

Stream Statistics Over Sliding Window

Anil Maheshwari

anil@scs.carleton.ca
School of Computer Science
Carleton University
Canada

Introduction

Algorithm

Sum Problem

Trends

References

Introduction

Main Problem

The input is an endless stream of binary bits. At any time, among the last N bits received, we are interested in queries that seek an (approximate) count of the number of 1's in the stream among the last k bits, where $k \leq N$.

Result: A data structure of size $O(\frac{1}{\epsilon} \log^2 N)$ that can report the approximate count of the number of 1s within a factor of $1 \pm \epsilon$

Reference: Maintaining stream statistics over sliding windows by Datar, Gionis, Indyk, and Motwani, SIAM JI. Computing 2002

Variants

1. A stream of positive numbers. The query consists of a value $k \in \{1, \dots, N\}$, and we want to know the (approximate) sum of the last k numbers in the stream.
2. A stream consisting of numbers from the set $\{-1, 0, +1\}$. We want to maintain the sum of last N numbers of the stream.
(Requires $\Omega(N)$ bits of storage to approximate the sum that is within a constant factor of the exact sum.)
3. What are the most popular movies in the last week?
4. What is trending in the last week?
5. ...

Main Problem

Main Problem

Report an approximate count of the number of 1's in the stream of binary bits among the last k bits, where $k \leq N$.

What about Exact Count?

Algorithm

Algorithm uses two data structures:

Time Stamps: To track the most recent N bits.

Buckets: $O(\log N)$ buckets maintain the 1's among the latest N bits.

Update

1 1 0 1 1 0 1 1 0 1 0 1 1 1 0 0

1 1 0 1 1 0 1 1 0 1 0 1 1 1 0 0 0

1 1 0 1 1 0 1 1 0 1 0 1 1 1 0 0 0 1

1 1 0 1 1 0 1 1 0 1 0 1 1 1 0 0 0 1 1

Complexity Analysis

1. **Space:** $O(\log^2 N)$ bits
2. **Total Time (per update):** $O(\log N)$

Query Problem

For any query value $k \in \{1, \dots, N\}$, report an approximate count of the number of 1's among the latest k bits of the stream.

1. Initialize $C := 0$
2. Traverse buckets from right to left. For each bucket of type B_i that is encountered in the traversal:
 - 2.1 B_i is completely contained in the window:
 $C := C + 2^i$
 - 2.2 B_i is completely outside the window:
 C remains unchanged
 - 2.3 Partially overlaps the window:
 $C := C + \frac{2^i}{2}$
3. Report C as an approximate count

2-factor approximation

Let C^* be the true count of number of 1s in the query window of size k .

Then, $\frac{1}{2} \leq \frac{C}{C^*} \leq 2$.

Proof: Let B_l be the last bucket type that partially overlaps the query window. Then, anywhere from $2^0 = 1$ to 2^l 1s may contribute to the count C^* , whereas for C we take the contribution of this bucket to be $\frac{2^l}{2} = 2^{l-1}$. Thus we may incur an error.

Note that all the previous buckets contribution is correctly accounted in the sum C . Let that contribution be α , and we know that $\alpha \geq \sum_{i=0}^{l-1} 2^i = 2^l - 1$ for C .

Note that $C = \alpha + 2^{l-1}$, $C^* \geq \alpha + 1$, and $\alpha \geq 2^l - 1$.

Thus $\frac{C}{C^*} \leq \frac{\alpha + 2^{l-1}}{\alpha + 1} \leq 1 + \frac{1}{2}$.

Similarly, $C^* \leq \alpha + 2^l$, and thus $\frac{C}{C^*} \geq \frac{\alpha + 2^{l-1}}{\alpha + 2^l} \geq 1 - \frac{1}{2}$.

□

Improvements

Let $r > 2$ be an integer parameter.

Maintain $r - 1$ or r copies of B_i for each $i \geq 1$

(B_0 and the largest bucket may have fewer than $r - 1$ copies)

At any time we exceed r copies of any type of buckets, we take the oldest two buckets and merge them to form a new bucket of the next size.

Answer queries as before.

Claim

For this setting, we have $1 - \frac{1}{r-1} \leq \frac{C}{C^*} \leq 1 + \frac{1}{r-1}$.

If $r = 1 + \frac{1}{\epsilon}$, we obtain a data structure of size $O(\frac{1}{\epsilon} \log^2 N)$ that approximates the count of the number of 1s within a factor of $1 \pm \epsilon$.

$$\text{True Count} \geq 1 + (r - 1) \sum_{i=1}^{j-1} 2^i$$

$$\text{Error} \leq 2^{j-1} - 1$$

$$\text{Therefore, } \frac{2^{j-1} - 1}{1 + (r-1) \sum_{i=1}^{j-1} 2^i} \leq \frac{1}{r-1}$$

Sum Problem

Computation of Sum

The Sum Problem

A stream of positive numbers. The query consists of a value $k \in \{1, \dots, N\}$, and we want to know the (approximate) sum of the last k numbers in the stream.

5 7 2 3 9 4 1 6 11 2 4 3

Approach I: Computation of Sum

If the next number in the stream is x , insert x 1's in the stream

5 7 2 3 9 4 1 6 11 2 4 3

Approach II: Computation of Sum

Assuming d -bit numbers. For each bit position i , maintain a stream. Let C_i be the value of approximate number of 1's in the stream i . Report approximate

sum as $\sum_{i=0}^{d-1} 2^i C_i$

	5	7	2	3	9	4	1	6	11	2	4	3
2^3	0	0	0	0	1	0	0	0	1	0	0	0
2^2	1	1	0	0	0	1	0	1	0	0	1	0
2^1	0	1	1	1	0	0	0	1	1	1	0	1
2^0	1	1	0	1	1	0	1	0	1	0	0	1

Trends

What is Trending?

Among the last 10^{12} movie tickets sold, list all *popular* movies?

Let $c := 10^{-3}$. Maintain (decaying) scores for movies whose threshold is at least $\tau \in (0, 1)$. For each new sale of ticket (say for Movie M):

1. For each movie whose score is being maintained, its new score is reduced by a factor of $(1 - c)$
2. If we have the score of M , add 1 to that score. Otherwise, create a new score for M and initialize it to 1
3. Remove any score that falls below τ

Questions

1. What is sum of all scores at any point of time?

Show that the sum total of all the scores is $\frac{1}{c}$

2. How many scores are maintained at any given time?

Assume $\tau = \frac{1}{2}$.

Show that at most $\frac{2}{c}$ movies are maintained, each with a score $\geq \tau$.

3. What are changes if $\tau = \frac{1}{3}$.

Variants

1. Min/Max
2. Stream with \pm numbers
3. Lower Bounds: Results are more-or-less optimal up to constant factors
4. ...

References

Main References:

1. Maintaining stream statistics over sliding windows, by Datar, Gionis, Indyk, and Motwani, SIAM J. Computing 2002.
2. Chapter in MMDS book (mmds.org)
3. Chapter on Data Streams in My Notes on Topics in Algorithm Design