

Optimal Algorithm for Profiling Dynamic Arrays with Finite Values

Dingcheng Yang, Wenjian Yu, Junhui Deng

BNRist, Dept. Computer Science & Tech.,

Tsinghua Univ., Beijing, China

ydc15@mails.tsinghua.edu.cn, yu-wj@tsinghua.edu.cn,

deng@tsinghua.edu.cn

Shenghua Liu

CAS Key Lab. Network Data Science & Tech.,

Inst. Computing Technology, Chinese Academy of

Sciences, Beijing, China

liushenghua@ict.ac.cn

ABSTRACT

How can one quickly answer the most and top popular objects at any time, given a large log stream in a system of billions of users? It is equivalent to find the mode and top-frequent elements in a dynamic array corresponding to the log stream. However, most existing work either restrain the dynamic array within a sliding window, or do not take advantages of only one element can be added or removed in a log stream. Therefore, we propose a profiling algorithm, named S-Profile, which is of $O(1)$ time complexity for every updating of the dynamic array, and optimal in terms of computational complexity. With the profiling results, answering the queries on the statistics of dynamic array becomes trivial and fast. With the experiments of various settings of dynamic arrays, our accurate S-Profile algorithm outperforms the existing methods, showing at least 2X speedup to the heap based approach and 13X speedup to the balanced tree based approach.

1 INTRODUCTION

Many online systems, especially with billions of users, are generating a large stream of logs [6], recording users' dynamics in the systems, e.g. users (un)follow other users, "(dis)like" objects, enter (exit) live video channels, and click objects. Then, a question is raised:

How can we efficiently know the most popular objects (including users), i.e. mode, top-K popular ones, and even the distribution of frequency in a fast and large log stream at any time?

Mathematically, the questions converge to calculating and updating the statistics of a dynamic array of finite values. Thus, the existing fast algorithms on the statistics are as follows:

Mode of an array. The mode of an array and its corresponding frequency can be calculated by sorting the array (if it's of numeric value) and scanning the sorted array in $O(n \log n)$ time, where n is the length of array [7]. Notice that through judging the frequency of mode we can solve the element distinctness problem, which was proven to have to be solved with $\Omega(n \log n)$ time complexity [12, 15]. Therefore, calculating the mode of an array has the lower bound of $\Omega(n \log n)$ as well. If the elements of array can only take finite values, the complexity of calculating the mode can be reduced. Suppose they can only take m values. One can use m buckets to store the frequency of each distinct element. Then, the mode can be calculated in $O(n + m)$ time by scanning the m buckets.

The problem of range mode query, which calculates the mode of a sub-array $A[i \dots j]$ for a given array A and a pair of indices (i, j) , has also been investigated [4, 10, 13]. The array with finite

values was considered. With a static data structure, the range mode query can be answered in $O(\sqrt{n/\log n})$ time [4].

Majority and frequency approximation. The majority is the element whose frequency is more than half of n . An algorithm was proposed to find majority in $O(n)$ time and $O(1)$ space [3]. Many work on the statistics like frequency count and quantiles, are under a setting of sliding window [1, 2, 5, 8, 11]. They consider the most recently observed data elements (within the window) and calculate the statistics. Space-efficient algorithms were proposed to maintain the statistics over the sliding window on a stream.

However, those existing work slow down their algorithms without considering that the increase and decrease of object frequency are always 1 at a time in log streams. Therefore, we propose an algorithm S-Profile to keep profiling the dynamic array. With such a profile, we can answer the queries of the statistics: mode, top-K and frequency distributions.

In summary, S-Profile has the following advantages:

- **Optimal efficiency:** S-Profile needs $O(1)$ time complexity and $O(m)$ space complexity to profile dynamic arrays, where m is the maximum number of objects.
- **Querying Statistics:** With the profiling, we have sorted frequency-object pairs, and can simply answer the queries on mode, top-K, majority and other statistics in $O(1)$.
- **Applicable:** Our S-Profile can be plugged into most of log streams in many systems, and profiling objects of interest.

In experiments, S-Profile compares with the existing methods in various settings of dynamic arrays, and shows its performance and robustness.

2 AN $O(1)$ -COMPLEXITY ALGORITHM FOR UPDATING THE MODE AND STATISTICS

We define tuples (x_i, c_i) as a log stream, where x_i and c_i is the object id and action in the i -th tuple. Action c_i can be either "add" or "remove", which, for example, can indicate object x_i is "liked" or "disliked", or user x_i is followed or unfollowed. Conceptually, we could imagine a dynamic array A of objects associated with a log stream, by appending object x_i into A if c_i is "add", and deleting object x_i from A if c_i is "remove". Dynamic array A is not necessarily generated and stored, which is defined for convenient description of our algorithm.

Therefore, our problem can be described as follows:

INFORMAL PROBLEM 1 (PROFILING DYNAMIC ARRAY). **Given:** a log stream of tuples (x_i, c_i) adding and removing an object each time,

- **To build:** a data structure profiling the dynamic array A of objects associated with the log stream,
- **Such that:** at any time answering the queries on mode, top-K and other statistics of objects is trivial and fast.

Let m be the maximum number of distinct objects in a log stream or dynamic array A . Without loss of generality, we assume

© 2019 Copyright held by the owner/author(s). Published in Proceedings of the 22nd International Conference on Extending Database Technology (EDBT), March 26-29, 2019, ISBN 978-3-89318-081-3 on OpenProceedings.org. Distribution of this paper is permitted under the terms of the Creative Commons license CC-by-nc-nd 4.0.

id $x_i \in [1, m]$, i.e. integers between 1 and m . For any m distinct objects, we can map them into the integers from 1 to m as ids.

We can use m buckets to store the frequency of each distinct object. Let F be such a frequency array with length m . $F[i] \in F$ is the frequency of object with id i . With F , most statistics of A can be calculated without visiting A itself. For example, the mode of A refers to locations in F where the element has the maximum value. Although updating F with each tuple of a log stream trivial and costs $O(1)$, finding the maximum value in F at each time is still time consuming.

Therefore, we first introduce a proposed data structure of a profile, named “block set”, which can answer the statistical queries in a trivial cost. And we show later that we can maintain such a profile in $O(1)$ time complexity and $O(m)$ space complexity.

2.1 Proposed data structure for profiling

In order to find the mode of A , we just need to care about the maximum in F . If only integers are added to A , the maximum element of F can be easily updated. However, in Problem 1 removing integer is also allowed. This complicates the calculation of mode and other statistics of A . So, the sorted array T must be employed and maintained. To facilitate the queries, T can be implemented as a binary tree. The heap and balanced tree are two kinds of binary tree, and are widely used for efficiently maintaining an sorted array. Upon a modification on A , they both can be updated in $O(\log m)$ time. The root node of a heap is the array element with the extreme value. This means the heap is only suitable for producing the A 's elements with either maximum frequency (the mode) or the minimum frequency. The balanced tree is good at answering the query of median of A , and can also output the mode and top-K elements, etc. It should be pointed out, these general algorithms do not take the particularity of Problem 1 into account (the modification on F is restricted to plus 1 or minus 1). By the way, no one can maintain the sorted array under an arbitrary modification with a time complexity below $O(\log m)$, because it can be regarded as a sorting problem which has been proven to have $\Omega(m \log m)$ time complexity.

We use Figure 1 as an example to illustrate the proposed data structure for maintaining T . Suppose T has frequency values in ascending order. In order to locate the index of the i -th element of T in F and vice versa, two conversion arrays are defined: $TtoF$ and $FtoT$. In other words, we have $T[i] = F[TtoF_i]$ and $F[i] = T[FtoT_i]$. Here, we use both the subscript and bracket notation to specify an element of array. As shown in Figure 1(c), we can partition T into nonoverlapped segments according to its elements. Each such segment is called *block* here and represented by an integer triple (l, r, f) , where l and r are starting and ending indices respectively, and f is the element value (frequency). So, a block $b = (l, r, f)$ always satisfies:

- $1 \leq l \leq r \leq m$
- $T_i = f, \quad \forall i \in [l, r]$
- $T_l > T_{l-1}, \quad \text{if } l > 1$
- $T_r < T_{r+1}, \quad \text{if } r < m$

There are at most m blocks, which form a *block set*. It fully captures the information in the sorted array T . An array of pointers called *PtrB* is also needed to make a link from each element in T to its relevant block. According to the definition of block, we always have:

$$T_i = PtrB[i].f, \text{ and } PtrB[i].l \leq i \leq PtrB[i].r. \quad (1)$$

Here we use “ l ” to denote the member l of a block, and so on.

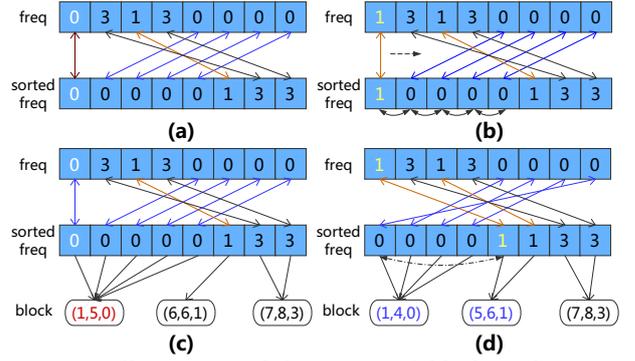


Figure 1: Illustration of the proposed *block set* for maintaining array T . (a) The initial F and T . (b) When “1” is added to A , a brute-force approach to maintain T includes four swaps of “1” rightwards and updates of $FtoT$ and $TtoF$. (c) The initial F and T , and the *block set*. F and T are up to be modified. (d) With the information of *block*, the swapping destination in T can be easily determined.

The block set B represents the sorted frequency array T , while with arrays $FtoS$, $StoF$ and $PtrB$ we no longer need to store F and T . These proposed new data structure well profile the dynamic array A . The remaining thing is to maintain them and answer the statistical queries on A in an efficient manner.

2.2 S-Profile: the $O(1)$ -Complexity Updating Algorithm

We first consider the situation where an integer is added to A . As shown in Figure 1(a) and 1(b), a brute-force approach to update T is swapping the updated frequency to its right-hand neighbor one by one, until T is in the appropriate order again. Now, with the proposed $PtrB$ and the block it points to, we can easily determine the index of T which is the destination of swapping the updated frequency. Then, we can update the relevant two blocks and pointer arrays (see Figure 1(d)).

Now, based on the situation shown as Figure 1(d), we assume a “4” is removed from A . As shown in Figure 2, we first locate the updated element in T . Then, with the information in its corresponding block we know with which it should be swapped. We further check if the updated frequency exists in T before. If it does not we need to create a new block (the case in Figure 2(b)), otherwise another block is modified.

The whole details of the algorithm for updating the data structure and returning the mode of A is described as Algorithm 1. We assume the data structures (B , $FtoS$, $StoF$ and $PtrB$) have been

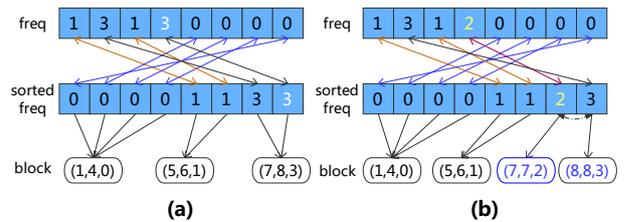


Figure 2: Illustration of the proposed data structure for a “remove” action on A . (a) The initial F and T , and the *block set*. A “4” is going to be deleted from A . (b) With the information of *block* and the pointer arrays, the *block set* can be easily updated to reflect the ordered T .

initialized, while Algorithm 1 responds to an event in the log stream and returns the updated mode and frequency.

Algorithm 1 S-Profile for updating the mode of array

Input: A tuple (x, c) in log stream, block set B , pointer arrays $FtoT$, $TtoF$ and $PtrB$, the length of sorted frequency array m

Output: The mode M , and the frequency v .

```

1: rank ← FtoT[x]
2: b ← PtrB[rank]
3: l ← b.l; r ← b.r;
4: if c is an "add" action then
5:   b.r ← b.r + 1
6:   if b.r < b.l then
7:     Delete b
8:   end if
9:   if r < m and b.f + 1 = PtrB[r + 1].f then
10:    PtrB[r] ← PtrB[r + 1]
11:    PtrB[r].l ← PtrB[r].l - 1
12:   else
13:    Create a new block in B and assign it to PtrB[r]
14:    PtrB[r] ← (r, r, b.f + 1)
15:   end if
16: else /* It's a "remove" action */
17:   b.l ← b.l - 1
18:   if b.r < b.l then
19:     Delete b
20:   end if
21:   if l > 1 and b.f - 1 = PtrB[l - 1].f then
22:    PtrB[l] ← PtrB[l - 1]
23:    PtrB[l].r ← PtrB[l].r + 1
24:   else
25:    Create a new block in B and assign it to PtrB[l]
26:    PtrB[l] ← (l, l, b.f - 1)
27:   end if
28: end if
29: M ← TtoF[PtrB[m].l ... PtrB[m].r]
30: v ← PtrB[m].f

```

As the proposed data structure maintains the sorted frequency array, it can be utilized to calculate the object with the minimum frequency (maybe a negative number) as well. We just need to replace Step 29 and 30 in Algorithm 1 with the following steps.

```

29a: M ← TtoF[PtrB[1].l ... PtrB[1].r]
30a: v ← PtrB[1].f

```

We can observe that the time complexity of the S-Profile algorithm is $O(1)$, as there is no iteration at all. The space complexity is $O(m)$, where m is the maximum number of objects in the log stream. Precisely, it needs $3m$ integers to store the pointer arrays and an additional storage for B . In the worst case B includes m blocks, but usually this number is much smaller than m .

Other queries on statistics of objects can also be answered. For example, the top- K order element is that whose frequency is the K -th largest. We can just use $PtrB[m - K + 1]$ to locate the block. Then, the frequency and object id can be obtained with block's member and the $TtoF$ array. Especially, the median in frequency can be located with the $K/2$ -th element of the $PtrB$ array.

2.3 Possible Applications

For some mission-critical tasks (e.g., fraud detection) in big graphs, the efficiency to make decisions and infer interesting patterns is

crucial. As a result, recent years have witnessed an increasing interest in heuristic "shaving" algorithms with low computational complexity [9, 14]. A critical step of them is to keep finding low-degree nodes at every time of shaving nodes from a graph. Thus, S-Profile can be plugged into such algorithms for further speedup, by treating a node as an object and its degree as frequency.

Furthermore, S-Profile can also deal with a sliding window on a log stream, by letting every tuple (x_i, c_i) outdated from the window be a new incoming tuple (x_i, \bar{c}_i) , where \bar{c}_i is the opposite action of c_i .

3 EXPERIMENTAL RESULTS

We have implemented the proposed S-Profile algorithm and its counterparts in C++, and tested them with randomly generated log streams. The streams are produced with the following steps. We first randomly generate an "add" or "remove" action, with 70% and 30% probabilities respectively. Then, for each "add" action we randomly choose an object id according to a probability distribution (called *posPDF*). For each "remove" action another distribution (called *negPDF*) is used to randomly choose an object id. With this procedure, we obtained three test log streams:

- Stream1: both *posPDF* and *negPDF* are uniform random distribution on $[1, m]$.
- Stream2: both *posPDF* and *negPDF* are normal distributions with $\mu = 2m/3, m/3$ and $\sigma = m/6, m/6$, respectively.
- Stream3: *posPDF* is a normal distribution ($\mu = 4m/5, \sigma = m$), while *negPDF* is a lognormal distribution ($\mu = 3m/5, \sigma = m$).

In the following subsections, we first compare the proposed S-Profile with the heap based approach, for updating the mode and frequency. Then, the comparison with the balanced tree is presented for calculating the median. All experiments are carried out on a Linux machine with Intel Xeon E5-2630 CPUs (2.30 GHz). The CPU time (in second) of different algorithms are reported.

3.1 Comparison with the Heap

Heap is a kind of binary tree where the value in parent node must be larger or equal to the values in its children. Used to maintain the sorted frequency array, it is easy to obtain the mode (the root has the largest frequency). Notice that the balanced tree is inferior to the heap for calculating the mode.

In the experiment, a log stream is the input and the tested algorithm calculates the mode every time a tuple arrives. In Figure 3, we show the CPU times consumed by the heap based method and our S-Profile for the three log streams with varied length. The x-axis means the number of processed tuples (n). From the

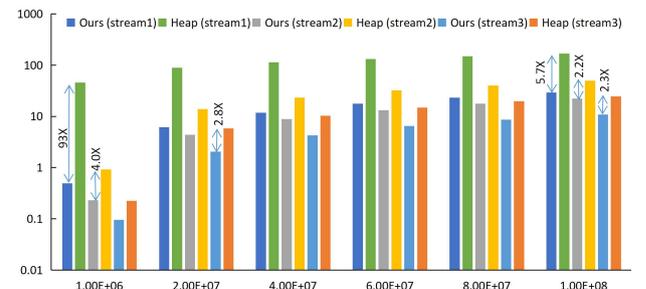


Figure 3: CPU times (in second) of the heap based method and our method for calculating the mode for n times ($m = 10^8$).

results we see that our method is at least 2.2X faster than the heap based method. Another experiment is carried out with fixed $n = 10^8$ and varied m . The results shown in Figure 4 also reveal that our S-Profile is at least 2X faster.

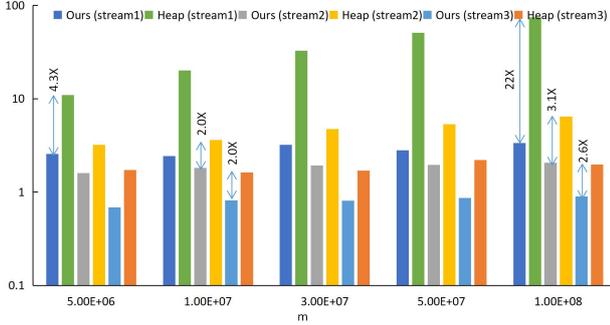


Figure 4: CPU times (in second) of the heap based method and our method for calculating the mode for $n = 10^8$ times.

For different kinds of log streams, the performance of the heap based method varies a lot. For the worst case updating the heap needs $O(\log m)$ time, despite this rarely happens in our tested streams. On the contrary, S-Profile needs $O(1)$ time for updating the data structure. This advantage is validated by the rather flat trend shown in Figure 5.

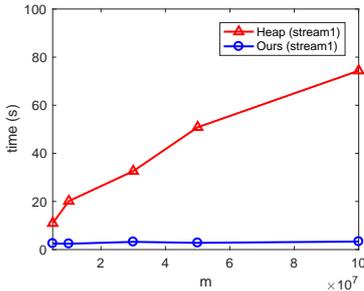


Figure 5: The trends of CPU time for varied m ($n = 10^8$).

It should be emphasized that, in addition to the speedup to the heap based method, our S-Profile possesses the advantage of wider applicability. Our method is not restricted to calculating the mode and corresponding frequency. As it well profiles the sorted frequency array, with it answering the queries on top-K and other statistics of objects is trivial and fast.

3.2 Comparison with the Balanced Tree

The proposed S-Profile can also calculate the median of the dynamic array. We compare it with the balanced tree based method implemented in the GNU C++ PBDS [16], which is more efficient than our implementation of balanced tree. The trends of CPU time are shown in Figure 6. They show that the runtime of the proposed S-Profile increases much less than that of the balanced tree based method when m increases. We can observe that the time of S-Profile is linearly depends on n , the number of modifications on array A , and hardly varies with different m . On the contrary, the balanced tree based method exhibits superlinear increase whether with n or m . Overall, the test results show that S-Profile is from 13X to 452X faster than the balanced tree based method on updating the median of a dynamic array.

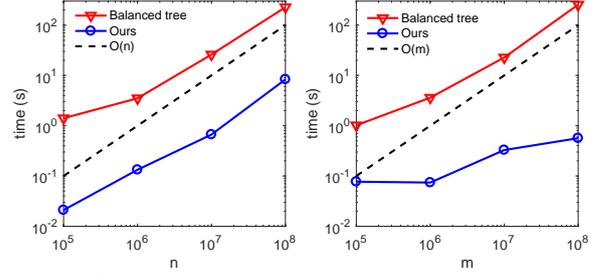


Figure 6: Comparison of the balanced tree based method and our method for calculating the median. Left: CPU time vs. n ($m = 10^6$). Right: CPU time vs. m ($n = 10^6$).

4 CONCLUSIONS

We propose an accurate algorithm, S-Profile, to fast keep profiling the dynamic array from online systems. It has the following advantages:

- **Optimal efficiency:** S-Profile needs $O(1)$ time complexity for every updating of a dynamic array, and totally linear complexity in memory.
- **Querying Statistics:** With profiling, at any time we can answer the statistical queries in a trivial and fast way.
- **Applicable:** S-Profile can be plugged into most of log streams, and heuristic graph mining algorithms.

ACKNOWLEDGMENTS

This work is supported by NSFC under grant Nos. 61872206 and 61772498, and in part by Beijing NSF No. 4172059.

REFERENCES

- [1] A. Arasu and G. S. Manku. 2004. Approximate counts and quantiles over sliding windows. In *Proceedings of the 23 ACM Symposium on Principles of Database Systems (PODS)*. 286–296.
- [2] B. Babcock, M. Datar, and R. Motwani. 2002. Sampling from a moving window over streaming data. In *Proceedings of the 13th annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*. 633–634.
- [3] R. S. Boyer and J. S. Moore. 1991. MJRTY—A fast majority vote algorithm. In *Automated Reasoning*. Springer, 105–117.
- [4] T. M. Chan, S. Durocher, K. G. Larsen, J. Morrison, and B. T. Wilkinson. 2014. Linear-space data structures for range mode query in arrays. *Theory of Computing Systems* 55, 4 (2014), 719–741.
- [5] M. Datar, A. Gionis, P. Indyk, and R. Motwani. 2002. Maintaining stream statistics over sliding windows. *SIAM J. Comput.* 31, 6 (2002), 1794–1813.
- [6] M. Dietz and G. Pernul. 2018. Big log data stream processing: Adapting an anomaly detection technique. In *International Conference on Database and Expert Systems Applications*. 159–166.
- [7] D. Dobkin and J. I. Munro. 1980. Determining the mode. *Theoretical Computer Science* 12, 3 (1980), 255–263.
- [8] P. B. Gibbons and S. Tirthapura. 2002. Distributed streams algorithms for sliding windows. In *Proceedings of the 14th annual ACM Symposium on Parallel Algorithms and Architectures (SPAA)*. 63–72.
- [9] B. Hooi, H. A. Song, A. Beutel, N. Shah, K. Shin, and C. Faloutsos. 2016. Fraudar: Bounding graph fraud in the face of camouflage. In *Proceedings of International Conference on Knowledge Discovery and Data Mining (KDD)*. 895–904.
- [10] D. Krizanc, P. Morin, and M. Smid. 2005. Range mode and range median queries on lists and trees. *Nordic Journal of Computing* 12, 1 (2005), 1–17.
- [11] X. Lin, H. Lu, J. Xu, and J. X. Yu. 2004. Continuously maintaining quantile summaries of the most recent n elements over a data stream. In *Proceedings of the 20th International Conference on Data Engineering (ICDE)*. 362.
- [12] A. Lubiwi and A. Rácz. 1991. A lower bound for the integer element distinctness problem. *Information and Computation* 94, 1 (1991), 83–92.
- [13] H. Petersen and S. Grabowski. 2009. Range mode and range median queries in constant time and sub-quadratic space. *Inform. Process. Lett.* 109, 4 (2009), 225–228.
- [14] K. Shin, B. Hooi, J. Kim, and C. Faloutsos. 2017. DenseAlert: Incremental dense-subtensor detection in tensor streams. In *Proceedings of International Conference on Knowledge Discovery and Data Mining (KDD)*. 1057–1066.
- [15] J. M. Steele and A. C. Yao. 1982. Lower bounds for algebraic decision trees. *Journal of Algorithms* 3, 1 (1982), 1–8.
- [16] A. Tavory, V. Dreizin, and B. Kosnik. [n. d.]. Policy-Based Data Structures. https://gcc.gnu.org/onlinedocs/libstdc++/ext/pb_ds/