

Space-Code Bloom Filter for Efficient Per-Flow Traffic Measurement

Abhishek Kumar¹
akumar@cc.gatech.edu

Li (Erran) Li²
erranli@bell-labs.com

Oliver Spatschek³
spatsch@research.att.com

Jia Wang³
jiawang@research.att.com

Jun (Jim) Xu¹
jx@cc.gatech.edu

Per-flow traffic measurement is critical for usage accounting, traffic engineering, and anomaly detection. Previous methodologies are either based on random sampling (e.g., Cisco’s NetFlow), which is inaccurate, or only account for the “elephants”. We introduce a novel technique for measuring per-flow traffic approximately, for all flows regardless of their sizes, at very high-speed (say, OC768). The core of this technique is a novel data structure called Space Code Bloom Filter (SCBF).

SCBF is an approximate representation of a *multiset*; each element in this multiset is a traffic flow and its multiplicity is the number of packets in the flow. The SCBF synopsis data-structure is updated upon each packet arrival, so that it will not fail to record the presence of any flow, small or large. The synopsis data-structure is paged to disk from time-to-time when it becomes “full”. A query concerning the size of a flow can be made to a SCBF page stored on the disk by providing the flow-label. The result of the query is the approximate number of packets in the flow during the measurement epoch recorded by that SCBF page. The aggregate size of a flow, independent of measurement epochs, can be obtained by taking the sum of flow-size estimates obtained by querying pages corresponding to consecutive epochs.

SCBF processes a query in two stages. First, a raw observation is obtained from the synopsis data-structure. This observation is then used to estimate the size of the flow through either of two mechanisms – Maximum Likelihood Estimation (MLE) or Mean Value Estimation (MVE). The final estimate returned is an approximation of the actual size of the flow. Our design goal for SCBF is to achieve “constant relative error tolerance”. In other words, SCBF returns estimates that are within a constant factor of the actual flow-size with high probability. In our evaluation of SCBF on real traces, estimates of flow size

This work was supported in part by the National Science Foundation under Grant ANI-0113911 and under NSF CAREER Award Grant ANI-0238315. Names of non-Student authors are in alphabetical order.

¹College of Computing, Georgia Institute of Technology

²Bell Labs, Lucent Technologies

³AT&T Labs - Research

are within 20 percent of the actual value for 80 to 95 the time. Through parameter tuning, SCBF allows for graceful tradeoff between measurement accuracy and computational and storage complexity.

SCBF also contributes to the foundation of data streaming by introducing a new paradigm called blind streaming. Since the data-structure is write only, the processing for individual packets is independent of the state maintained in the data-structure. This enables a fast and simple hardware implementation.




Fig. 1. Original vs. estimated flow size. Note that both axes are on logscale.

Experiments with a software implementation of SCBF and both estimation algorithms on traffic traces from a Tier-1 ISP backbone agree very well with our theoretical analysis. Figure 1 shows the scatter diagram of flow size estimated using MLE (y axis) vs. actual flow size (x axis) in terms of the number of packets in the flow. The fact that all the points are concentrated within a narrow band of fixed width along the $y = x$ line indicates that our estimates are consistently within a constant factor of the actual flow size.