

Diamond Sketch: Accurate Per-flow Measurement for Real IP Streams

Tong Yang*, Siang Gao*, Zhouyi Sun*, Yufei Wang*, Yulong Shen[†], Xiaoming Li*

*Department of Computer Science, Peking University, China

[†]Xidian University, China

Abstract—Existing sketches often have low memory efficiencies when performing per-flow measurement tasks on skewed IP streams. In this paper, we propose **Diamond Sketch**, a novel sketch that dynamically assigns an appropriate number of atom sketches to each flow on demand, improving the accuracy considerably while keeping a comparable speed.

I. INTRODUCTION

Per-flow measurement is a critical issue in computer networks. It provides information for anomaly detection, capacity planning, accounting and billing, and service provision [1], [2]. One fundamental problem in per-flow measurement is to estimate the flow size, which is the number of packets in each flow. A flow is often identified by a certain combination of fields in the five-tuple in the packet’s header: source IP address, destination IP address, source port number, destination port number, and protocol type. Real network flows have two characteristics: *high-speed* and *non-uniform distribution*. On the one hand, the speed of network traffic is so high that it is very hard to make a precise record of sizes of flows [3]. On the other hand, the sizes of network flows are usually non-uniformly distributed [4]. A small part of flows have very large sizes (called *elephant flows*) while most flows have small sizes (called *mice flows*). A typical distribution of this kind of network traffic is Zipfian [5], and network traffic following Zipfian distribution are called *skewed* traffic.

Due to the high speed of network traffic, approximately recording and estimating data with sketches has gained popularity [4]. Sketches are probabilistic data structures, and can achieve small memory footprints, high accuracy, and fast speeds of insertion and query. Unfortunately, most existing sketches cannot work well on skewed network traffic. Conventional sketches (such as CM [6], CU [7], Count [8], CSM [9] sketches) use counters that are of the same size to store the number of packets. Elephant flows are often considered to be more important than mice flows, thus the counters need to be large enough to represent the largest size of the elephant flows. However, the large quantity of mice flows means that most counters will just represent a small value, and the higher bits in these counters are all 0, leading to a waste of memory.

In this paper, we propose a new sketch, namely the Diamond sketch, as it takes on a diamond shape. The *key idea* of Diamond is to assign atom sketches to flows on demand. Specifically, an appropriate number of atom sketches are assigned dynamically and automatically to record the sizes of elephant and mice flows, thus optimizing memory efficiency.

For mice flows, we use only one or two atom sketches to record their sizes. For elephant flows, we dynamically increase the number of atom sketches to record their sizes.

II. DIAMOND SKETCH

A. Data Structure

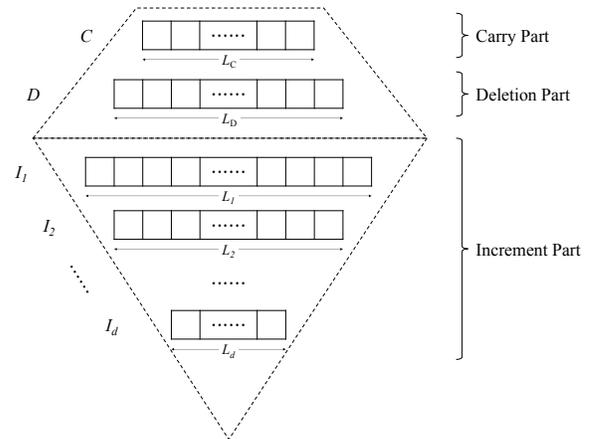


Fig. 1. Data structure of the Diamond Sketch.

As shown in Figure 1, a Diamond Sketch consists of three parts: increment part, carry part, and deletion part. **Increment part** together with carry part records the size of each flow that has been inserted. It consists of d atom sketches, where we denote the i^{th} atom sketch with I_i . Each atom sketch I_i is composed of L_i counters, and each counter contains w_1 bits. We denote the j^{th} counter of the i^{th} atom sketch with $I_i[j]$. $\{L_1, L_2, \dots, L_d\}$ is a decreasing sequence of numbers. **Carry part** records the *overflow depth*¹ of each flow. It is an atom sketch composed of L_C counters, each of which contains w_2 bits ($2^{w_2} \geq d$ is a requisite). We denote this atom sketch with C , and the j^{th} counter with $C[j]$. **Deletion part** is used to support deletions. It is an atom sketch composed of L_D counters, each of which contains w_3 bits. We denote this atom sketch with D , and the j^{th} counter with $D[j]$.

Each atom sketch I_i , C , and D is associated with k_1 , k_2 , and k_3 hash functions, whose output is uniformly distributed in the range $[1, L_i]$, $[1, L_C]$, and $[1, L_D]$. And we denote the j^{th} hash function with $h_j^i(\cdot)$, $h_j^C(\cdot)$, and $h_j^D(\cdot)$, respectively.

¹Suppose a flow causes overflows in atom sketch I_1 and I_2 when being inserted, and is successfully inserted into I_3 without overflows, then the deepest sketch this flow has reached is I_3 , and 3 is called the *overflow depth*.

B. Operations

The **initialization** for the Diamond Sketch is simply setting all counters $I_i[j]$, $C[l]$, and $D[m]$ to 0, where $1 \leq i \leq d$, $1 \leq j \leq L_i$, $1 \leq l \leq L_C$, and $1 \leq m \leq L_D$. Next, we will discuss the insertion, query, and deletion operation.

The **insertion** of a packet with flow ID e requires update in first the increment part and second the carry part. We compute the k_1 hash functions of the first atom sketch I_1 and increment the smallest one(s) of the k_1 counters $I_1[h_1^1(e)]$, $I_1[h_2^1(e)]$, ..., $I_1[h_{k_1}^1(e)]$ (we call them k_1 mapped counters for short) by 1. If more than one counters have the smallest value, we increment them all by 1. If all the k_1 counters hold the value $2^{w_1} - 1$, which is the largest number that w_1 bits can represent, then incrementing them by 1 will result in *overflows*. When the k_1 mapped counters in I_i overflow, the following steps need to be taken: **1)** Set the values of all k_1 mapped counters to 0; **2)** Compute the k_1 hash functions of I_{i+1} and increment the smallest one(s) of the k_1 counters $I_{i+1}[h_1^{i+1}(e)]$, $I_{i+1}[h_2^{i+1}(e)]$, ..., $I_{i+1}[h_{k_1}^{i+1}(e)]$ by 1; **3)** If all the k_1 mapped counters in I_{i+1} overflow, repeat step 1 and 2 until the *termination condition* is satisfied, and we record the overflow depth dep . There are two termination conditions: there is no overflow in I_i ($1 \leq i \leq d$) (we set dep to i) or the last atom sketch I_d overflows (we set dep to d); **4)** Compute the k_2 hash functions of \mathcal{C} , and get the smallest value of the k_2 counters: $C[h_1^C(e)]$, $C[h_2^C(e)]$, ..., $C[h_{k_2}^C(e)]$. If the smallest value is smaller than $dep - 1$, we need to update the carry part to guarantee that the smallest value equals $dep - 1$. Therefore, for each of the k_2 counters, if its value is less than $dep - 1$, we set its value to $dep - 1$; otherwise, we do nothing.

The **deletion** operation is not requisite in per-flow measurement, but can be added if needed in other scenarios. To delete a packet e , the Diamond Sketch only updates its deletion part. We compute the k_3 hash functions of \mathcal{D} and check the corresponding k_3 counters $D[h_1^D(e)]$, $D[h_2^D(e)]$, ..., $D[h_{k_3}^D(e)]$. If all k_3 mapped counters are $2^{w_3} - 1$, which is the biggest number that w_3 bits can represent, we do nothing; otherwise, we increment the smallest one(s) among the k_3 mapped counters by 1.

To **query** the size of a flow with ID e , we need to check the status of all the three parts. **1)** We check the carry part. Specifically, we compute the k_2 hash functions of \mathcal{C} and get the smallest value v_C among $C[h_1^C(e)]$, $C[h_2^C(e)]$, ..., $C[h_{k_2}^C(e)]$. **2)** We check the increment part. Let $dep = 1 + v_C$. We need to check dep atom sketches I_1, I_2, \dots, I_{dep} . For each atom sketch I_i , we compute k_1 hash functions and get the smallest value V_i of those k_1 mapped counters $I_i[h_1^i(e)]$, $I_i[h_2^i(e)]$, ..., $I_i[h_{k_1}^i(e)]$. Once all dep smallest values have been calculated, the query result determined by increment part and carry part is computed by the following formula: $V_{insert} = \sum_{i=1}^{dep} V_i \cdot (2^{w_1})^{i-1}$. **3)** We check the deletion part. Specifically, we compute the k_3 hash functions of \mathcal{D} and return the smallest value (V_{delete}) among $D[h_1^D(e)]$, $D[h_2^D(e)]$, ..., $D[h_{k_3}^D(e)]$. **4)** We compute $V_{insert} - V_{delete}$ as the query result.

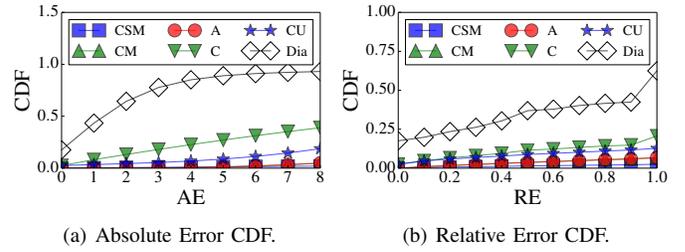


Fig. 2. AE-CDF and RE-CDF on IP streams. *Dia* stands for *Diamond sketch*.

III. EVALUATION

We use real IP trace streams as experimental datasets, which are captured by the main gateway of our campus. We compare Diamond Sketch with CU [7], CM [6], Augmented [10], Count [8], and CSM [9] sketches. The metrics used are absolute error (AE) and relative error (RE)².

As shown in Fig. 2, the percentage of flows whose AE is no bigger than 1 of our Diamond sketch is 43.39%, which is 124.23, 2018.3, 1823.25, 5.3, and 12.5 times higher than the corresponding percentages of CSM, CM, A, C, CU sketches, respectively; the percentage of flows whose RE is less than 1.0 of our Diamond sketch is 62.46%, which is 24.6, 9.29, 9.29, 3.02, and 4.91 times higher than the corresponding percentages of CSM, CM, A, C, CU sketches, respectively. The results show that the performance of Diamond is far better than other sketches in terms of accuracy.

REFERENCES

- [1] Xenofontas Dimitropoulos, Paul Hurley, and Andreas Kind. Probabilistic lossy counting: an efficient algorithm for finding heavy hitters. *ACM SIGCOMM Computer Communication Review*, 38(1):5–5, 2008.
- [2] MyungKeun Yoon, Tao Li, Shigang Chen, and J-K Peir. Fit a spread estimator in small memory. In *Proc. IEEE INFOCOM 2009*.
- [3] Yuliang Li, Rui Miao, Changhoon Kim, and Minlan Yu. Flowradar: a better netflow for data centers. In *Proc. USENIX NSDI*, 2016.
- [4] Graham Cormode. Sketch techniques for approximate query processing. *Synopses for Approximate Query Processing: Samples, Histograms, Wavelets and Sketches, Foundations and Trends in Databases*. NOW publishers, 2011.
- [5] David MW Powers. Applications and explanations of zipf’s law. In *Proceedings of the joint conferences on new methods in language processing and computational natural language learning*, pages 151–160. Association for Computational Linguistics, 1998.
- [6] Graham Cormode and Shan Muthukrishnan. An improved data stream summary: the count-min sketch and its applications. *Journal of Algorithms*, 55(1):58–75, 2005.
- [7] Cristian Estan and George Varghese. New directions in traffic measurement and accounting: Focusing on the elephants, ignoring the mice. *ACM Transactions on Computer Systems (TOCS)*, 21(3):270–313, 2003.
- [8] Moses Charikar, Kevin Chen, and Martin Farach-Colton. Finding frequent items in data streams. In *International Colloquium on Automata, Languages, and Programming*, pages 693–703. Springer, 2002.
- [9] Tao Li, Shigang Chen, and Yibei Ling. Per-flow traffic measurement through randomized counter sharing. *IEEE/ACM Transactions on Networking (TON)*, 20(5):1622–1634, 2012.
- [10] Pratanu Roy, Arijit Khan, and Gustavo Alonso. Augmented sketch: Faster and more accurate stream processing. In *Proc. ACM SIGMOD*, pages 1449–1463, 2016.

²Let f_e be the real size of a packet e , and \hat{f}_e be the estimated size reported by the sketch. Then the *absolute error* (AE) of flow e is calculated as $|f_e - \hat{f}_e|$, and the *relative error* (RE) is calculated as $|f_e - \hat{f}_e|/f_e$.