

On the Amortized Complexity of Approximate Counting

Ishaq Aden-Ali   




University of California, Berkeley, CA, USA

Yanjun Han   

New York University, NY, USA

Jelani Nelson   

University of California, Berkeley, CA, USA

Huacheng Yu   

Princeton University, NJ, USA

Abstract

Naively storing a counter up to value n would require $\Omega(\log n)$ bits of memory. Nelson and Yu [9], following work of Morris [8], showed that if the query answers need only be $(1 + \epsilon)$ -approximate with probability at least $1 - \delta$, then $O(\log \log n + \log \log(1/\delta) + \log(1/\epsilon))$ bits suffice, and in fact this bound is tight. Morris' original motivation for studying this problem though, as well as modern applications, require not only maintaining one counter, but rather k counters for k large. This motivates the following question: for k large, can k counters be simultaneously maintained using asymptotically less memory than k times the cost of an individual counter? That is to say, does this problem benefit from an improved *amortized* space complexity bound?

We answer this question in the negative. Specifically, we prove a lower bound for nearly the full range of parameters showing that, in terms of memory usage, there is no asymptotic benefit possible via amortization when storing multiple counters. Our main proof utilizes a certain notion of “information cost” recently introduced by Braverman, Garg and Woodruff [2] to prove lower bounds for streaming algorithms.

2012 ACM Subject Classification Theory of computation → Lower bounds and information complexity

Keywords and phrases streaming, approximate counting, information complexity, lower bounds

Digital Object Identifier 10.4230/LIPIcs.APPROX/RANDOM.2024.33

Category RANDOM

Funding *Ishaq Aden-Ali*: supported by ONR DORECG award N00014-17-1-2127.

Yanjun Han: work done while at UC Berkeley, supported by the Berkeley-Simons Research Fellowship and Norbert Wiener postdoctoral fellowship.

Jelani Nelson: supported by NSF award CCF-1951384, ONR grant N00014-18-1-2562, and ONR DORECG award N00014-17-1-2127.

Acknowledgements We thank Sidhanth Mohanty for very enlightening discussions on unpredictable paths, half Cauchy random variables, and stochastic processes in general that ultimately led to the discovery of the first version of our main lower bound. We also thank Mark Sellke for answering a certain question regarding stochastic processes. Lastly, we thank Greg Valiant for raising the question of the amortized space complexity of approximate counting.

1 Introduction

Maintaining a counter subject to increments is one of the most basic data structural problems in computer science. If the counter value stays below n , then the naive solution of maintaining the counter explicitly consumes $O(\log n)$ bits of memory. In 1978, Morris devised a new



© Ishaq Aden-Ali, Yanjun Han, Jelani Nelson, and Huacheng Yu;
licensed under Creative Commons License CC-BY 4.0

Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2024).

Editors: Amit Kumar and Noga Ron-Zewi; Article No. 33; pp. 33:1–33:17



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

approximate counting algorithm which could maintain such a dynamic counter using much fewer bits in expectation [8]. His algorithm uses a number of random bits which itself is a random variable, whose expectation is $O(\log \log n + \log(1/\epsilon) + \log(1/\delta))$. At query time, this *Morris Counter* returns the correct counter value up to $(1+\epsilon)$ -multiplicative error with failure probability at most δ . Recently, Nelson and Yu [9] showed that a slight alteration to Morris' algorithm improves the dependence on δ , consuming only $O(\log \log n + \log(1/\epsilon) + \log \log(1/\delta))$ bits of memory. They also proved that this bound is asymptotically optimal.

Morris' original motivation for developing his approximate counter in the mid 1970s was, for each of 26^3 possible trigrams (sequences of three characters), to count the number of occurrences of that trigram amongst words in some text; these counts were then used by **typo**, an early Unix spellchecker (see more on this history in a survey by Lumbroso [7]). Approximate counters still find use in modern-day applications. For example, Redis is a highly popular in-memory database which is often used as a cache, and one built-in cache replacement policy is Least Frequently Used. To implement LFU, each item stored in cache must have an associated counter, counting the number of queries to that item over some time window. These counters are currently implemented in Redis using a variant of the Morris Counter [6].

Both the original application of Morris as well as modern-day applications of approximate counting, such as that described in Redis, need not store just one counter but many. Even in the 1970s, Morris' computer would have had enough memory to count the number of occurrences of a single trigram, but the difficulty was keeping counts of all 26^3 counters in memory simultaneously. Similarly, modern-day computers and mobile devices clearly also contain enough memory to store a single counter for any reasonable counter value, but one may wish to minimize the space per counter in an application where many counters must be stored simultaneously, as in the case of Redis being used to cache a large number of items. Thus even the original motivation of Morris inspires the following question:

*What is the **amortized** space complexity of approximate counting? Specifically, for k large, is it possible to maintain k approximate counters using less memory than k times that of an individual approximate counter?*

Typically such lower bounds are obtained via a direct sum argument in combination with an information complexity lower bound [3]. Specifically, a common way to carry out proving such lower bounds in the streaming context is as follows [1]: first, devise a communication game that you believe is hard (requires a lot of communication) to solve, and which reduces to the streaming problem (i.e., a low-memory streaming algorithm would imply a low-communication protocol). Suppose for illustration that this communication game is one-way, with Alice receiving input X who sends a single message Π to Bob who holds input Y . Here we use distributional complexity, in which (X, Y) is drawn from some distribution \mathcal{D} , and the goal is to devise a protocol in which Bob computes some $f(X, Y)$ with success probability at least $1 - \delta$. Now, the so-called *external information complexity* of Π is defined as $I(\Pi; X, Y)$, which is at most $|\Pi|$ and thus it suffices to lower bound $I(\Pi; X, Y)$. Once one establishes a lower bound on this quantity, then for a sequence $\{(X_i, Y_i)\}_{i=1}^k$ of k of independently drawn such inputs, we have the following inequality [3, Theorem 1.5]: $I(\Pi; X_1, Y_1, X_2, Y_2, \dots, X_k, Y_k) \geq \sum_{i=1}^k I(\Pi; X_i, Y_i)$. Note though that if we want a lower bound for Π solving the k -fold problem, then we cannot immediately invoke the single-instance lower bound to bound each summand on the right hand side (since this Π is a protocol solving the k -fold problem, not a single instance!). One then typically completes the argument by showing how to efficiently embed the single-instance communication game into that of k parallel instances, to then lift the single-instance lower bound to a lower bound for the k -fold problem as outlined above.

The trouble with obtaining our desired lower bound for approximate counting via the above paradigm is that the memory lower bound of [9] was not proven via information complexity, and thus it is not clear how to execute the above direct sum argument. Rather, the lower bound proof there followed an approach similar to the proof of the pumping lemma for regular languages [10].

Our Contribution. We answer the above question in the negative for nearly the full range of parameters (for technical reasons in our analysis, we require the restriction $\epsilon < 1/\log \log(1/\delta)$). Specifically, for this range of ϵ we prove a strong memory lower bound demonstrating that there is no asymptotic benefit from amortization when maintaining multiple approximate counters. Our main approach is to first establish a novel information complexity type lower bound for a single approximate counter, which we then lift to multiple counters using a standard direct sum argument, by constructing an embedding of the single-counter problem into the k -fold problem.

Below we formally define the problem we solve, then state our main result. The approximate counting problem for a single counter is to maintain a multiplicative approximation to a count N (initially 0) that undergoes a sequence of increment operations. Specifically, the algorithm should support two operations: `increment()` increments N by 1, and `query()` returns a value \hat{N} such that $\mathbb{P}(|\hat{N} - N| \geq \epsilon N) \leq \delta$; that is, at any point in time the data structure should be able to provide a $(1 + \epsilon)$ multiplicative approximation with probability at least $1 - \delta$. Our generalization to approximating k counters is as follows:

► **Definition 1** ((k, ϵ, δ) -approximate counter). *Let N_1, \dots, N_k be k counters all initially set to 0. We say a randomized algorithm \mathcal{A} is a (k, ϵ, δ) -approximate counter if it supports two operations: `increment(i)` increments N_i by 1, and `query(i)` returns a value \hat{N}_i such that*

$$\mathbb{P}\left(|\hat{N}_i - N_i| \geq \epsilon N_i\right) \leq \delta.$$

When $k = 1$, we will simply call this an (ϵ, δ) -approximate counter.

The following is then our main theorem, when here and henceforth we use n in lower bounds to denote an upper bound on the number of `increment` operations:

► **Theorem 2.** *For any $\delta < c_1$ and $\epsilon < c_2/\log \log(1/\delta)$, if $k < c_3 n$ then after a sequence of at most n updates any (k, ϵ, δ) -counter must use $\Omega(k \min\{\log(n/k), \log \log(n/k) + \log(1/\epsilon) + \log \log(1/\delta)\})$ bits of memory in expectation. Else if $k > c_3 n$, then a lower bound of $\Omega(n \log(2k/(c_3 n)))$ holds; $c_1, c_2, c_3 \geq 0$ are universal constants. Furthermore, for both ranges of k , these lower bounds are tight up to constant factors.*

1.1 Preliminaries and notation

We write $X \sim \mathcal{D}$ to represent a random variable X sampled from the probability distribution \mathcal{D} . If the distribution of X has not been explicitly defined, we write P_X to the corresponding probability distribution of X . For two probability distributions P and Q defined on the same domain, we write $\text{TV}(P, Q) = \int |dP - dQ|$ to be their total variation distance and $\text{KL}(P||Q) = \int dP \log(dP/dQ)$ their Kullback-Leibler divergence. We frequently make use of Pinsker's inequality $\text{TV}(P, Q) \leq \sqrt{\text{KL}(P||Q)}/2$. For random variables X and Y , we write $H(X) = \mathbb{E}_{x \sim P_X}[\log(1/P_X(x))]$ to denote the (Shannon) entropy and $H(X | Y) = \mathbb{E}_{(x,y) \sim P_{XY}}[\log(1/P_{X|Y}(x | y))]$ to be the conditional (Shannon) entropy. The mutual

33:4 On the Amortized Complexity of Approximate Counting

information between two random variables X and Y is $I(X; Y) := H(X) - H(X | Y) = H(Y) - H(Y | X)$. The conditional mutual information is $I(X; Y | Z) := H(X | Z) - H(X | Y, Z)$. We frequently make use of the following inequality:

$$I(X; Y | Z) = \mathbb{E}_{X, Z} [\text{KL}(P_{Y|Z} \| P_{Y|X, Z})].$$

We also use the following well known facts about conditional mutual information [4]:

► **Proposition 3** (Chain rule). *For random variables X_1, X_2, Y, Z we have*

$$I(X_1, X_2; Y | Z) = I(X_1; Y | Z, X_2) + I(X_2; Y | Z).$$

► **Proposition 4** (Superadditivity). *Let X_1, \dots, X_n, Y and Z be random variables such that X_1, \dots, X_n are conditionally independent given Z . Then*

$$\sum_{i=1}^n I(X_i; Y | Z) \leq I(X_1, \dots, X_n; Y | Z).$$

Lastly, $c, c_1, c_2, \dots > 0$ represent universal constants that may change from statement to statement.

1.2 Proof overview

Now we present an overview of our lower bound proof. It turns out that the terms $\Omega(k(\log \log(n/k) + \log(1/\epsilon)))$ in the lower bound can be proved using a similar argument to the single counter case, which holds even in the offline setting, and the main challenge is to prove the dependence on δ , $\Omega(k \log \log(1/\delta))$. The previous proof, for single counter, of $\Omega(\log \log(1/\delta))$ is based on a pumping-lemma argument, which crucially uses the fact that all updates are exactly the same, i.e., incrementing the counter. However, this no longer holds with multiple counters – we may increment any one of the k counters each time, and there are k different updates we can perform. The previous argument fails to generalize.

As we mentioned in the introduction, we first (re)prove an *information theoretic* lower bound for single counter, then apply the *superadditivity of mutual information* for independent variables to obtain the direct-sum result. For simplicity, in this overview we will focus on the case where $\delta = 2^{-\Theta(n/k)}$, and prove an $\Omega(k \log \log(1/\delta)) = \Omega(k \log(n/k))$ lower bound when $\epsilon \leq O(1/\log(n/k))$. This case captures all the main ideas, and generalizing to the full lower bound is straightforward.


To facilitate the argument, we first slightly reformulate the problem:¹ Consider a stream with T batches of updates (think $T = (n/k)^{0.1}$), and in each batch i , the inputs are k nonnegative integers $X_{i,c}$ for $c \in [k]$, which are the number of increments we apply to each counter in this batch. Since the batch number i takes only $O(\log T)$ bits to maintain, we may assume without loss of generality that it is given to the algorithm for free. Clearly if there is an algorithm that approximately maintains k counters, this reformulation also admits a solution using the same space, provided $\sum_{i,c} X_{i,c} \leq n$. After the reformulation, the single counter problem simply has T nonnegative numbers X_i as the input stream, and we would like to approximate their sum using small space provided that the sum is at most poly T , again assuming that i is given to the algorithm for free.

We will design a hard distribution over the input streams, and analyze the failure probability and measure “information” with respect to this distribution. It is crucial that the notion of information cost we use for streaming algorithms is chosen carefully. The


¹ For simplicity, the overview uses slightly different notation and parameters than the actual proof.

information cost defined in [2] turns out to be the right measure for our problem. Fix a streaming algorithm, and let M_i be its memory state immediately after processing X_i (see Figure 1). When all numbers X_i are independent, which will be the case for our distribution, the *information cost* of this algorithm is defined as follows (see Figure 2)

$$\text{IC} := \sum_{i=1}^T \sum_{j=1}^i I(M_i; X_j | M_{j-1}).$$



■ **Figure 1** A depiction of the evolution of the memory state of a randomized streaming algorithm \mathcal{A} . X_1, \dots, X_n are the inputs the streaming algorithm receives and R_1, \dots, R_n are the independent random bits \mathcal{A} uses as it processes the stream. M_0 is the (possibly random) initial state and M_i is the state of \mathcal{A} after processing X_i . Note that M_i is a deterministic function of the previous memory state M_{i-1} , the i th input X_i , and the random bits R_i . When the inputs are random variables, this figure also depicts the dependence structure of the joint distribution of the random variables $(X_1, \dots, X_n, R_1, \dots, R_n, M_0, M_1, \dots, M_n)$.



■ **Figure 2** An illustration of a single inner sum of the information cost $\sum_{j \leq i} I(M_i; X_j | M_{j-1})$ for $i = 4$. To simplify the presentation we do not include the random bits that the streaming algorithm uses while processing the stream in this illustration.

As it was shown in [2], this quantity lower bounds $\sum_{i=1}^T |M_i|$, i.e., T times the memory consumption, and it satisfies the direct-sum property: solving k independent instances of the problem needs exactly k times IC. Thus, it suffices to prove an IC lower bound of $\Omega(T \log T) = \Omega(T \log(n/k))$ for a single approximate counter. Intuitively, such a lower bound means that the algorithm must spend $\Omega(1)$ bits remembering each bit of the sum (recall that we ensure $\sum X_i \leq \text{poly } T$).

Let us first focus on the lowest bit, i.e., the parity of the sum. Note that one does not have to know the lowest bit in order to return an approximation of the sum. Nevertheless, we will show that the algorithm has to constantly pay attention to this bit in order to output a good approximation with very high probability. To this end, let us consider the distribution where the X_i 's are independent and uniform in $\{0, 1\}$. Let us focus on the terms in IC with $j = i$, i.e., those of the form $I(M_i; X_i | M_{i-1})$. Intuitively, under this distribution, this term represents how much attention the algorithm is paying to the evolution of the parity at batch i . This is because X_i is simply the difference between the parities of the first $i - 1$ and first i inputs.

Suppose that $I(M_i; X_i \mid M_{i-1}) \ll 1$ for a constant fraction of i . Then, we can show that the algorithm will make an error of $\Omega(T)$ with at least δ probability. Roughly speaking, this is because for each such i , conditioned on M_i and M_{i-1} , the distribution of X_i is still close to uniform (as X_i is uniform conditioned on M_{i-1}). Therefore, if we condition on all M_0, M_1, \dots, M_T , most X_i can still take both values 0 or 1 with constant probability, and all X_i are still independent by the Markov property of the algorithm. In particular, by setting all these X_i 's to 0 or setting all to 1, we reach the same final memory state M_T , but in the two cases, the total sum differs by $\Omega(T)$. Since both happen with probability $2^{-O(T)} \gg \delta$ given the final memory state M_T , the algorithm must make an error of at least $\Omega(T)$ with probability at least δ .

We can extend this argument to any specific bit of the sum by considering a stream with $\Theta(T/B^l)$ independent random increments that are uniform in $\{0, B^l\}$ for some constant B and $l \in [L]$, where $L = \log_B T$. Our final hard distribution interleaves L such streams, which we call the scales. For each scale $l \in [L]$, we evenly spread the $\Theta(T/B^l)$ random increments in the whole stream with a gap of $\Theta(B^l)$ batches. Note that now the sum of all inputs is always at most $O(T \log T)$. For the sum in the definition of IC, we will only focus on L terms for each i : $I(M_i; X_{\lfloor i \rfloor_l} \mid M_{\lfloor i \rfloor_l - 1})$, where $X_{\lfloor i \rfloor_l}$ is the closest scale l input before M_i . If $\text{IC} \ll T \cdot L = O(T \log T)$, then there must exist some scale l^* such that $I(M_i; X_{\lfloor i \rfloor_{l^*}} \mid M_{\lfloor i \rfloor_{l^*} - 1}) \ll 1$ for most i . Then we apply the above argument, and show that conditioned on the memory states right before each scale- l^* input, most scale- l^* inputs can still take values 0 or B^{l^*} with constant probability, and the scale- l^* inputs are still independent. We further observe that for most of them, the inputs between $X_{\lfloor i \rfloor_{l^*}}$ and M_i are coming from the lower scales $l < l^*$. The standard deviation of their sum is much smaller than B^{l^*} , and we can show that they do not affect the sum by too much as we alter $X_{\lfloor i \rfloor_{l^*}}$. Thus, by setting these scale- l^* inputs to 0 or B^{l^*} , the entire sum will again differ by $\Omega(T)$, but they lead to the same final memory state, i.e., the algorithm does not distinguish between the two cases. Since the sum is $O(T \log T)$, such a difference is more than ϵ times the sum when $\epsilon < O(1/\log T)$. A more careful analysis shows that this happens with probability at least $2^{-O(T)} > \delta$, leading to a contradiction.

Discussion on the choice of the hard distribution

We note here that the independent uniform $\{0, 1\}$ distribution, by itself, is *not* hard for the information cost. One solution with low information cost is to divide X_i 's into blocks of size $S = \Theta(\epsilon^{-2} \log T)$, and maintain the exact sum within the current block. If *all* blocks have sums $(1 \pm \epsilon)S/2$, then we simply remember this fact and use $T/2$ as the final output. Otherwise, the algorithm finds a block whose sum is not in this range, then it maintains the exact sum for all future blocks. Since all previous blocks have sums $(1 \pm \epsilon)S/2$, in particular, $S/2$ is a $(1 + \epsilon)$ -approximation for them, the algorithm can also return a $(1 + \epsilon)$ -approximation of the whole sum. Now note that this case only happens with $1/\text{poly}(T)$ probability, and the total information cost is at most T times the expected memory usage. The expected memory is at most $O(\log(1/\epsilon) + \log \log T)$ bits for maintaining the sum in the current block, plus $O(T^{-\Theta(1)} \cdot \log T)$ bits in expectation for maintaining the entire exact sum. When $\epsilon = \Theta(1/\log T)$, the information cost is only $O(T \log \log T) \ll T \log T$.

The above strategy works since the sum of block is concentrated around the expectation, hence, we need extra space to maintain the exact sum only with very small probability. One can also show that the above strategy also applies to any i.i.d. distributions with some concentration.² Our hard distribution is inspired by the *discrete half-Cauchy distribution*,

² For example, it suffices to have finite $\mathbb{E}[|X|^{1.01}]$.

which has probability $\Theta(1/(x+1)^2)$ at integers $x \geq 0$. This distribution does not have an expectation, hence, there is no concentration for blocks of any size. We also have a more involved proof of our main result that uses this distribution instead; the main property that proof relies on is that for every W random variables, roughly one of them takes value $\Theta(W)$. The hard distribution we actually use in this paper is a bounded distribution that can be viewed as extracting this useful property of the half-Cauchy distribution, but which can be analyzed more simply.

2 Lower Bounds

In this section we will prove our lower bound for the space complexity of (k, ϵ, δ) -approximate counting. We split the proof into two parts. We first focus on proving the difficult part of the lower bound that depends on the failure probability δ . To do so, we derive a lower bound for the space complexity of (ϵ, δ) -approximate counting by appealing to an information theoretic argument. By using a good definition of information cost together with an appropriately chosen hard distribution, we can prove that any accurate algorithm remember a lot of information about many different parts of the stream, i.e. it incurs a high information cost. This immediately gives us a lower bound on the memory size. We then use this result to prove a space lowerbound for (k, ϵ, δ) -approximate counting via a direct sum argument. We conclude the section by proving the portion of the lower bound that depends on the approximation error ϵ and the total sum of all counters n by generalizing the argument used in the single counter case.

2.1 Information lower bound for a single counter

Recall that for a randomized streaming algorithm \mathcal{A} we define M_i to be the state of \mathcal{A} after processing the i th input X_i together with some additional independent random bits R_i , i.e. M_i is a deterministic function of M_{i-1} , X_i and R_i (equivalently, a deterministic function of $X_{\leq i}$ and $R_{\leq i}$). The following is a notion of information cost for streaming algorithms originally defined by Braverman, Garg, and Woodruff [2].

► **Definition 5.** *Let \mathcal{A} be a randomized streaming algorithm. Given a distribution \mathcal{D} over input sequences of length s , we define the information cost of algorithm \mathcal{A} on input $(X_1, \dots, X_s) \sim \mathcal{D}$ to be*

$$\text{IC}(\mathcal{A}, \mathcal{D}) := \sum_{i=1}^s \sum_{j=1}^i I(M_i; X_j \mid M_{j-1}).$$

The above definition of the information cost is motivated by the following chain of inequalities:

$$\begin{aligned} \mathbb{E}|M_i| &\geq H(M_i) \text{ (source coding theorem}^3) \\ &\geq I(M_i; X_{\leq i}, M_{< i}) \\ &= \sum_{j=1}^i I(M_i; X_j, M_{j-1} \mid X_{< j}, M_{< j-1}) \\ &= \sum_{j=1}^i I(M_i; M_{j-1} \mid X_{< j}, M_{< j-1}) + I(M_i; X_j \mid X_{< j}, M_{< j}) \\ &= \sum_{j=1}^i I(M_i; X_j \mid X_{< j}, M_{< j}) \\ &= \sum_{j=1}^i I(M_i; X_j \mid M_{j-1}). \end{aligned}$$

33:8 On the Amortized Complexity of Approximate Counting

This implies that $\text{IC}(\mathcal{A}; \mathcal{D}) \leq \sum_{i=1}^s \mathbb{E} |M_i|$. The main technical part of this paper is proving the following lower bound for a single counter using this notion of information cost.

► **Lemma 6.** *Let \mathcal{A} be a (ϵ, δ) -approximate counter with parameters $\delta \in (0, c_1)$ and $\epsilon \in (0, \frac{c_2}{\log \log(1/\delta)})$ that uses $|M|$ bits of space. There is a distribution \mathcal{D} over inputs such that the information cost of \mathcal{A} on $(X_1, \dots, X_n) \sim \mathcal{D}$ satisfies*

$$\text{IC}(\mathcal{A}; \mathcal{D}) = \Omega(n \log \log(1/\delta)),$$

where $\mathbb{P}[\sum_{i=1}^n X_i \leq n] = 1$ and $n \geq c_3 \log^{c_4}(1/\delta)$. This implies the space lower bound

$$\frac{1}{n} \sum_{i=1}^n \mathbb{E} |M_i| = \Omega(\log \log(1/\delta)) = \Omega(\min\{\log n, \log \log(1/\delta)\}).$$

Proof. We construct the distribution \mathcal{D} as follows: let $B \in \mathbb{N}$ be a large integer constant to be specified later, and T is the largest power of B such that $T < \log_{32}(1/\delta)$. When $\delta < c_1$ where $c_1 = c_1(B)$ is a sufficiently small constant, we have $T \geq B$ and so $L := \log_B T \in \mathbb{N}$. The distribution \mathcal{D} is an interleaving of L distributions \mathcal{D}_ℓ on L different scales, where for each scale $\ell \in [L]$, the distribution \mathcal{D}_ℓ is a product distribution $\prod_{i=1}^T \mathcal{D}_{\ell,i}$:

$$\text{under } \mathcal{D}_{\ell,i}, \quad Y_{\ell,i} \begin{cases} \sim \text{Unif}(\{0, B^\ell\}) & \text{if } B^\ell \text{ divides } (i-1), \\ = 0 & \text{otherwise.} \end{cases}$$

For notational simplicity, we also denote the non-zero entries of $Y_{\ell,i}$ by $Z_{\ell,j} = Y_{\ell, B^\ell(j-1)+1}$ for $j \in [T/B^\ell]$. The stream under distribution \mathcal{D} is then generated by interleaving the $Y_{\ell,i}$ terms to form the sequence $X^n = (Y_{1,1}, \dots, Y_{L,1}, Y_{1,2}, \dots, Y_{L,2}, \dots, Y_{L,T})$, where $n := TL$ is the length of the sequence. The total value of the counter under this stream is at most

$$\sum_{\ell=1}^L B^\ell \cdot \frac{T}{B^\ell} = n.$$

Finally, note that $n = TL \geq c_3 \log^{c_4}(1/\delta)$ for an appropriate choice of constants $c_3, c_4 \geq 0$.

Let \mathcal{A} be an (ϵ, δ) -approximate counter, and assume towards contradiction that $\text{IC}(\mathcal{A}; \mathcal{D}) < \alpha n L$ where α is a small constant. For any index $i \in [n]$ in the stream, let $[i]_\ell := B^\ell L \lfloor (i-1)/(B^\ell L) \rfloor + \ell$ be the index of the closest $Z_{\ell,j}$ to the left of X_i in the stream, i.e. $X_{[i]_\ell} = Z_{\ell,j}$. Similarly, for $j \in [T/B^\ell]$, we define $\lceil j \rceil_\ell := B^\ell L(j-1) + \ell$ to be the index of the j th non-zero entry of scale ℓ in the stream, i.e. $X_{\lceil j \rceil_\ell} = Z_{\ell,j}$. The definition of the information cost tells us that

$$\begin{aligned} \alpha n L &> \text{IC}(\mathcal{A}; \mathcal{D}) \\ &= \sum_{i=1}^n \sum_{j \leq i} I(M_i; X_j \mid M_{j-1}) \\ &\geq \sum_{i=1}^n \sum_{\ell=1}^L I(M_i; X_{[i]_\ell} \mid M_{\lceil [i]_\ell / B^\ell \rceil_\ell}) \end{aligned}$$

³ The source coding theorem holds for any prefix code. In general we may lose a factor of 2 in this inequality: we have both $\mathbb{E}[|M_i|] \geq H(M_i \mid |M_i|)$ and $\mathbb{E}[|M_i|] = \sum_{n \geq 1} n p_n = \sum_{n \geq 1} p_n \log(1/2^{-n}) = \sum_{n \geq 1} p_n \log(p_n/2^{-n}) + H(|M_i|) \geq H(|M_i|)$ for $p_n := \mathbb{P}(|M_i| = n)$; consequently $2\mathbb{E}[|M_i|] \geq H(M_i \mid |M_i|) + H(|M_i|) \geq H(M_i)$.

$$\begin{aligned}
&= \sum_{\ell=1}^L \sum_{j=1}^{T/B^\ell} \sum_{i=[j]_\ell}^{[j+1]_\ell-1} I(M_i; X_{[j]_\ell} \mid M_{[j]_\ell-1}) \\
&\geq \sum_{\ell=1}^L \sum_{j=1}^{T/B^\ell} B^\ell L \cdot I(M_{[j+1]_\ell-1}; X_{[j]_\ell} \mid M_{[j]_\ell-1}) \\
&\geq \min_{\ell \in [L]} B^\ell L^2 \sum_{j=1}^{T/B^\ell} I(M_{[j+1]_\ell-1}; X_{[j]_\ell} \mid M_{[j]_\ell-1}) \\
&= \min_{\ell \in [L]} B^\ell L^2 \sum_{j=1}^{T/B^\ell} I(M_{[j+1]_\ell-1}; Z_{\ell,j} \mid M_{[j]_\ell-1})
\end{aligned}$$

where the second last inequality is due to the data processing inequality. By Markov's inequality, there exist $\ell_0 \in [L]$ and $J_0 \subseteq [T/B^{\ell_0}]$ with $|J_0| \geq T/(2B^{\ell_0})$ and

$$I(M_{[j+1]_{\ell_0}-1}; Z_{\ell_0,j} \mid M_{[j]_{\ell_0}-1}) \leq 2\alpha, \quad \forall j \in J_0.$$

For ease of presentation, we abuse notation and write the above inequality as $I(M_j; Z_j \mid M_{j-1}) \leq 2\alpha$ for $j \in J_0$. We shall also keep in mind that the stream between M_{j-1} and M_j contains $(Z_j^<, Z_j, Z_j^>)$, where $Z_j^<$ is the collection of all non-zero inputs $\{Z_{\ell',j'}\}$ inside this window with a lower scale $\ell' < \ell$, and $Z_j^>$ is the counterpart with a higher scale $\ell' > \ell$ (see Figure 3).

Next we define several good events for the sake of analysis. The first good event $E_{j,1}$ characterizes the behavior of the contribution of $Z_j^<$ and is formally defined as

$$\mathbb{1}(E_{j,1}) := \mathbb{1} \left(\left| \text{sum}(Z_j^<) - \mathbb{E}[\text{sum}(Z_j^<) \mid M_{j-1}, M_j] \right| \leq \frac{B^{\ell_0}}{4} \right), \quad j \in J_0.$$

By Chebyshev's inequality, it is clear that

$$\begin{aligned}
\mathbb{E}_{M_{j-1}, M_j} [\mathbb{P}(E_{j,1}^c \mid M_{j-1}, M_j)] &\leq \mathbb{E}_{M_{j-1}, M_j} \left[\frac{\text{Var}(\text{sum}(Z_j^<) \mid M_{j-1}, M_j)}{(B^{\ell_0}/4)^2} \right] \\
&\leq \frac{\text{Var}(\text{sum}(Z_j^<))}{(B^{\ell_0}/4)^2} \\
&= \frac{16}{B^{2\ell_0}} \sum_{\ell < \ell_0} B^{\ell_0-\ell} \cdot \frac{B^{2\ell}}{4} \leq \frac{4}{B-1}.
\end{aligned}$$

Consequently, it holds that

$$\begin{aligned}
&\mathbb{E}_{M_{j-1}, M_j} [\text{TV}(P_{Z_j}, P_{Z_j \mid M_{j-1}, M_j, E_{j,1}})] \\
&= \mathbb{E}_{M_{j-1}, M_j} [\text{TV}(P_{Z_j \mid M_{j-1}}, P_{Z_j \mid M_{j-1}, M_j, E_{j,1}})] \\
&\leq \mathbb{E}_{M_{j-1}, M_j} [\text{TV}(P_{Z_j \mid M_{j-1}}, P_{Z_j \mid M_{j-1}, M_j}) + \text{TV}(P_{Z_j \mid M_{j-1}, M_j}, P_{Z_j \mid M_{j-1}, M_j, E_{j,1}})] \\
&= \mathbb{E}_{M_{j-1}, M_j} [\text{TV}(P_{Z_j \mid M_{j-1}}, P_{Z_j \mid M_{j-1}, M_j})] + \mathbb{E}_{M_{j-1}, M_j} [\mathbb{P}(E_{j,1}^c \mid M_{j-1}, M_j)] \\
&\leq \sqrt{\frac{1}{2} \mathbb{E}_{M_{j-1}, M_j} [\text{KL}(P_{Z_j \mid M_{j-1}} \parallel P_{Z_j \mid M_{j-1}, M_j})]} + \frac{4}{B-1} \\
&= \sqrt{I(Z_j; M_j \mid M_{j-1})} + \frac{4}{B-1} \\
&\leq \sqrt{\alpha} + \frac{4}{B-1},
\end{aligned}$$

33:10 On the Amortized Complexity of Approximate Counting

which can be made small by choosing $\alpha > 0$ small enough and $B \in \mathbb{N}$ large enough. Note that in the above inequality we have used the triangle inequality $\text{TV}(P, Q) \leq \text{TV}(P, R) + \text{TV}(Q, R)$, the conditioning relationship $\text{TV}(P, P|_E) = P(E^c)$, Pinsker's inequality $\text{TV}(P, Q) \leq \sqrt{\text{KL}(P||Q)}/2$, and Jensen's inequality $\mathbb{E}[\sqrt{X}] \leq \sqrt{\mathbb{E}[X]}$.

The next good event E_2 concerns the simultaneous occurrence of $\{E_{j,1}\}$ for a constant proportion of $j \in J_0$. Specifically, E_2 is the event that there exists some $J_1 \subseteq J_0$ such that:

1. $|J_1| \geq |J_0|/2 \geq T/(4B^{\ell_0})$;
2. event $E_{j,1}$ is true for all $j \in J_1$;
3. a small TV distance $\text{TV}(P_{Z_j}, P_{Z_j|M_{j-1}, M_j, E_{j,1}}) \leq 1/4$ (denoted by event $E_{j,2}$) for all $j \in J_1$.

Since $\{(Z_j^<, Z_j, Z_j^>)\}$ are conditionally independent given $\{M_j\}$,

$$\begin{aligned} & \mathbb{E}_{\{M_j\}} \left[\sum_{j \in J_0} \mathbb{1}(E_{j,1} \cap E_{j,2}) \right] \\ &= \sum_{j \in J_0} \mathbb{E}_{M_{j-1}, M_j} [\mathbb{1}(E_{j,1} \cap E_{j,2})] \\ &\geq \sum_{j \in J_0} \left(1 - \mathbb{E}_{M_{j-1}, M_j} [\mathbb{P}(E_{j,1}^c | M_{j-1}, M_j)] \right. \\ &\quad \left. - 4 \cdot \mathbb{E}_{M_{j-1}, M_j} [\text{TV}(P_{Z_j}, P_{Z_j|M_{j-1}, M_j, E_{j,1}})] \right) \\ &\geq \left(1 - \frac{4}{B-1} - 4 \left(\sqrt{\alpha} + \frac{4}{B-1} \right) \right) |J_0| \\ &\geq \frac{3}{4} |J_0|, \end{aligned}$$

by choosing α small enough and B large enough. Consequently, by Markov's inequality, we have $\mathbb{P}(E_2) \geq 1/2$ over the randomness of $\{M_j\}$ and $\{(Z_j^<, Z_j, Z_j^>)\}$.

Now we condition on E_2 and arrive at the desired contradiction. For a probability distribution P over the real line and $\Delta \geq 0$, define a quantity $f(P, \Delta)$ as follows:

$$f(P, \Delta) = \max\{\delta > 0 : \exists L \in \mathbb{R} \text{ such that } P((-\infty, L]) \geq \delta, P([L + \Delta, \infty)) \geq \delta\}.$$

Intuitively, a *small* $f(P, \Delta)$ implies that the distribution P assigns a lot of probability to *some* interval of length Δ . The following lemma summarizes some properties of $f(P, \Delta)$.

► **Lemma 7.** *Let P and Q be two probability distributions over \mathbb{R} , and $P \star Q$ denote their convolution. For $\Delta_1, \Delta_2, \Delta \geq 0$, it holds that*

$$\begin{aligned} f(P \star Q, \Delta_1 + \Delta_2) &\geq f(P, \Delta_1) f(Q, \Delta_2), \\ f(P \star Q, \Delta) &\geq f(P, \Delta)/2. \end{aligned}$$

Proof. For the first claim, suppose that

$$\begin{aligned} \min\{P((-\infty, L_1]), P([L_1 + \Delta_1, \infty))\} &\geq f(P, \Delta_1), \\ \min\{Q((-\infty, L_2]), Q([L_2 + \Delta_2, \infty))\} &\geq f(Q, \Delta_2). \end{aligned}$$

Then the first inequality follows from

$$\begin{aligned} P \star Q((-\infty, L_1 + L_2]) &\geq P((-\infty, L_1]) Q((-\infty, L_2]) \geq f(P, \Delta_1) f(Q, \Delta_2), \\ P \star Q([L_1 + L_2 + \Delta_1 + \Delta_2, \infty)) & \\ &\geq P([L_1 + \Delta_1, \infty)) Q([L_2 + \Delta_2, \infty)) \geq f(P, \Delta_1) f(Q, \Delta_2). \end{aligned}$$

The second claim is a direct consequence of the first claim and $f(Q, 0) \geq 1/2$. ◀

To apply Lemma 7, we consider the conditional distribution $P_{S|\{M_j\}}$, where $S = \sum_{i=1}^n X_i$ is the total number of counter updates, and $\{M_j\}$ are the memory states at scale ℓ_0 defined before. Since the counter \mathcal{A} is (ϵ, δ) -approximate, in expectation $P_{S|\{M_j\}}$ must have probability mass at least $1 - \delta$ on an interval of size $2\epsilon n$. This implies that

$$\mathbb{E}_{\{M_j\}} [f(P_{S|\{M_j\}}, 2\epsilon n)] \leq \delta. \quad (1)$$

On the other hand, since $\{(Z_j^<, Z_j, Z_j^>)\}$ are conditionally independent given $\{M_j\}$, we may invoke Lemma 7 (first part for J_1 and second part for J_1^c) to arrive at

$$\begin{aligned} \mathbb{E}_{\{M_j\}} [f(P_{S|\{M_j\}}, 2\epsilon n)] &\geq \mathbb{P}(E_2) \mathbb{E}_{\{M_j\}} [f(P_{S|\{M_j\}}, 2\epsilon n) \mid E_2] \\ &\geq \frac{1}{2} \cdot \mathbb{E}_{\{M_j\}} \left[\prod_{j \in J_1} f(P_{\text{sum}(Z_j^<, Z_j, Z_j^>)} | M_{j-1}, M_j, \frac{2\epsilon n}{|J_1|}, \left(\frac{1}{2}\right)^{\frac{T}{B^{\ell_0}} - |J_1|} \mid E_2 \right] \\ &\geq \frac{1}{2T} \cdot \mathbb{E}_{\{M_j\}} \left[\prod_{j \in J_1} f(P_{\text{sum}(Z_j^<, Z_j, Z_j^>)} | M_{j-1}, M_j, 8\epsilon L B^{\ell_0}) \mid E_2 \right]. \end{aligned}$$

Conditioned on the event E_2 , the event $E_{j,1}$ implies that the deviation of $\text{sum}(Z_j^<)$ to its posterior mean is at most $B^{\ell_0}/4$, and the event $E_{j,2}$ implies that the posterior (marginal) distribution of Z_j puts at least $1/4$ probability mass on both 0 and B^{ℓ_0} . Moreover, $\text{sum}(Z_j^>)$ is always an integral multiple of B^{ℓ_0+1} . Now we prove that

$$f(P_{\text{sum}(Z_j^<, Z_j, Z_j^>)} | M_{j-1}, M_j, E_{j,1}, E_{j,2}, \frac{B^{\ell_0}}{3}) \geq \frac{1}{16}.$$

We distinguish into two cases:

1. Case I: $\mathbb{P}(\text{sum}(Z_j^>) \neq \text{median}(\text{sum}(Z_j^>))) \geq 1/8$. As $\text{sum}(Z_j^>)$ is always an integral multiple of B^{ℓ_0+1} , this implies that with probability at least $\frac{1}{8}$, $|\text{median}(\text{sum}(Z_j^>)) - \text{sum}(Z_j^>)|$ is at least $B^{\ell_0+1}/2$. By symmetry, without loss of generality we may assume that $\text{median}(\text{sum}(Z_j^>)) - \text{sum}(Z_j^>)$ is at least $B^{\ell_0+1}/2$ with probability at least $(1/8)/2 = 1/16$. Moreover, the range of Z_j is B^{ℓ_0} , and the range of $\text{sum}(Z_j^<)$ is at most $B^{\ell_0}/2$ under $E_{j,1}$. Consider the interval

$$\left[\text{median}(\text{sum}(Z_j^<, Z_j, Z_j^>)), \text{median}(\text{sum}(Z_j^<, Z_j, Z_j^>)) + \frac{B^{\ell_0}}{3} \right]$$

of length $B^{\ell_0}/3$, it is clear that

$$\begin{aligned} \mathbb{P}(\text{sum}(Z_j^<, Z_j, Z_j^>) \leq \text{median}(\text{sum}(Z_j^<, Z_j, Z_j^>))) &\geq \frac{1}{2}, \\ \mathbb{P}\left(\text{sum}(Z_j^<, Z_j, Z_j^>) \geq \text{median}(\text{sum}(Z_j^<, Z_j, Z_j^>)) + \frac{B^{\ell_0}}{3}\right) \\ &\geq \mathbb{P}\left(\text{sum}(Z_j^>) \geq \text{median}(\text{sum}(Z_j^>)) + \frac{3B^{\ell_0}}{2} + \frac{B^{\ell_0}}{3}\right) \geq \frac{1}{16}, \end{aligned}$$

as long as $B^{\ell_0+1}/2 \geq 11B^{\ell_0}/6$, or equivalently $B \geq 4$.

2. Case II: $\mathbb{P}(\text{sum}(Z_j^>) = \text{median}(\text{sum}(Z_j^>))) \geq 7/8$. In this case, we argue that each of the following two probabilities is at least $1/16$:

$$\mathbb{P}\left(\text{sum}(Z_j^<) + Z_j + \text{sum}(Z_j^>) \leq \mathbb{E}[\text{sum}(Z_j^<)] + \frac{B^{\ell_0}}{3} + \text{median}(\text{sum}(Z_j^>))\right),$$

33:12 On the Amortized Complexity of Approximate Counting

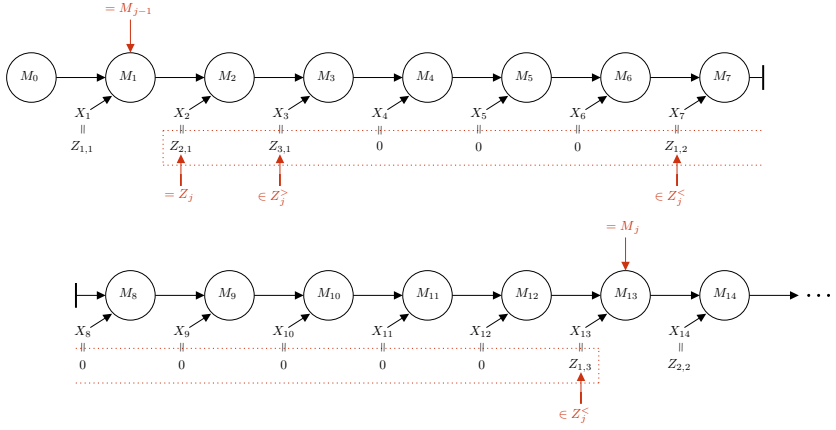
$$\mathbb{P} \left(\text{sum}(Z_j^<) + Z_j + \text{sum}(Z_j^>) \geq \mathbb{E}[\text{sum}(Z_j^<)] + \frac{2B^{\ell_0}}{3} + \text{median}(\text{sum}(Z_j^>)) \right).$$

By symmetry we only prove the first claim, where the event occurs whenever $Z_j = 0$ and $\text{sum}(Z_j^>)$ is equal to its median. By the union bound, this happens with probability at least $1 - (1/8 + 3/4) = 1/8$.

Consequently, as long as $\epsilon \leq 1/(24L) = O(1/\log \log(1/\delta))$,

$$\mathbb{E}_{\{M_j\}} [f(P_{S|\{M_j\}}, 2\epsilon n)] \geq \left(\frac{1}{32} \right)^T. \quad (2)$$

However, as $T < \log_{32}(1/\delta)$, inequalities (1) and (2) are contradictory to each other. Thus, the assumption that $\text{IC}(\mathcal{A}; \mathcal{D}) < \epsilon n L$ cannot be true. \blacktriangleleft



■ **Figure 3** A simple example that illustrates the interleaving process for our random stream. We set the parameters $B = 2$ and $T = 8$ so that we sample from $L = 3$ scales to get a stream of total length $n = 24$. We do not illustrate the entire stream to save space. In this example we consider $\ell_0 = 2$ and demonstrate what a typical Z_j looks like. The sequence of random variables inside the red dotted box/window consists of the terms $(Z_j^<, Z_j, Z_j^>)$ and we highlight the memory states M_{j-1}, M_j .

2.2 Amortized space complexity via direct sum

We will now “lift” the information lower bound for any (ϵ, δ) -approximate counter to a lower bound for any (k, ϵ, δ) -approximate counter from which we can derive a lower bound on the memory size. The proof follows from a simple direct sum argument applied to the information cost.

► **Lemma 8.** *Any (k, ϵ, δ) -approximate counter with parameters $\delta \in (0, c_1)$ and $\epsilon \in (0, \frac{c_2}{\log \log(1/\delta)})$ that maintains k counters with total sum at most $n \geq k(c_3 \log^{c_4}(1/\delta))$ must use*

$$\Omega(k \log \log(1/\delta)) = \Omega(k \min\{\log(n/k), \log \log(1/\delta)\})$$

bits of memory in expectation.

Proof. We prove the lower bound via a reduction to the information lower bound for a single counter. That is, we will embed updates from the hard distribution \mathcal{D} used to prove the single counter lower bound in Lemma 6 into the k counter problem. Fix a (k, ϵ, δ) -approximate counter \mathcal{A}_k . Since the distribution \mathcal{D} is a product distribution, we can write $\mathcal{D} = \prod_{i=1}^{n'} \mathcal{D}_i$ for $n' := n/k$. We define $X_i^k := (X_{i,1}, \dots, X_{i,k}) \sim (\mathcal{D}_i)^k$ for each $i \in [n']$. We consider X_i^k to be the i th input to \mathcal{A}_k where $X_{i,j}$ is the update applied to j th counter. The overall input to \mathcal{A}_k is

$$(X_1^k, \dots, X_{n'}^k) \sim \prod_{i=1}^{n'} (\mathcal{D}_i)^k,$$

and the total sum of all counts is at most $n = kn' \geq k(c_3 \log^{c_4}(1/\delta))$ (see Lemma 6). Define $\mathcal{D}_{\text{int}} := \prod_{i=1}^{n'} (\mathcal{D}_i)^k$. Let M_i be the memory state of the algorithm after processing the i th input X_i^k . We can lower bound the information cost incurred by \mathcal{A}_k :

$$\text{IC}(\mathcal{A}_k, \mathcal{D}_{\text{int}}) = \sum_{i=1}^{n'} \sum_{j=1}^i I(M_i; X_j^k \mid M_{j-1}) \geq \sum_{u=1}^k \sum_{i=1}^{n'} \sum_{j=1}^i I(M_i; X_{j,u} \mid M_{j-1}),$$

where the inequality follows from the superadditivity of conditional mutual information (Proposition 4).

We now explain how we embed the single counter problem into the k counter problem using \mathcal{A}_k . Given the input $X^{n'} = (X_1, \dots, X_{n'}) \sim \mathcal{D}$ for a single counter, we pick an index $u \in [k]$ uniformly at random and proceed as if the updates are applied to u th counter of \mathcal{A}_k . Denote by U this uniformly random index. For the other counters, we simulate “fake” updates from the same distribution and apply them to \mathcal{A}_k as if they were received as inputs. Denote by \mathcal{A}' the resulting approximate counter that maintains $X^{n'}$. We can upper bound the information cost of \mathcal{A}' by

$$\begin{aligned} \text{IC}(\mathcal{A}', \mathcal{D}) &= \sum_{i=1}^{n'} \sum_{j=1}^i I(M_i; X_{j,U} \mid M_{j-1}, U) \\ &= \frac{1}{k} \sum_{u=1}^k \sum_{i=1}^{n'} \sum_{j=1}^i I(M_i; X_{j,u} \mid M_{j-1}) \\ &\leq \frac{1}{k} \sum_{i=1}^{n'} \sum_{j=1}^i I(M_i; X_j^k \mid M_{j-1}) \\ &= \frac{\text{IC}(\mathcal{A}_k, \mathcal{D}_{\text{int}})}{k}. \end{aligned}$$

Combining the above with Lemma 6 provides a lower bound on $\text{IC}(\mathcal{A}', \mathcal{D})$, which further implies the space lower bound

$$\frac{1}{n'} \sum_{i=1}^{n'} \mathbb{E}[|M_i|] \geq \Omega(k \min\{\log(n/k), \log \log(1/\delta)\})$$

for the claimed range of ϵ and δ . ◀

2.3 Offline lower bound

We now state and prove the remaining part of the lower bound.

33:14 On the Amortized Complexity of Approximate Counting

► **Lemma 9.** For any $\epsilon \in (0, c_1)$, $\delta \in (0, c_2)$ and $1 < k \leq n$, any (k, ϵ, δ) -approximate counter \mathcal{A} that maintains k counters with total sum at most n must use

$$|M| = \Omega(\min\{k \log(n/k), k \log(1/\epsilon) + k \log \log(n/k)\})$$

bits of memory.

Proof. Let $N(i) = \lceil (e^{4i\epsilon} - 1) \cdot \epsilon^{-1} \rceil$. Some simple calculations show

$$\begin{aligned} & (1 - \epsilon) \cdot N(i + 1) - (1 + \epsilon) \cdot N(i) \\ & \geq \epsilon^{-1}(1 - \epsilon)(e^{4(i+1)\epsilon} - 1) - \epsilon^{-1}(1 + \epsilon)(e^{4i\epsilon} - 1) - (1 + \epsilon) \\ & = \epsilon^{-1}((1 - \epsilon)e^\epsilon - 1 - \epsilon)e^{4i\epsilon} + 1 - \epsilon \\ & \geq \epsilon^{-1}((1 - \epsilon)e^{4\epsilon} - 1 - \epsilon)e^{4i\epsilon} \\ & \geq \epsilon^{-1}((1 - \epsilon)(1 + 4\epsilon) - 1 - \epsilon) \cdot 1 \\ & = 2 - 4\epsilon > 0. \end{aligned}$$

So for any $i \neq i'$, $N(i)$ and $N(i')$ must have a $(1 \pm \epsilon)$ multiplicative gap. We will consider k counters that take on such values, i.e. k counters that receive an (arbitrary) sequence of increments leading to counts $N(i_1), N(i_2), \dots, N(i_k)$ respectively. It is easy to see that if $i_r \leq (4\epsilon)^{-1} \cdot \ln(1 + n\epsilon/k)$ for all $r \in [k]$, the total sum of all the counters will be at most n . Let $q := \lfloor (4\epsilon)^{-1} \cdot \ln(1 + n\epsilon/k) \rfloor$ and define the set of possible counts $\mathbf{N} := \{N(1), N(2), \dots, N(q)\}$. We can represent the counts of all k counters as vectors in \mathbf{N}^k .

Consider a “large” collection of vectors $V \subseteq \mathbf{N}^k$ such that for every pair of counts $u, v \in V$, u and v differ in at least at least 90% of the coordinates. Notice that this implies that u and v differ multiplicatively in at least 90% of the coordinates by definition. Such a V is equivalent to an error correcting code.

► **Definition 10.** An error correcting code C of length k over a finite alphabet Σ is a subset of Σ^k . The elements of C are called code words. The distance of the code C , denoted $\Delta(C)$, is defined as the minimum hamming distance between any two code words $c_1, c_2 \in C$, i.e.

$$\Delta(C) := \min_{\substack{c, c' \in C \\ c \neq c'}} \Delta(c, c'),$$

where $\Delta(c, c') := |\{i : c_i \neq c'_i\}|$ is the hamming distance between two vectors.

In the language of error correcting codes, we want our collection of counts V to be a large error correcting code with a minimum distance of $0.9k$. Fortunately, the Gilbert-Varshamov bound immediately implies the existence of such a V that is large enough for our purposes.

► **Lemma 11 (Gilbert-Varshamov bound).** For any alphabet size $q > 1$, code length k and distance $d \leq k$, there exists an error correcting code C with size,

$$|C| \geq \frac{q^k}{\sum_{i=0}^{d-1} \binom{k}{i} (q-1)^i}.$$

Consequently, for $d = 0.9k$ and any q larger than a universal constant, there is a code C with size $|C| \geq q^{0.05k}$.

Thus, when q is a large enough constant (which can be achieved by making for ϵ smaller than some universal constant c), the Gilbert-Varshamov bound tells us there is a choice of V such that $|V| \geq q^{0.05k}$. Fix a (k, ϵ, δ) -approximate counter \mathcal{A} . For any $v \in V$ and $i \in [k]$, \mathcal{A} must

accurately approximate v_i with probability at least $1 - \delta$. Denote the event that \mathcal{A} correctly approximates v_i by $E_{v,i}$. Since $\mathbb{E}[\sum_{i=1}^k \mathbb{1}(E_{v,i}^c)] \leq k \cdot \delta < k/20$, by Markov's inequality we can conclude the existence of a subset $I_v \subseteq [k]$ such that:

1. $|I_v| \geq 0.9k$,
2. $\mathbb{P}[\cap_{i \in I_v} E_{v,i}] > 1/2$. Put in words, with probability at least $1/2$, for every $i \in I_v$, the algorithm \mathcal{A} outputs a $(1 + \epsilon)$ -approximation for v_i .

Define the event $E_v := \cap_{i \in I_v} E_{v,i}$. Since $\mathbb{E}[\sum_{v \in V} \mathbb{1}(E_v^c)] \leq |V|/2$, a standard averaging argument implies the existence of a fixed choice for the random seed of \mathcal{A} and a subset of counts $V' \subseteq V$ such that:

1. $|V'| > |V|/2$,
2. and the algorithm is correct on all $v \in V'$ in the sense of the event E_v on this random seed.

Fix such a random seed and denote the now deterministic algorithm \mathcal{A}' . For two counts $u, v \in V'$ define the set $D_{u,v} := \{i \in [k] : u_i \neq v_i\}$. By construction we have that $|D_{u,v}| \geq 0.9k$. We have

$$\begin{aligned} |D_{u,v} \cap I_u \cap I_v| &\geq k - |D_{u,v}^c| - |I_u^c| - |I_v^c| \\ &\geq k - 0.3k = 0.7k > 1. \end{aligned}$$

Thus, there is at least one index $i^* \in [k]$ such that $u_{i^*} \neq v_{i^*}$ and \mathcal{A}' $(1 + \epsilon)$ -approximates both u_{i^*} and v_{i^*} , so \mathcal{A}' arrives at a different state for u and v . Since this holds for every pair in V' , we can conclude that \mathcal{A}' must arrive at a different state for every count in V' . We can now conclude that

$$2^{|M|} \geq |V'| \geq 0.5 \cdot (\lfloor (4\epsilon)^{-1} \cdot \ln(1 + n\epsilon/k) \rfloor)^{0.05k} = \left(\Omega\left(\frac{\ln(1 + n\epsilon/k)}{\epsilon}\right) \right)^{0.05k},$$

or

$$|M| \geq 0.05k \log\left(\frac{\ln(1 + n\epsilon/k)}{\epsilon}\right) - O(k).$$

We distinguish into three cases:

1. If $\epsilon < k/n$, then $\frac{\ln(1+n\epsilon/k)}{\epsilon} = \Omega(n/k)$, which implies $|M| \geq 0.05k \log(n/k) - O(k)$.
2. If $k/n \leq \epsilon < \sqrt{k/n}$, we have

$$\begin{aligned} |M| &\geq 0.05k \log(1/\epsilon) - O(k) \\ &\geq 0.05k \log(1/\epsilon) + k \log \log(n/k) - O(k \log \log(1/\epsilon)). \end{aligned}$$

3. If $\epsilon > \sqrt{k/n}$, we have

$$\begin{aligned} |M| &\geq 0.05k \log(1/\epsilon) + 0.05k \log \log(en/k) - O(k) \\ &\geq 0.05k \log(1/\epsilon) + 0.05k \log \log(n/k) - O(k). \end{aligned}$$

Putting all the above bounds together gives us

$$\begin{aligned} |M| &\geq \min\{0.05k \log(n/k) - O(k), 0.05k \log(1/\epsilon) + 0.025k \log \log(n/k) \\ &\quad - O(k \log \log(1/\epsilon))\} \\ &= \Omega(\min\{k \log(n/k), k \log(1/\epsilon) + k \log \log(n/k)\}). \end{aligned}$$

◀

3 Upper bounds

We now state matching upper bounds (k, ϵ, δ) -approximate counting that follow immediately from the single counter case. We give upper bounds in two regimes: the case $k \leq N$, and the case $k > N$. We start by analyzing the first case.

Recall that the space usage of an approximate counter is typically a random variable, and the goal is to then prove an upper bound on the *expected* space (or, say, a high probability bound). The work of Nelson and Yu [9] provided the following bound on expected space usage for a single counter:

► **Theorem 12** ([9]). *For any $\epsilon, \delta \in (0, 1/2)$, there is an (ϵ, δ) -approximate counter with expected space usage $O(\log \log N + \log \log(1/\delta) + \log(1/\epsilon))$ bits.*

The following is then a very simple corollary of Theorem 12.

► **Corollary 13.** *For any $\epsilon, \delta \in (0, 1/2)$ and $1 \leq k \leq N$, there is a (k, ϵ, δ) -approximate counter that uses $O(k(\log \log(2N/k) + \log \log(1/\delta) + \log(1/\epsilon)))$ bits in expectation.*

Proof. We simply instantiate k independent copies of the data structure from Theorem 12 to provide k independent approximate counters, one per actual counter. For each $1 \leq i \leq k$, let S_i denote the (random) number of bits of memory used to store the approximate counter representing the i th counter N_i and recall $N := \sum_i N_i$. Then the total expected space is

$$\begin{aligned} \mathbb{E} \left[\sum_{i=1}^k S_i \right] &= \sum_{i=1}^k \mathbb{E}[S_i] \\ &\leq C \sum_{i=1}^k (\log \log(N_i) + \log \log(1/\delta) + \log(1/\epsilon)) \quad (\text{Theorem 12}) \\ &= Ck(\log \log(1/\delta) + \log(1/\epsilon)) + \sum_{i=1}^k \log \log(N_i) \\ &\leq Ck(\log \log(4N/k) + \log \log(1/\delta) + \log(1/\epsilon)) \quad (\text{Jensen}), \end{aligned}$$

where the last inequality uses concavity of the function $x \in (1, \infty) \mapsto \log \log(x)$. Note that we inject a constant 4 in the iterated logarithm so that the $\log \log$ term stays nonnegative. ◀

We now turn to the case $k > N$. In this case, there is a clear lower bound of $\Omega(N \log(Ck/N))$ bits, since $\log(\sum_{j=1}^N \binom{k}{j})$ bits of memory are needed to simply remember which counters are non-zero, and the logarithm of this sum is $\Theta(N \log(Ck/N))$ [5, Exercise 9.42]. We claim that this is also an upper bound. Specifically, we can use $O(N \log(Ck/N))$ bits to remember the locations of the $t \leq N$ non-zero counters. Now we have reduced to the case $k = t \leq N$ and can apply Corollary 13 to approximately remember their values.

References

- 1 Ziv Bar-Yossef, T. S. Jayram, Ravi Kumar, and D. Sivakumar. An information statistics approach to data stream and communication complexity. *J. Comput. Syst. Sci.*, 68(4):702–732, 2004.
- 2 Mark Braverman, Sumegha Garg, and David P. Woodruff. The coin problem with applications to data streams. In *Proceedings of the 61st IEEE Annual Symposium on Foundations of Computer Science (FOCS)*, pages 318–329, 2020.

- 3 Amit Chakrabarti, Yaoyun Shi, Anthony Wirth, and Andrew Chi-Chih Yao. Informational complexity and the direct sum problem for simultaneous message complexity. In *Proceedings of the 42nd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 270–278, 2001.
- 4 Thomas M. Cover and Joy A. Thomas. *Elements of information theory*. Wiley, 2 edition, 2006. URL: <http://www.elementsofinformationtheory.com/>.
- 5 Ronald L. Graham, Donald E. Knuth, and Oren Patashnik. *Concrete Mathematics: A Foundation for Computer Science, 2nd Ed.* Addison-Wesley, 1994.
- 6 Elena Kolevska. What happens with Redis runs out of memory, December 2018. URL: <https://www.youtube.com/watch?v=Xjq5XL2u3po>.
- 7 Jérémie Lumbroso. The story of HyperLogLog: How Flajolet processed streams with coin flips. *CoRR*, abs/1805.00612, 2018.
- 8 Robert H. Morris. Counting large numbers of events in small registers. *Commun. ACM*, 21(10):840–842, 1978.
- 9 Jelani Nelson and Huacheng Yu. Optimal bounds for approximate counting. In *Proceedings of the 41st ACM International Conference on Principles of Database Systems (PODS)*, pages 119–127, 2022.
- 10 Michael Sipser. *Introduction to the theory of computation*. PWS Publishing Company, 1997.