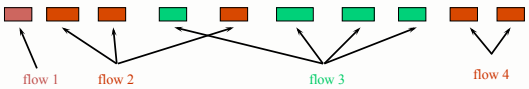

- Set of packets that "belong together"
  - Source/destination IP addresses and port numbers
  - Same protocol, ToS bits, …
  - Same input/output interfaces at a router (if known)
- Packets that are "close" together in time
  - Max spacing between packets (e.g., 15 sec, 30 sec)
  - Example: flows 2 and 4 are different flows due to time

11

Netflow: Traffic Data

- Packet header information
  - Src/dst IP, src/dst port numbers, ToS bits, …
- Summary statistics
  - Start and finish times
  - Number of bytes and packets
  - TCP flags (logical OR over all packets)

4 packets
1436 bytes
SYN, ACK, & FIN

12

Netflow: Routing Information

- Routing information
  - Input and output ports
  - Source and destination prefix
  - Source and destination Autonomous System

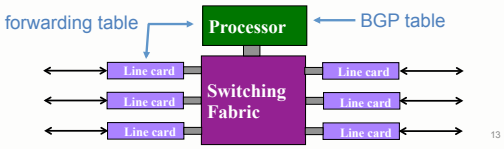

13

Netflow: Measuring in Passing

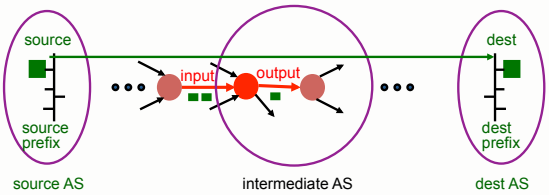

Source and destination: IP header
Source and dest prefix: forwarding table or BGP table
Source and destination AS: BGP table

14

Netflow: Implementation

- Maintain a cache of active flows
  - Storage of byte/packet counts, timestamps, etc.
- Compute a key per incoming packet
  - Concatenation of source, destination, port #s, etc.
- Index into the flow cache based on the key
  - Creation or updating of an entry in the flow cache

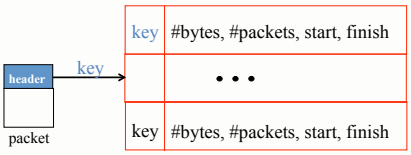

15

Netflow: Implementation

- Flow timeout
  - Remove flows that have not received a packet recently
  - Periodic sequencing through the cache to time out flows
- Cache replacement
  - Remove flow(s) when the flow cache is full
  - Evict existing flow(s) upon creating a new cache entry
  - Apply eviction policy (LRU, random flow, etc.)
- Long-lived flows
  - Remove flow(s) that persist for a long time (e.g., 30 min)
  - … otherwise flow statistics don't become available
  - … and the byte and packet counters might overflow

16

Sampled Netflow

- Packet sampling
  - Perform operations on 1/m packets
- Reducing overhead
  - Avoid per-packet overhead on (m-1)/m packets
  - Avoid creating records for the many small flows
- May split some long flows

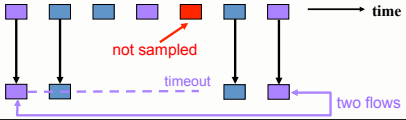

17

Netflow vs. sFlow

- Netflow
  - Aggregates (sampled) IP packets into flows
  - (Data-plane overhead of storing flow state)
  - Originally only on Cisco routers
- sFlow
  - Exports packet samples directly
  - Measures layer-two packets (ARP, DHCP)
  - Polls on-switch counters

http://blog.sflow.com/2011/10/comparing-sflow-and-netflow-in-vswitch.html

18

Sketches

19

Streaming Algorithms

- Processing data streams
  - Input items presented one at a time
  - Each item examined (say) only once
- Limited resources
  - Processing
  - Memory
- Approximate answer
  - Based on a summary or "sketch" of the data
  - Provable space-accuracy trade-off

20

Bloom Filter

- Set-membership problem
  - Was element x in the input stream?
- Solution (with false positives)
  - Per-item: compute s hash functions, set bit to1
  - Query: compute s hashes, check if all bits are 1

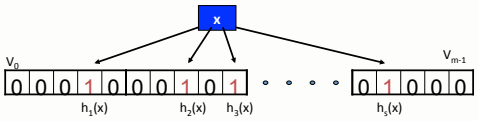

21

Count Min Sketch

- Counting problem
  - Count number of occurrences of item x
- Solution
  - Per-item: compute s hashes, increment counts
  - Query: compute s hashes, select the min value

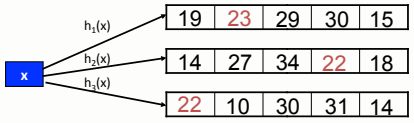

22

SDN Measurement

- OpenFlow 1.0
  - Rules with byte and packet counters
  - Sending packets to a collector or controller
- Existing measurement techniques
  - E.g., sFlow support
- What measurement support do we want?

23