Counter Braids

Yi Lu
Department of Electrical Engineering
Stanford University
yi.lu@stanford.edu

Andrea Montanari
Departments of EE and Statistics
Stanford University
montanar@stanford.edu

Balaji Prabhakar
Departments of EE and CS
Stanford University
Email: balaji@stanford.edu

***Abstract*— In this extended abstract the authors summarize recent work they have done on the design of a novel counter architecture for estimating flow sizes in high-speed networks, the algorithms and the theory that goes along with it. This note will provide a description of the problem and our approach. It serves as a pointer to papers ([2] and [3]) which cover the design, the algorithms and the theory of Counter Braids in more detail.**

I. BACKGROUND: NETWORK MEASUREMENT

There is an increasing need for fine-grained network measurement to aid the management of large networks [1]. Network measurement consists of counting the size of a logical entity called "flow," at an interface such as a router. A flow is a sequence of packets that satisfies a common set of rules. For instance, packets with the same source (destination) address constitutes a flow. Measuring flows of this type gives the volume of upload (download) by a user and is useful for accounting and billing purposes. Measuring flows with a specific 5-tuple in the packet header gives more detailed information such as routing distribution and types of traffic in the network. Such information can help greatly with traffic engineering and bandwidth provisioning. Flows can also be defined by packet classification; for example, ICMP Echo packets used for network attacks form a flow. Measuring such flows is useful during and after an attack for anomaly detection and network forensics.

II. PROBLEM STATEMENT

This is the central problem: Let F_i be the size of flow i measured in packets or bytes. The F_i are drawn from some common distribution, which is heavy-tailed (typically Pareto for Internet flows). The packets of n flows pass through a router in a measurement interval. The problem is to measure the values F_i using counters.

This problem is challenging because of the following constraints: (i) The *size* of the memory needs to be large to store all the counters and to ensure accurate counting, and (ii) the *speed* of memory needs to be very high because packets arrive very rapidly (roughly once every 40 nanosecs). Since very fast and very large are prohibitively expensive, the common practice in research and in the industry has been to obtain approximate flow counts; that is, to detect large *elephant* flows quickly and measure their sizes.

III. OUR APPROACH

Recently the authors and their collaborators have developed [2] a novel method for per-flow traffic measurement that is fast, highly memory efficient and accurate. At the core of this method is a novel counter architecture called "Counter Braids."

We now briefly describe this architecture. Figure 1 shows a naive counter architecture that stores five flow sizes in counters of the same depth. The least significant bit (LSB) is the one closest to the flow node, the MSB is furthest out. Darkened circles represent a 1, hollow circles a 0. Since the depth of each counter has to exceed the size of the largest flow, all counters are equal in size. This leads to an enormous wastage of space for the short flows.

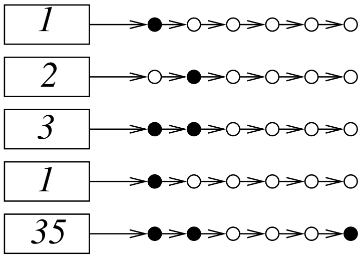

Fig. 1. A simple counter structure: to each flow size we associate its binary representation (filled circle = 1, empty circle = 0).

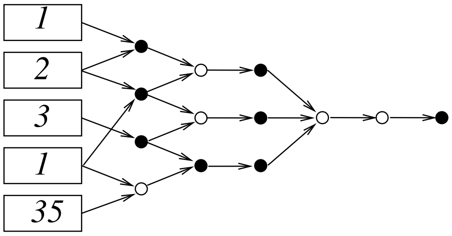

Fig. 2. Counter braids.

Figure 2 shows Counter Braids where the flows *share* the more significant bits of the counters. The rule for updating the bits is simple: when the count corresponding to a flow increases by 1, a 1 is added to the counter nodes on the outgoing edges; addition is modulo 2 (in general, it is modulo the size of the counter). If a counter node overflows, its value is set to 0 and a 1 is added to the counter nodes on its outgoing edges. This is recursed through the structure until no more overflows occur. As a result of the MSBs being shared, counter space is greatly reduced.

Designing Counter Braids thus has two components: (i) Choosing the requisite number of counters at each level and the interconnection network between the various levels, and (ii) a decoding algorithm for recovering flow sizes from counter values.

IV. MAIN RESULTS

Counter Braids is a novel, space-efficient, and large-scale counter architecture. It has the following features:

1) Incremental compression of flow sizes as packet arrive; only a small number (e.g. 3) of counters are incremented at each packet arrival.
2) Counter Braids is an optimal, incremental compressor: the number of bits needed to store flow sizes is equal to the entropy lower bound. Even though the decoder used to achieve optimality (the Maximum Likelihood Decoder) is NP-hard in the worst-case, this result is surprising because counter values are essentially a *linear* transformation of flow sizes. In fact, linear transformations are not known to be optimal compressors, in general.
3) A linear-complexity message-passing decoding algorithm that recovers flow sizes from the compressed version with vanishing error at link speed. This is interesting because, in general, message-passing algorithms are known to be exact and converge only on trees; in particular, they need not converge on loopy graphs. We not only get convergence here, but also exactness (convergence to the correct answer).
4) The total counter space used by the message-passing algorithm is comparable to the optimal. Further, the message-passing algorithm is fully analyzable, enabling the choice of design parameters for different hardware requirements.

V. RELATED LITERATURE

There is much prior work on network measurement and there is also related work in the areas of Compressed Sensing and Sparse Random Graph Codes. A detailed review of this literature is presented in [2] and [3]. We refer the interested reader to these papers.

REFERENCES

[1] M. Grossglauser and J. Rexford, "Passive traffic measurement for IP operations," *The Internet as a Large-Scale Complex System,* Oxford University Press, 2002.

[2] Y. Lu, A. Montanari, B. Prabhakar, S. Dharmapurikar, A. Kabbani, "Counter Braids: A novel counter architecture," to appear in *Proc. of ACM Sigmetrics,* 2008.

[3] Y. Lu, A. Montanari, B. Prabhakar, "Detailed Network Measurements Using Sparse Graph Counters: The Theory," *Proc of the Annual Allerton Conference*, 2007.