

Discount Counting for Fast Flow Statistics on Flow Size and Flow Volume

Abstract:

A complete flow statistics report should include both flow size (the number of packets in a flow) counting and flow volume (the number of bytes in a flow) counting. Although previous studies have contributed a lot to the flow size counting problem, it is still a great challenge to well support the flow volume statistics due to the demanding requirements on both memory size and memory bandwidth in monitoring device. In this paper, we propose a DIScount COunting (DISCO) method, which is designed for both flow size and flow bytes counting. For each incoming packet of length l , DISCO increases the corresponding counter assigned to the flow with an increment that is less than l . With an elaborate design on the counter update rule and the inverse estimation, DISCO saves memory consumption while providing an accurate unbiased estimator. The method is evaluated thoroughly under theoretical analysis and simulations with synthetic and real traces. The results demonstrate that DISCO is more accurate than related work given the same counter sizes. DISCO is also implemented on the **network** processor Intel IXP2850 for a performance test. Using only one microengine (ME) in IXP2850, the throughput can reach up to 11.1 Gb/s under a traditional traffic pattern. The throughput increases to 39 Gb/s when employing four MEs.