Space-Code Bloom Filter for Efficient Traffic Flow Measurement

Abhishek Kumar, Jun (Jim) Xu
Networking and Telecommunications Group
College of Computing
Georgia Institute of Technology
{akumar,jx}@cc.gatech.edu

Li (Erran) Li
Bell Laboratories
erranlli@bell-labs.com

Jia Wang
AT&T Labs - Research
jiawang@research.att.com

Problem Statement

Goal: To keep track of the total number of packets belonging to **each** flow at high speed links.

Applications like traffic characterization, anomaly detection, per-flow QoS etc., need to know the size of all flows.

Definition of Flow: All packets with the same flow-label.
The flow-label can be defined as any combination of fields from the IP header, e.g <Source IP, source Port, Dest. IP, Dest. Port, Protocol>.

Why is per-flow measurement hard?

- Majority of the packets belong to large flows, yet a majority of the flows are small.
- High cost of maintaining per-flow data-structures. Amortization is difficult.
- No clear definition of the “end” of a flow.

Related Approaches

Sampling Sample packets with a fixed probability p and trace/process headers of sampled packets. This is the approach used by Cisco Netflow.

- Flow-sizes can be inferred from sampled data.
- Space-intensive.
- Inaccurate for small flows.

Keep track of elephants

- Fast algorithm to filter packets from large flows. [Estan and Varghese, 2002]
- Maintain counters for large flows only.
- Success in tracking the largest few flows (e.g. carrying $\geq 1\%$ of the total traffic) with limited memory.

Our Solution – Space-Code Bloom Filter (SCBF)

- Tracks **all** flows – from mice to elephants.
- Provides approximate estimate of flow-size.
- The relative error in estimation is the same for all flow sizes.
- The approx. estimates are close to the actual flow-size with high probability.

Operation of Space-Code Bloom Filter (SCBF) – Insertion Phase

- Measurement proceeds in epochs (e.g. 10 second).
- Maintain an aggregate synopsis data-structure.
- Update the data-structure on every packet arrival.
- Write-only data structure → fast updates, low hardware complexity.
- Copies of the synopsis are paged to disk periodically.

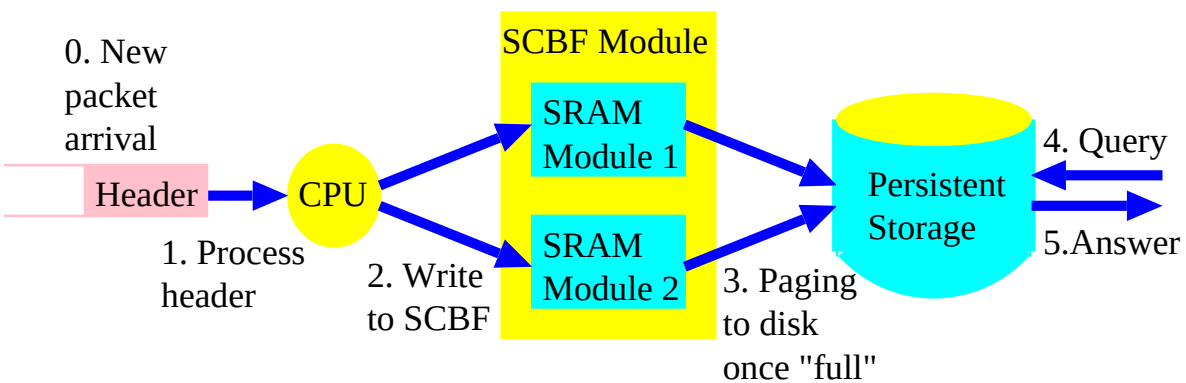

Operation of Space-Code Bloom Filter (SCBF) – Query Phase

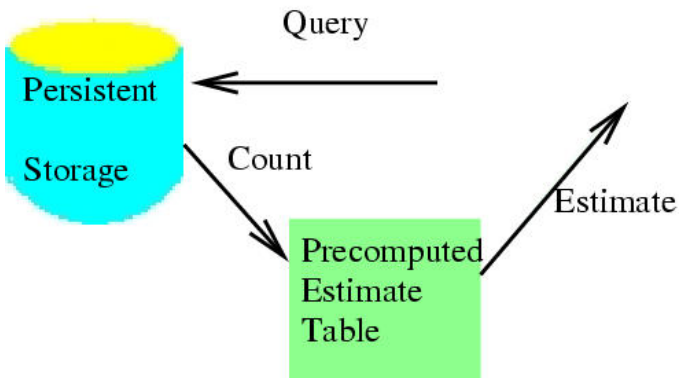

- Queries provide a flow-label and ask for its size.
- Obtain a “count” from the data-structure and then lookup a precomputed table to return approximate size of the flow.
- This provides approximate estimates that have low relative error with high probability.

Design of the aggregate data-structure

- Bloom-filters answer set-membership questions with high accuracy.
- Space-Code Bloom Filter answers multiset-membership questions with high accuracy.
- Use a number of “virtual” Bloom-filters, thus spreading the multiplicity information over **space**.
- Hash functions allow us to “isolate” flows from each other, thus spreading the multiplicity information over **code**.
- A **Space-Code** Bloom filter represents a large number of statistical estimators in parallel.

Performance of SCBF - complexities

- Computational complexity – compute 5 hash-fuctions and write 5 bits per packet.
- Space complexity – 4 bits of storage required for each packet.
- Can operate at OC768 (40 Gbps) with 5 ns SRAM.
- More than 80% responses are within $\pm 25\%$ of the actual value.

Accuracy of SCBF

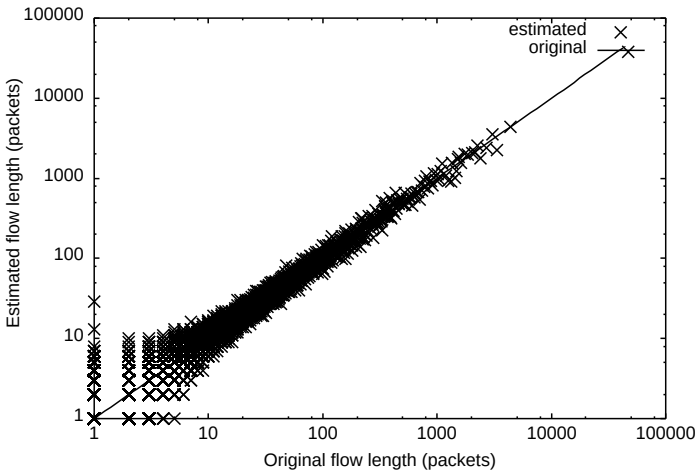

(a) Original vs. estimated flow size. Note that both axes are on logscale.

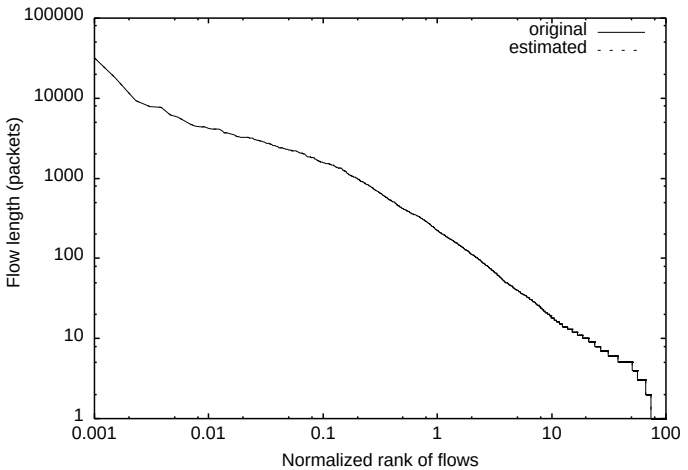

(b) Distribution of the original and estimated flow size.

Conclusions

- Space-Code Bloom Filters can track the **approximate** size of every flow.
- Per-flow accounting without per-flow state.
- The relative error in approximation is same for all flow-sizes.
- Very fast (upto OC768) implementations possible due to “write-only” nature of updates.
- Design parameters of SCBF can be tuned to trade storage space and CPU cycles for accuracy.

Acknowledgments

We thank Oliver Spatschek for providing us with the traffic traces.

Questions ???

Accuracy of SCBF using Maximum Likelihood Estimation (MLE)

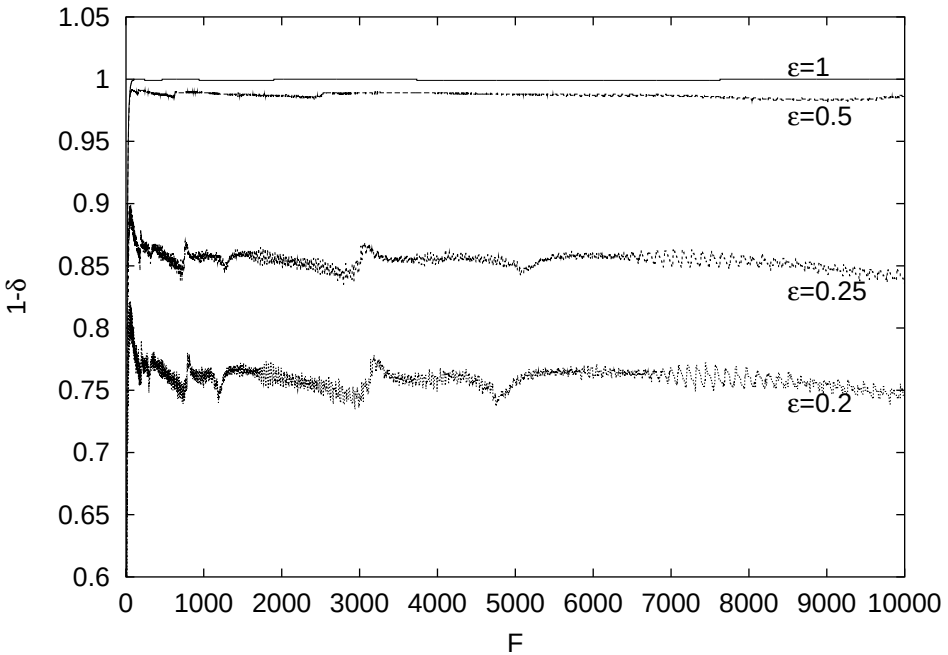

(c) Theoretical accuracy of MLE using 32 groups.

Accuracy of SCBF using Maximum Likelihood Estimation

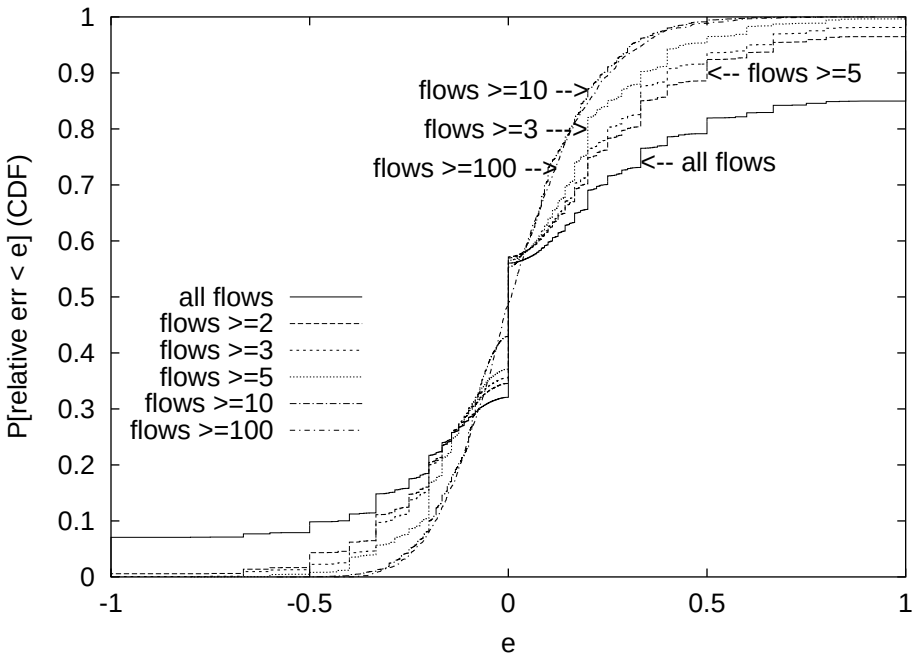

(d) CDF of relative error for flows of various size