Analyzing Count Min Sketch with Conservative Updates[1]

Younes Ben Mazziane, Sara Alouf, Giovanni Neglia

[a] *Université Côte d'Azur, Inria, Sophia Antipolis, France*

Abstract

Count-Min Sketch with Conservative Updates (CMS-CU) is a popular algorithm to approximately count items' appearances in a data stream. Despite CMS-CU's widespread adoption, the theoretical analysis of its performance is still wanting because of its inherent difficulty. In this paper, we propose a novel approach to study CMS-CU and derive new upper bounds on both the expected value and the CCDF of the estimation error under an i.i.d. request process. Our formulas can be successfully employed to derive improved estimates for the precision of heavy-hitter detection methods and improved configuration rules for CMS-CU. The bounds are evaluated both on synthetic and real traces.

Keywords: Count-Min Sketch, Conservative Updates, Networking

1. Introduction

Counting how many times a given item appears in a data stream is a basic step common to a variety of applications spanning different domains including network management. For example, routers and servers often routinely count the number of packets in each flow for troubleshooting, traffic monitoring [1], detection of denial of service attacks [2], etc. Similarly, caching policies often rely on content popularity estimates [3]. Counting is a deceptively simple operation: in many applications the available memory does not permit to instantiate a counter for each possible item, because the number of items is huge (e.g., catalogs of cacheable objects in content delivery networks) or because counters are updated frequently and then require expensive fast memories (e.g., for high-rate inline packet flow processing). As a consequence, these applications rely on approximate counting techniques such as sketch-based algorithms [4], among which a popular one is the Count-Min Sketch (CMS) [5]. Many recent sketch algorithms are variations of CMS [6, 7, 8, 9].

CMS achieves significant memory reduction by mapping different items to the same counters through hash functions. As different items may increment the same counter, CMS suffers from overestimation errors. When counters are only incremented, a slight modification to CMS operation, referred to as Conservative Update [10] or Minimal Increment [11], can reduce the estimation error. The Count-Min sketch with Conservative Updates (CMS-CU) is successfully employed for caching [3], heavy flows detection [12], telemarketing call detection [13], and natural language processing [14].

Although conservative updates are a minor modification to CMS operation, they entangle the growth of the different counters, making CMS-CU much more difficult to study than CMS. As CMS-CU reduces CMS estimation errors, it is still possible to maintain upper bounds originally proposed for CMS [5, 15]. This approach has been adopted in some papers, for example to study CMS-CU's trade-off between memory and accuracy [12, 16], but it fails to capture the specific advantages offered by CMS-CU. To the best of our knowledge, only three papers

[1] Accepted Manuscript, Computer Networks, DOI: 10.1016/j.comnet.2022.109315.

ventured to study CMS-CU [17, 18, 19]. Bianchi et al. relied on a fluid approximation under the assumption that all counters are equally likely to be updated at each step [17]. This assumption may be satisfied only for a large number of counters and a large number of items with similar popularities. In [18], Einziger and Friedman modeled CMS-CU as a stack of Bloom filters [20] and derived bounds for the error's Complementary Cumulative Distribution Function (CCDF) when requests follow the Independent Reference Model (IRM) [21]. Unfortunately, the CCDF computation in [18] requires an iterative procedure, whose time complexity grows quadratically with the target error value (and then in general with the stream length). Moreover, the analysis in both papers holds for families of strongly k-universal hash functions [22], only for sufficiently large values of k. Such families are however incompatible with memory-constrained applications that need CMS-CU, since memory requirements and computation time grow with k [23]. More recently, the authors of [19] proposed a statistical estimator for CMS-CU's error, but it requires to have access to the counters associated to one hash function for a representative stream. Last, the bounds derived in [17, 18, 19] are the same for all items regardless of their popularity, even though it has been observed in [17] that the most popular items are better estimated than less popular ones.

In this paper, we extend and strengthen our earlier work [24] that studies CMS-CU. We propose a novel analysis of CMS-CU that leads to new upper bounds on both the CCDF and the expected value of the estimation error under an IRM request process. Our methodology diverges from related work as it quantifies the error on a per-item basis, which is particularly suitable for data streams with heterogeneous items' popularities. The analysis also overcomes the limitations of the previous studies as (i) it holds for pairwise independent hash functions, and (ii) it provides CCDF expressions with time complexity independent of the error's value. We show that our formulas can be successfully employed to derive both improved estimates for the precision of heavy-hitter detection methods and improved configuration rules for CMS-CU. We compare our new bounds both qualitatively and quantitatively to those in [5], [15], [17], [18], and our preliminary bounds in [24].

The rest of the paper is organized as follows. In Section 2, we provide the background, review the state-of-the-art studies and introduce the notation. The theoretical analysis is carried out in Section 3. Section 4 presents numerical experiments both on synthetic and real world traces. Section 5 concludes the paper.

2. Background, Notation and Assumptions

2.1. Data Stream Model

A data stream is a sequence $S_t = (Z(s))_{s=1,\dots,t}$, where $Z(s)$ is an item from a universe $I = \{1, \dots, N\}$ [25]. In general, we want to compute a function of the sequence, $\mathcal{F}(S_t)$, for example the number of occurrences of a given item, the set of heavy-hitters (items whose number of requests exceeds a given threshold), or the top-k most frequent items. Streaming algorithms aim to compute the function of interest using a few passes through the data stream (only one for the applications we consider) with an amount of memory which is sublinear in the universe's size N and the data stream size t. Even for the simple quantities mentioned above, an exact computation requires a linear amount of memory and then streaming algorithms need to settle for approximate results. In the next section, we present two popular streaming algorithms for approximate counting.

In what follows, we denote the set of integer numbers between 1 and $d \in \mathbb{N}$ by $[d]$. Moreover, to lighten the notation, we do not append the sketch name to the symbols. We believe there will be no ambiguity as each sketch is presented and analyzed in a separate section.

2.2. Count-Min Sketch (CMS)

A Count-Min sketch is a two dimensional array with d rows, each with w counters. An item i is mapped to d counters, one per row, via d hash functions $\{h_r\}_{r\in[d]}$ chosen uniformly at random from a family of strongly 2-universal hash functions. We note that, once selected, the hash functions do not change during the processing of the stream S_t. We model the association between items and counters as a bipartite undirected graph $G = (I, O, E)$, where O is the set of counters and $E \triangleq \{(i, h_r(i)) : i \in I, r \in [d]\}$ is the set of edges. We denote the open neighbourhood of item i in the graph as $N_G(i) \triangleq \{c : (i, c) \in E\}$ (we naturally have $|N_G(i)| = d$, for any item i). We denote the value at time t of the counter in row r corresponding to item i as $c_i^r(t)$, with $c_i^r(0) = 0$. When item i is requested at time t, the counters $\{h_r(i)\}_{r\in[d]}$ are incremented by 1. Namely,

$$c_i^r(t) = c_i^r(t-1) + 1, \ \forall r \in [d]. \tag{1}$$

Let $n_i(t)$ denote the number of occurrences of item i in the stream up to time t. Note that $c_i^r(t)$ is updated not only by new requests for item i, but also by requests for all items that are also mapped by h_r to the same counter $h_r(i)$, i.e., by all items in the set $\{j \in I : h_r(j) = h_r(i)\}$. These items are said to *collide* with i. It follows that $c_i^r(t) = \sum_{j:\, h_r(j)=h_r(i)} n_j(t)$. As such, $c_i^r(t)$ upper bounds $n_i(t)$. We denote the error resulting from using $c_i^r(t)$ for estimating $n_i(t)$ as $e_i^r(t)$, i.e., $e_i^r(t) \triangleq c_i^r(t) - n_i(t)$. Since all counters' values $\{c_i^r(t)\}_{r\in[d]}$ upper bound $n_i(t)$, their minimum also upper bounds $n_i(t)$. This minimum is the estimate of $n_i(t)$ provided by CMS and we denote it as $\hat{n}_i(t)$,

$$\hat{n}_i(t) \triangleq \min_{r\in[d]} c_i^r(t) \ . \tag{2}$$

The estimation error is then

$$e_i(t) \ \triangleq \hat{n}_i(t) - n_i(t) = \min_{r\in[d]} e_i^r(t) \ . \tag{3}$$

We also introduce $\delta_{i,j}^r(s)$ to represent the contribution of item $j \neq i$ to counter $h_r(i)$ at time s. We have:

$$\delta_{i,j}^r(s) \triangleq \mathbb{1}\big(Z(s) = j,\ h_r(i) = h_r(j)\big) \ , \tag{4}$$

$$e_i^r(t) = \sum_{s\in[t]} \sum_{j\in I\setminus\{i\}} \delta_{i,j}^r(s) \ . \tag{5}$$

All quantities we defined are random variables due to the initial random choice of the hash functions. From (5) and the definition of strongly 2-universal hash functions [22], one can immediately conclude that $\mathbb{E}\left[e_i^r(t)\right] = \frac{\sum_{j\neq i} n_j(t)}{w} \leq \frac{t}{w}$. Applying (3), we obtain the following upper bound on the expected estimation error:

$$\mathbb{E}\left[e_i(t)\right] \leq \frac{t}{w} \quad \Leftrightarrow \quad \mathbb{E}\left[\frac{e_i(t)}{t}\right] \leq \frac{1}{w} \ . \tag{6}$$

Moreover, the random variables $\{e_i^r(t)\}_{r\in[d]}$ being i.i.d., an application of the Markov inequality leads to the following upper bound on the CCDF of $e_i(t)$:

$$\Pr\left(\frac{e_i(t)}{t} \geq x\right) \leq \left(\frac{1}{wx}\right)^d \ . \tag{7}$$

Cormode and Muthukrishnan proved this result in [5, Theorem 1] for the specific value $x = \frac{\mathrm{e}}{w}$.

2.3. Count-Min Sketch with Conservative Updates (CMS-CU)

The conservative update [10] or minimal increment [11] is an optimization of CMS that consists in incrementing only the counters that attain the minimum value. The update procedure when item i is requested at time t becomes

$$c_i^r(t) = \max\left(c_i^r(t-1), \min_{f\in[d]} c_i^f(t-1) + 1\right), \ \forall r \in [d] \ . \tag{8}$$

The error $e_i^r(t)$ in each row r, the count estimate $\hat{n}_i(t)$, and the estimation error $e_i(t)$, all depend on $c_i^r(t)$ in the same way as in CMS. Equations (2) and (3) hold with CMS-CU. The quantities $\{\delta_{i,j}^r(s)\}_{s\in[t],j\neq i}$ are now defined as

$$\delta_{i,j}^r(s) \triangleq \mathbb{1}\big(Z(s) = j, h_r(i) = h_r(j), \hat{n}_j(s-1) = c_i^r(s-1)\big). \tag{9}$$

Equation (5) holds for CMS-CU. With respect to (4), (9) captures the additional condition that counter $h_r(i)$ is updated by a request for j at time s only if its current value $c_i^r(s-1)$ coincides with the current estimate $\hat{n}_j(s-1)$. Because of this additional condition, CMS-CU enjoys always a smaller error than CMS. Therefore, CMS upper bounds on the expectation (6) and on the CCDF (7) also hold for CMS-CU.

2.4. State of the art

When evaluating our analysis in Section 4 with synthetic and real traces, we will compare our results to the seminal paper [5] and its follow-up by the same authors [15] as well as to the more recent [17, 18, 19] which provide, to the best of our knowledge, state-of-the-art analyses for CMS-CU.

Bianchi et al. consider in [17] a particular case where, at each step, all counters are equally likely to be updated. The corresponding error is called the error floor and denoted here as $\epsilon_f(t)$. Experimental observations suggest that the error floor bounds the expected estimation error for any stream process, i.e.

$$\mathbb{E}\left[e_i(t)\right] \lesssim \epsilon_f(t) \quad \Leftrightarrow \quad \mathbb{E}\left[\frac{e_i(t)}{t}\right] \lesssim \frac{\epsilon_f(t)}{t} \ . \tag{10}$$

The error floor (denoted here as $\epsilon_f(t)$) is approximated using the following formula

$$\epsilon_f(t) \approx \frac{\overline{g}}{w \cdot d} \cdot t \ , \tag{11}$$

$$\text{with} \quad \overline{g} = \lim_{t \to +\infty} \frac{1}{t} \sum_{s=1}^{t} g(s) \ , \tag{12}$$

$$g(s) = \mathbb{E}\left[\left|\{r : \ c^r_{Z(s)}(s-1) = \hat{n}_{Z(s)}(s-1)\}\right|\right] . \tag{13}$$

Equation (13) says that $g(s)$ is the expected number of increments made at time s. $\overline{g}$ can be approximated by $\frac{1}{t_0}\sum_{s=1}^{t_0} g(s)$, for t_0 large enough. Each value $g(s)$ can be estimated through Montecarlo simulations or solving numerically an opportune differential equation. The authors show that $\overline{g}$ depends on d but not on w.

Building on (10), we derive an approximate probabilistic bound on the error by a direct use of the Markov inequality:

$$\Pr\left(\frac{e_i(t)}{t} \geq x\right) \lesssim \frac{\epsilon_f(t)}{xt} \ . \tag{14}$$

Einziger and Friedman modeled in [18] CMS-CU as a stack of Bloom filters. They proposed approximate bounds on the CCDF of the error $e_i(t)$ under the IRM model as follows:

$$\Pr\left(\frac{e_i(t)}{t} \geq x\right) \lesssim \begin{cases} \mathrm{PFP}(A_{\lceil xt \rceil}), & \text{if } \mathbb{E}\left[n_i(t)\right] \geq 1, \\ \mathrm{FP}(A_{\lceil xt \rceil}), & \text{otherwise,} \end{cases} \tag{15}$$

where $A_k \triangleq \mathbb{E}\left[|\{j \in I : \ \hat{n}_j(t) \geq k\}|\right]$, $\mathrm{FP}(n)$ is the false positive probability of a regular Bloom filter after n insertions [20], and $\mathrm{PFP}(n)$ is its average over all the past insertions, i.e., $\mathrm{PFP}(n) = \frac{1}{n}\sum_{k=1}^{n}\mathrm{FP}(k)$. The false positive probability can be approximated as $\mathrm{FP}(n) \approx \left(1 - \mathrm{e}^{-n/w}\right)^d$, and A_k for $k \in \mathbb{N}$ can be recursively computed using the following formula,

$$A_k \approx D_k + \sum_{j=1}^{k-1} (D_j - D_{j+1}) \cdot \mathrm{PFP}(A_{k-j}), \tag{16}$$

where $D_k \triangleq |\{j \in I : \mathbb{E}[n_j(t)] \geq k\}|$. Note that the expected number of arrivals $\mathbb{E}[n_j(t)]$ is known under the IRM assumption.

Peiqing et al. [19] proposed a method to estimate the error for CMS-CU by leveraging the knowledge of the sketch counters $(K_j^r(t))_{j\in[w]}$ associated to the hash function h_r. Under this assumption, they propose an estimator $(\hat{g}(\delta,t))$ for the $(1-\delta)$-quantile of the error $e_i(t)$, that is for $g_i(\delta,t) \triangleq \inf\{q \in [0,t] : \Pr(e_i(t) > q) = \delta\})$. Assuming the values $(K_j^r(t))_{j\in[w]}$ are sorted in decreasing order, the estimator $\hat{g}(\delta,t)$ has the following expression:

$$\hat{g}(\delta,t) = K^r_{\lceil \delta^{1/d} w \rceil}(t) \ . \tag{17}$$

We observe that in practice the knowledge of $(K_j^r(t))_{j\in[w]}$ requires to simulate the whole CMS-CU evolution with a computational cost proportional to the stream size t. On the contrary, the computational cost of our method does not depend on the stream size.

2.5. Our Assumptions

We will assume in our analysis that the request process follows the Independent Reference Model (IRM) [21]; in other words, $\{Z(s)\}_{s\in[t]}$ are i.i.d. categorical random variables with $\Pr(Z(s) = i) = p_i$ for $i \in I$, and $\sum_{i\in I} p_i = 1$. We refer to p_i as the *popularity* of item i. Without loss of generality, we number items in I according to their popularity rank, hence $p_i \geq p_{i+1}$, for $i \in [N-1]$. Note that there are two sources of randomness in our setting: the hash functions' selection and the request process S_t. From now on, the expectation $\mathbb{E}[.]$ and the probability $\Pr(.)$ take both kinds of randomness into account.

3. Theoretical Analysis of CMS-CU

Under the IRM model, we first prove a tighter upper bound on the CCDF of $e_i(t)$ for CMS, then we upper bound both the CCDF and the expectation of $e_i(t)$ for CMS-CU.

3.1. CMS: CCDF of the Estimation Error

In this section we derive a bound for CMS error under the IRM assumption that is tighter than (7). As discussed in Section 2.3, our new bound applies as well to the CMS-CU error.

Proposition 1 (Upper bound on the CCDF of $e_i(t)/t$). *The CCDF of the estimation error $e_i(t)$, when using CMS, verifies*

$$\Pr\left(\frac{e_i(t)}{t} \geq x\right) \leq \mathcal{A}(x)^d \ , \tag{18}$$

where

$$\mathcal{A}(x) \triangleq \min_{k=0,\dots,w-1} \mathcal{A}_k(x) \ , \quad \mathcal{A}(0) = 1 \ , \tag{19}$$

and

$$\mathcal{A}_k(x) \triangleq \frac{1}{x(w-k)} \sum_{\substack{j>k \\ j\neq i}} p_j + \frac{k}{w} \ . \tag{20}$$

Proof 1. *See Appendix A.*

Proposition 1 extends known results in the literature. In particular, upper bounding the right-hand side of (18) by $(\mathcal{A}_0(x))^d$ yields (7), and then replacing $x = \mathrm{e}/w$, we obtain [5, Theorem 1]. Proposition 1 also recovers [15, Theorem 5.1], by considering the particular case where items are requested according to a Zipf distribution with parameter $\alpha > 1$ and upper bounding the right-hand side of (18) by $\mathcal{A}_{w/3}(x)^d$ and the tail $\sum_{j>w/3} p_j$ by $(w/3)^{1-\alpha}$. In our experimental evaluation in Section 4, we will use for comparison purposes a combination of [5, Theorem 1] and [15, Theorem 5.1], namely,

$$\Pr\left(\frac{e_i(t)}{t} \geq x\right) \leq \min\left(\mathcal{A}_0(x)^d, \mathcal{A}_{w/3}(x)^d\right) . \tag{21}$$

In order to highlight the relevance of this proposition, we present an example where the improvement of (18) over (7) is evident.

Example 1. *We consider a setting with a small set of k popular items collecting a fraction α of the requests and $N-k$ unpopular ones. More specifically,*

$$p_i = \begin{cases} \frac{\alpha}{k}, & \text{if } i \leq k, \\ \frac{1-\alpha}{N-k}, & \text{otherwise.} \end{cases} \tag{22}$$

For $x = \frac{\beta}{w}$ for some $\beta \in (1-\alpha, 1]$, bounding the right-hand side of (18) by $(\mathcal{A}_k(x))^d$, yields a bound asymptotically equivalent[2] to $\left(\frac{1-\alpha}{\beta}\right)^d$ as $w \to +\infty$, in sharp contrast with (7) which provides the trivial bound 1.

3.2. CMS-CU: CCDF of the Estimation Error

We consider now CMS-CU and derive an upper bound on the CCDF of the estimation error $e_i(t)/t$. While the bound that we derived in Proposition 1 when CMS is used holds also with CMS-CU, our objective in this section is to derive a bound that makes use of the enhancement that CMS-CU brings over CMS. An important step to this end is to characterize the random variable $\delta^r_{i,j}(s)$ defined in (9). We establish first a preliminary result on the expectation of $\delta^r_{i,j}(s)$ that proved to be useful in the following.

Lemma 1 (Upper bound on $\mathbb{E}\left[\delta^r_{i,j}(s)\right]$). *The expected contribution of item j to item i's count at row r at time s satisfies (23).*

$$\exists \alpha_{i,j} > 0, \beta_{i,j} \geq 0 : \quad \mathbb{E}\left[\delta^r_{i,j}(s)\right] \leq \frac{p_j}{w}\left(\gamma_{i,j} + \beta_{i,j}\mathrm{e}^{-\alpha_{i,j}(s-1)}\right), \tag{23}$$

with

$$\gamma_{i,j} \triangleq \begin{cases} 1, & \forall j \leq i, \\ \min\left(\mathcal{A}(p_i - p_j)^{d-1}, 1\right), & \forall j > i, \end{cases} \tag{24}$$

and $\mathcal{A}(x)$ given in (19).

Proof 2. *See Appendix B.*

---

[2] f(w) is asymptotically equivalent to g(w) when $w \to +\infty$ if $\lim_{w\to+\infty} \frac{f(w)}{g(w)} = 1$

It is interesting to observe the structure of the bound in (23). The term p_j/w is simply $\mathbb{E}\left[\delta^r_{i,j}(s)\right]$ with CMS as can readily be seen from (4). Therefore, the coefficient of p_j/w in (23) is an attenuation term that captures the effect of using the conservative update procedure. As $s \to \infty$, this attenuation term converges to $\gamma_{i,j}$. The larger the difference between p_i and p_j, the smaller the term $\gamma_{i,j}$. This is expected as, the larger the difference in popularity between any two items i and j, the likelier that $c^r_i(t) > \hat{n}_j(t)$, and then the lesser item j is able to interfere with item i's estimation.

We now state the main result of the paper.

Proposition 2 (Upper bound on the CCDF of $e_i(t)/t$). *The CCDF of the estimation error $e_i(t)/t$, when using CMS-CU, is upper bounded as follows:*

$$\Pr\left(\frac{e_i(t)}{t} \geq x\right) \leq \min\left(\mathcal{A}(x)^d, \mathcal{C}_k(x,i,t), k = 0, \ldots, w-1\right) \tag{25}$$

where

$$\mathcal{C}_k(x,i,t) \triangleq \frac{1}{xw\left(1-\left(\frac{k}{w}\right)^d\right)}\left(\sum_{\substack{j>k\\ j\neq i}} p_j\gamma_{i,j} + \frac{B_i(k)}{t}\right) + \left(\frac{k}{w}\right)^d \tag{26}$$

and

$$B_i(k) \triangleq \sum_{\substack{j>k\\ j\neq i}} \frac{p_j\beta_{i,j}}{1-\mathrm{e}^{-\alpha_{i,j}}}\ . \tag{27}$$

$\alpha_{i,j}$ and $\beta_{i,j}$ are the constants in Lemma 1 and $\gamma_{i,j}$ is given in (24).

Proof 3. *See Appendix C.*

In practical situations, the data stream size t is large enough to allow us to ignore the constants $B_i(k)$ in (26). Consequently, the upper bound on the CCDF depends solely on x and i for t large enough.

Remark 1. *In our earlier work [24], the upper bound we derived on the CCDF of the error was different from* (25) *as it used another function than $\mathcal{C}_k(x,i,t)$. While deriving the bound in [24] we had considered the partitioning used for proving Proposition 1. As we will compare numerically in Section 4 our new bound* (25) *to that in [24], we include it here for completeness.*

$$\Pr\left(\frac{e_i(t)}{t} \geq x\right) \leq \min\left(\mathcal{A}(x)^d, \mathcal{B}(x,i,t)\right)\ , \tag{28}$$

where

$$\mathcal{B}(x,i,t) \triangleq \min_{k=0,\ldots,w-1} \frac{1}{x}\left(\frac{1}{w-k}\sum_{\substack{j>k\\ j\neq i}} p_j\gamma_{i,j} + \frac{B'_i(k)}{t}\right) + \frac{k}{w}, \tag{29}$$

and $B'_i(k)$ are constants that depend only on item i and on k. (We have $B'_i(k) = B_i(k)/(w-k)$.)

We use the popularity distribution (22) of Example 1 to illustrate the improvement brought by Proposition 2 with respect to the bounds (7), (14), and (15), that are respectively proposed/deduced from [5], [17], and [18].

Example 2. *Consider the popularity distribution in* (22)*,* $w, N \to +\infty$*, and* $w = o(N)$*.*[3] *In this setting, we show in Appendix D that* (15) *results in the trivial bound* 1 *(this particular result holds for any popularity distribution). From Proposition 2, if we neglect the constants* $B_i(k)$ *and upper bound the CCDF by* $\mathcal{C}_k(x,i,t)$ *and* $\gamma_{i,j}$ *by* $(\mathcal{A}_0(p_1 - p_{k+1}))^{d-1}$*, we obtain, for any* $i \leq k$*,*

$$\Pr\left(\frac{e_i(t)}{t} \geq x\right) \lesssim \frac{1-\alpha}{x\left(1-\left(\frac{k}{w}\right)^d\right) w^d (p_1 - p_{k+1})^{d-1}} + \left(\frac{k}{w}\right)^d. \tag{30}$$

Choosing x *such that* $\frac{1}{x} = o\left(w^d\right)$ *and* $x \leq \frac{\overline{g} \cdot d}{w}$ *(*$\overline{g}$ *is defined in* (12)*), the bound in* (30) *is* $o(1)$*, whereas* (7) *and* (14) *result in the trivial bound* 1*.*

In Section 4 we will provide evidence that the qualitative difference seen in Example 2 also exists for realistic values of the parameters and Zipf popularity distributions. For example, Figure 1 shows that our analysis correctly predicts that estimates for item with rank 100 are off by at most of few units, while state-of-the-art approaches estimate that the error should be at least 10 times larger.

Remark 2. *For the least popular items, the CCDF bound in Proposition 2 could have been derived for CMS. In other words, our analysis for those items fails to capture the benefit of the conservative update procedure. Experimentally, we observed that it is possible to improve these results using* (10) *by modifying the functions* $\mathcal{C}_k(x,i,t)$ *in* (26) *as follows:*

$$\mathcal{C}_k(x,i,t) = \frac{\epsilon_f(t)}{xt\left(1-\left(\frac{k}{w}\right)^d\right)} \sum_{\substack{j>k \\ j\neq i}} p_j + \left(\frac{k}{w}\right)^d, \tag{31}$$

removing the dependency on the rank of item i *(*$\epsilon_f(t)$ *is defined in* (11)*).*

We highlight the usefulness of Proposition 2 in Section 3.4 where we estimate the precision achievable by CMS and CMS-CU methods applied to the heavy-hitters detection problem.

We complete our formal analysis of CMS-CU with the derivation of a bound on the expectation of the estimation error, which is the object of the next section.

3.3. CMS-CU: Expected Estimation Error

The results obtained so far suggest two possible bounds on the expectation of $e_i(t)/t$. One bound can be derived using (3), (5) and Lemma 1. Another bound comes from writing the expectation as the sum of probability tails and using Proposition 2. We now detail the derivation of these bounds.

Because of (3) (that holds for CMS-CU), an upper bound on $\mathbb{E}\left[e_i^r(t)/t\right]$ suffices. Such a bound is readily found by linearity of the expectation, using (5) and Lemma 1. We find

$$\begin{aligned}\mathbb{E}\left[\frac{e_i^r(t)}{t}\right] &\leq \frac{1}{t} \sum_{s\in[t]} \sum_{j\in I\setminus\{i\}} \frac{p_j}{w}\left(\gamma_{i,j} + \beta_{i,j} \mathrm{e}^{-\alpha_{i,j}(s-1)}\right) \\ &\leq \frac{1}{w} \sum_{j\in I\setminus\{i\}} p_j \gamma_{i,j} + \frac{1}{wt} \sum_{j\in I\setminus\{i\}} \frac{p_j \beta_{i,j}}{1-\mathrm{e}^{-\alpha_{i,j}}}\ .\end{aligned}$$

[3] $\mathrm{f(w)=o(g(w))}$ if $\lim_{w\to+\infty} \frac{f(w)}{g(w)} = 0$.

The second summation is nothing but $B_i(0)$ (recall (27)). We obtain the first bound (which we derived in our earlier work [24])

$$\mathbb{E}\left[\frac{e_i^r(t)}{t}\right] \leq \frac{1}{w}\left(\sum_{j\in I\setminus\{i\}} p_j\gamma_{i,j} + \frac{B_i(0)}{t}\right). \tag{32}$$

Incidentally, the bound in (32) is simply $\mathcal{C}_0(1,i,t)$. It is interesting to notice that the term $B_i(0)$ can be ignored as $t\to\infty$, leading to a bound that is insensitive to t. As discussed earlier, $\gamma_{i,j}$ is a decreasing function of $p_i - p_j$, thus, of j. A necessary and sufficient condition for (32) to improve over (6) for a given item i is then to have $\gamma_{i,N} < 1$. At the same time, $\gamma_{i,j}$ is an increasing function of i. Therefore, the more popular the item, the smaller the bound (32), which is always smaller than or equal to the bound (6) when neglecting $B_i(0)$.

To derive the second bound we write

$$\begin{aligned}\mathbb{E}\left[\frac{e_i(t)}{t}\right] &= \frac{1}{t}\sum_{n\in[t]} \Pr\left(e_i(t)\geq n\right) = \frac{1}{t}\sum_{n\in[t]}\Pr\left(\frac{e_i(t)}{t}\geq\frac{n}{t}\right)\\ &\leq \frac{1}{t}\sum_{n\in[t]}\min\left(\mathcal{A}\left(\frac{n}{t}\right)^d, \mathcal{C}_k\left(\frac{n}{t},i,t\right), k=0,\ldots,w-1\right)\end{aligned} \tag{33}$$

where we have used the upper bound on the CCDF found in (25) (Proposition 2).

Combining (32) and (33) yields the following proposition.

Proposition 3 (Upper bound on $\mathbb{E}\left[e_i(t)/t\right]$). *The error experienced by item i is upper bounded as follows*

$$\mathbb{E}\left[\frac{e_i(t)}{t}\right] \leq \min\left(\mathcal{C}_0(1,i,t), \frac{1}{t}\sum_{n\in[t]}\min\left(\mathcal{A}\left(\frac{n}{t}\right)^d, \mathcal{C}_k\left(\frac{n}{t},i,t\right), k=0,\ldots,w-1\right)\right). \tag{34}$$

While previous studies [17] and [18] bounded the error uniformly across items, Proposition 3 provides error bounds that depend on an item's popularity. In particular, our work is the first to support analytically the experimental evidence that the most popular items barely experience any error, as observed in [17] and also shown in Figure 3.

In order to highlight the importance of Proposition 3, we compare (34), more specifically (32), to (6) and (10) using the same popularity distribution as in Examples 1 and 2.

Example 3. *Recall the popularity distribution* (22). *We can show that for the k most popular items, if we neglect in* (32) *the constant $B_i(0)$ and let w diverge, the expectation of the error verifies*

$$\mathbb{E}\left[\frac{e_i(t)}{t}\right] \lesssim \frac{\alpha}{w}\left(1-\frac{1}{k}\right) + \Theta\left(\frac{1}{w^d}\right),\ \forall i\leq k. \tag{35}$$

The proof requires to upper bound $\gamma_{i,j}$ by $(\mathcal{A}_0(p_1-p_{k+1}))^{d-1}$ for $j>k$ and observe that $\mathcal{A}_0(p_1 - p_{k+1}) = \Theta\left(\frac{1}{w}\right)$.

The ratio of (35) *and* (6) *is $\alpha(1-1/k)+o(1)$ which is smaller than 1 at least for large w. Comparing* (35) *and* (10)*, it is not possible in general to conclude which bound is tighter. In practice, we will consider the minimum of the two. We stress though that our analysis is formal whereas that in [17], leading to* (10)*, is based on experimental observations.*

Remark 3. *In practice, metrics like Average Absolute Error:* $AAE = \frac{\sum_{i\in I}|e_i(t)|}{|I|}$*, Average Relative Error:* $ARE = \frac{1}{|I|}\sum_{i\in I}\frac{|e_i(t)|}{n_i(t)}$ *and Weighted Average Absolute Error:* $WAAE = \sum_{i\in I}\frac{n_i(t)}{t}\cdot|e_i(t)|$ *are used to evaluate the performance of a sketch, e.g., [26, 27, 7]. Approximating* $n_i(t)$ *by* $p_i \cdot t$ *allows one to evaluate the expected value of these metrics as follows:*

$$\mathbb{E}\left[AAE\right] = \frac{1}{|I|}\sum_{i\in I}\mathbb{E}\left[|e_i(t)|\right] \ , \tag{36}$$

$$\mathbb{E}\left[ARE\right] \approx \frac{1}{|I|}\sum_{i\in I}\frac{\mathbb{E}\left[|e_i(t)|\right]}{p_i t} \ , \tag{37}$$

$$\mathbb{E}\left[WAAE\right] \approx \sum_{i\in I} p_i \cdot \mathbb{E}\left[|e_i(t)|\right] \ , \tag{38}$$

then an upper bound can be derived using our result in Proposition 3.

3.4. Heavy-Hitters Use Case: Lower Bound on the Precision

Heavy-hitters are items whose request rate exceeds a given threshold ϕ. Detecting heavy-hitters in a stream can be done using a sketch (for instance CMS or CMS-CU), but sketches overestimate the number of requests and can then lead to "false positives," i.e., items with a rate smaller than ϕ can erroneously be classified as heavy-hitters. Let H be the set of heavy-hitters, $H = \{i: \ n_i(t) \geq \phi \cdot t\}$, and $\hat{H}$ be the set of items classified as heavy-hitters by the sketch, $\hat{H} = \{i: \ \hat{n}_i(t) \geq \phi \cdot t\}$. The "precision" is one metric used for assessing the performance of the sketch [4], and is defined as follows: $P = |H|/|\hat{H}|$. For the sake of simplicity, we assume that $n_i(t) \approx p_i t$, for large enough t; this is reasonable because of the law of large numbers. Under this approximation, $|H|$ is constant and we can write the expected value of the precision as:

$$\mathbb{E}\left[P\right] \approx \frac{|H|}{|H| + \sum_{i>|H|}\Pr\left(\frac{e_i(t)}{t} \geq \phi - p_i\right)} \ . \tag{39}$$

Combining (39) with Proposition 2 we obtain a lower bound on the expected precision when CMS-CU is used. This lower bound will be illustrated in Section 4.5 and compared to experimental values.

Again we highlight qualitatively the advantage of our bound (25) (Proposition 2) with respect to (7), (14), and (15) through an example.

Example 4. *For the heavy-hitter problem we consider a popularity distribution, slightly more complex than* (22)*, with an additional group of medium popular items. More specifically,*

$$p_i = \begin{cases} \alpha_1/k_1, & \text{if } i \leq k_1, \\ \alpha_2/k_2, & \text{if } k_1 < i \leq k, \\ (1-\alpha)/(N-k), & \text{otherwise}, \end{cases} \tag{40}$$

where $\frac{\alpha_1}{k_1} > \frac{\alpha_2}{k_2}$*,* $\alpha = \alpha_1 + \alpha_2$*, and* $k = k_1 + k_2$*.*

We consider $\frac{\alpha_2}{k_2} < \phi < \frac{\alpha_1}{k_1}$*, that is, the heavy-hitters coincide with the* k_1 *most popular items. The precision is then* $\mathbb{E}\left[P\right] = k_1/(k_1 + S_1 + S_2)$*, where* S_1 *and* S_2 *are defined as follows:*

$$\begin{aligned} S_1 &\triangleq (N-k)\cdot\Pr\left(\frac{e_N(t)}{t} \geq \phi - p_N\right) \\ S_2 &\triangleq k_2 \cdot \Pr\left(\frac{e_k(t)}{t} \geq \phi - p_k\right). \end{aligned} \tag{41}$$

The different approaches estimate in different way the probabilities appearing in (41)*. We study the regime where both N and w diverge with $w = o(N)$ and $N = o(w^d)$.*

If we use the bound (15) *in* (41)*, the arguments in Appendix D lead to conclude that S_1 is potentially unbounded (N diverges and $\Pr(e_N(t)/t \geq \phi - p_N)$ is upper-bounded by a positive constant). The conclusion is then trivial: the precision is lower bounded by* 0*. The bound* (14) *does not provide meaningful bounds for S_1 and the precision either.*

On the contrary, both (25) *and* (7) *guarantee that S_1 is arbitrarily small asymptotically. But the two bounds may draw different conclusions for S_2. In fact,* (7) *concludes that S_2 can be made smaller than ϵ for $\phi - p_k > \frac{1}{w}\sqrt[d]{\frac{k_2}{\epsilon}}$. The bound* (25) *can be relaxed to the simpler form*

$$\Pr\left(\frac{e_i(t)}{t} \geq x\right) \lesssim \frac{k_2^{d-1}(1-\alpha)}{w^d x \left(1 - \left(\frac{k}{w}\right)^d\right)\alpha_2^{d-1}} + \left(\frac{k}{w}\right)^d \, , \tag{42}$$

which is obtained similarly to (30)*, by considering only $\mathcal{C}_k(x,i,t)$ in the right hand side of* (25) *and upper bounding $\gamma_{i,j}$ by $(\mathcal{A}_0(p_i - p_j))^{d-1}$. From* (42)*, we conclude that there exists a constant a' such that S_2 is guaranteed to be smaller than ϵ when $\phi - p_k > \frac{a' \cdot k_2}{w^d \epsilon}$. In conclusion, if $\phi - p_k$ belongs to the interval $\left(\frac{a' \cdot k_2}{w^d \epsilon}, \frac{1}{w}\sqrt[d]{\frac{k_2}{\epsilon}}\right]$, our bound predicts that the precision is at least $1 - \epsilon$, while the bound* (7) *simply guarantees a precision inferior to $1 - \epsilon$.*

4. Experimental Evaluation and Numerical Analysis

4.1. Experimental Setting

To support our analysis, we have undertaken three series of experiments in which we simulated requests for items over time and used CMS-CU to count the requests for each item. In the first series of experiments, we generated 100 synthetic streams from a Zipf distribution with shape parameter $\alpha = 0.8$. Each stream contains 5 million requests for items in the set $I = [N]$ with $N =$ 1,000,000. For each stream, we employed a sketch with a default configuration of width $w =$ 10,000 and depth $d = 8$ and we selected different 'MurmurHash' hash functions [28] for each stream by choosing uniformly at random d different seeds in $[N]$. The experimental values reported for this setting are averaged over the 100 streams (or 100 independent trials). We also computed the 95% confidence intervals but do not report them as they are very narrow and would not be visible in the figures. The second series of experiments is identical to the first except for the shape parameter of the Zipf distribution which is $\alpha = 1.1$. In the third series of experiments, we used a trace log of accesses to Wikipedia pages in all languages during September 2007 [29]. The trace contains 10,628,125 requests. The number of distinct Wikipedia pages requested in this trace is 1,712,459. We extracted 10 non-overlapping chunks from this trace, each containing $N' =$1,000,000 requests, and discarded the rest. All theoretical values are computed assuming the IRM and that item popularities coincide with empirical frequencies over the first chunk, i.e., $p_i \approx n_i(N')/N'$. The reported experimental values are instead computed on the remaining 9 chunks of the trace. The CMS-CU in this setting is configured by default with a width $w = 5000$ and a depth $d = 5$.

4.2. Numerical Evaluation

For each series of experiments, we report the results obtained for the following metrics: (*i*) the CCDF of the sketch estimation error for both a popular item and a non-popular one, (*ii*) the expected sketch estimation error for each item along with the Average Absolute Error (AAE), the Average Relative Error (ARE), the Weighted Average Absolute Error (WAAE) , and (*iii*) the precision in the heavy-hitters detection problem. We compare our results to those of [5, 15], [17], [18], and our earlier work [24] which we refer to as 'CM05', 'BDLS12', 'EF15', and 'BAN22' respectively. We also carry out a limited comparison with the method in [19]—referred to as 'CWYJL21'—, which differs from our approach and all other methods listed above in that, in order to estimate the error at time

Table 1: Labels and equations used in the comparison.

Metric	[5, 15] CM05	[17] BDLS12	[18] EF15	[19] CWYJL21	Earlier work [24] BAN22	Current paper Ours
CCDF popular i	(21)	(14)	(15)	(17)	(28)	(25), (26)
non-popular i	(21)	(14)	(15)	(17)	(28)	(25), (31)
Sketch error	$\frac{1-p_i}{w}$	(10)	–	–	(32)	(34)
Precision	(39), (21)	(39), (14)	(39), (15)	–	(39), (28)	(39), (25), (26)

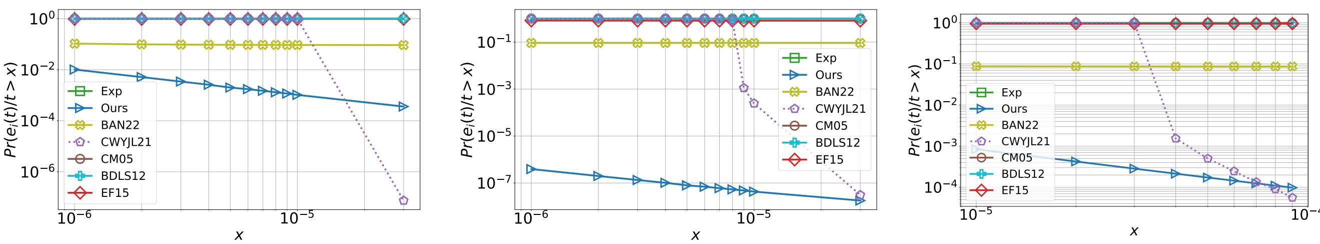

(a) synthetic, $\alpha = 0.8$ (b) synthetic, $\alpha = 1.1$ (c) Wikipedia

Figure 1: CCDF of error for a popular item. (a) and (b) Synthetic traces: N = 1,000,000, t = 5,000,000, w = 10,000, d = 8, trials = 100, item rank i = 100. (c) Wikipedia trace: N = 292,332, t = 1,000,000, w = 5000, d = 5, item rank i = 50.

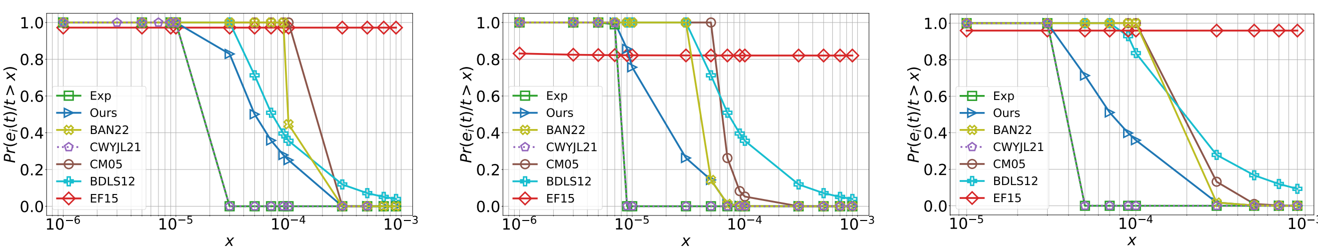

(a) synthetic, $\alpha = 0.8$ (b) synthetic, $\alpha = 1.1$ (c) Wikipedia

Figure 2: CCDF of error for a non-popular item. (a) and (b) Synthetic traces: N = 1,000,000, t = 5,000,000, w = 10,000, d = 8, trials = 100, item rank i = 100,000. (c) Wikipedia trace: N = 292,332, t = 1,000,000, w = 5000, d = 5, item rank i = 10,000.

t, it requires access to a sample of the counters at time t for a representative data stream (see Section 2). For completeness, we list the relevant formulas used in Table 1. In the following sections, the estimation of the ground truth using simulations is referred to as 'Exp.'

4.3. The CCDF of the Sketch Estimation Error

The results are shown in Figures 1 and 2 for popular and non-popular items, respectively. The ground truth CCDF (Exp) is not visible in Figure 1 as it is 0. Note that all existing approaches, apart from our earlier work BAN22, provide a single bound for the CCDF that is valid for all items and is essentially tailored on the unpopular items for which errors are larger. In all three series of experiments, our bounds are better at capturing the CCDF of the error for popular items (ranks 100 and 50), for which the other state-of-the-art approaches provide only rough estimations; see Figure 1. The interval identified analytically in Example 2 (over which our bounds decrease whereas most of the state of the art bounds are trivially 1) are clearly visible in Figure 1. In particular, in Figure 1b our upper bound in the interval $[10^{-6}, 10^{-5}]$ decreases and is below 10^{-6} while CM05, BDLS12 and EF15 are equal

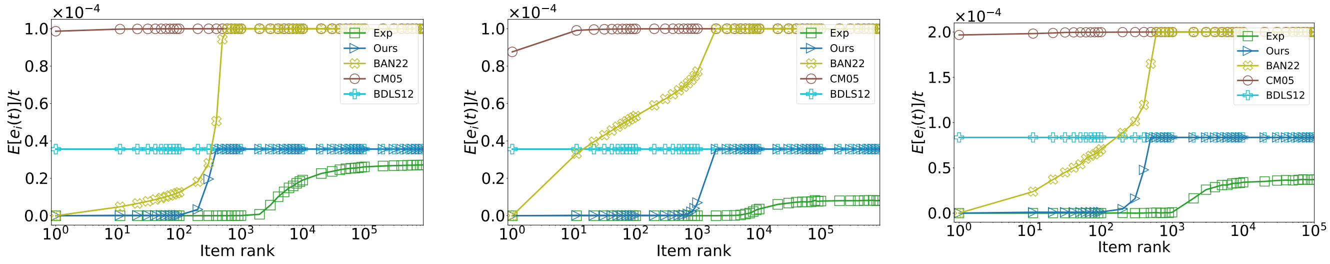

(a) synthetic, $\alpha = 0.8$ (b) synthetic, $\alpha = 1.1$ (c) Wikipedia

Figure 3: Estimation error for each item. (a) and (b) Synthetic traces: $N = 1{,}000{,}000$, $t = 5{,}000{,}000$, $w = 10{,}000$, $d = 8$, trials $= 100$. (c) Wikipedia trace: $N = 292{,}332$, $t = 1{,}000{,}000$, $w = 5000$, $d = 5$.

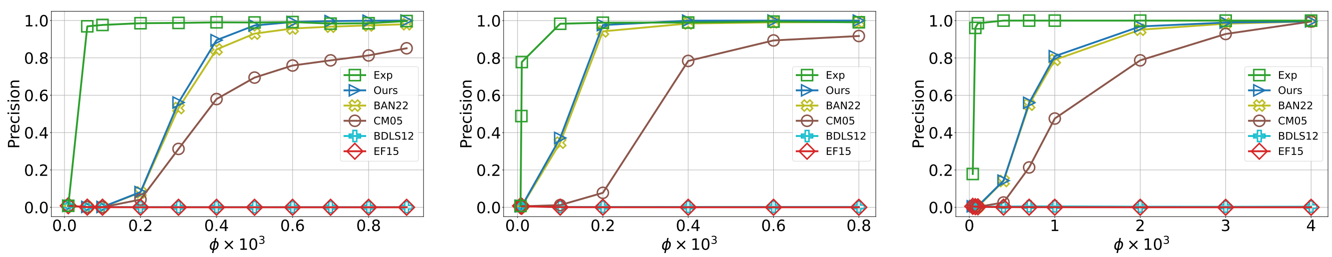

(a) synthetic, $\alpha = 0.8$ (b) synthetic, $\alpha = 1.1$ (c) Wikipedia

Figure 4: Precision as a function of the threshold ϕ. (a) and (b) Synthetic traces: $N = 1{,}000{,}000$, $t = 5{,}000{,}000$, $w = 10{,}000$, $d = 8$, trials $= 100$. (c) Wikipedia trace: $N = 292{,}332$, $t = 1{,}000{,}000$, $w = 5000$, $d = 5$.

to 1. The results for EF15 confirm our discussion in Appendix D. The only exception is CWYJL21, which improves on our CCDF bound for large values of the error. CWYJL21 is better at capturing the CCDF of the estimation error for non-popular items (ranks 100,000 and 10,000), with our method being the second best one; see Figure 2. Note that CWYJL21 requires to simulate CMS-CU over a trace drawn from the same distribution, and then its computational cost grows at least linearly with the length of the stream t.

4.4. The Expected Sketch Estimation Error

The results for the synthetic and Wikipedia traces are depicted in Figure 3. We observe that our analysis correctly predicts that different items experience a different error. It is evident that the bounds BAN22 and Ours improve over CM05. Our bound improves over BAN22 and BDLS12 in all three series of experiments, namely for the 300 most popular items when $\alpha = 0.8$, for the 1000 most popular items when $\alpha = 1.1$ and for the 400 most popular items for the Wikipedia trace. We notice that our bounds on the expectations are not always tight leaving room for improvement.

We also show in Table 2 other metrics—AAE, ARE and WAAE—commonly used to evaluate the performance of a sketch (see Remark 3). Our formulas are better than CM05 at predicting all three metrics but the improvement with respect to BDLS12 is only evident for WAAE. The reason is that our bounds improve BDLS12's ones only for the most popular items which constitute a small fraction of the whole catalogue. The difference is then marginal in terms of AAE and ARE which quantify average errors over the whole catalogue. On the contrary, WAAE is an average error per stream element and then errors on the most popular items are given a larger weight (i.e., proportional to their popularity). This metric shows then a clear difference between our approach and BDLS12.

Table 2: Average Absolute Error, Average Relative Error, Weighted Average Absolute Error. Synthetic traces ($\alpha = 0.8$, $\alpha = 1.1$): $N =$ 1,000,000, $t =$ 5,000,000, $w =$ 10,000, $d = 8$, trials = 100. Wikipedia trace: $N =$ 292,332, $t =$ 1,000,000, $w = 5000$, $d = 5$.

	Synthetic: $\alpha = 0.8$			**Synthetic: $\alpha = 1.1$**			**Wikipedia**		
	AAE	**ARE**	**WAAE**	**AAE**	**ARE**	**WAAE**	**AAE**	**ARE**	**WAAE**
CM05	500.0	260.0	500.0	500.0	1531.0	500.0	200.0	176.0	200.0
BDLS12	178.1	93.41696	178.1	178.1	545.1894	178.1	83.6	73.5783	83.6
Ours	178.0	93.41692	152.2	177.9	545.1890	54.3	83.5	73.5781	44.9
Exp	132.7	70.7	26.5	39.7	124.3	6.8	36.2	32.3	14.8

4.5. Precision in Detecting $\phi-$Heavy-Hitters

The results are presented in Figure 4. Our formulas outperform state of the art methods in bounding the precision for both synthetic and real world traces. The improvement is mainly due to the tightness of our CCDF bound for the most popular items. CM05 achieves high precision values only for large values of ϕ, as already qualitatively highlighted in Example 4. The poor lower bound on the precision achieved by EF15 was expected seeing its poor CCDF bound in Figures 1 and 2. It is interesting to observe that while the BDLS12 CCDF bound may be tighter than the CM05 CCDF bound for medium error values (like in Figures 1a and 2a for $x \in [3 \cdot 10^{-5}, 10^{-4}]$), its precision bound is much poorer. The reason is due to the sum of probability tails in the denominator of (39) being highly affected by the CCDF tail of non-popular items, and while the CM05 CCDF bound decreases exponentially fast, the BDLS12 CCDF bound decreases only inversely linearly (compare (14) with (7)).

4.6. Configuring CMS-CU with QoS Guarantees

The bounds we derived can also be used to configure the width w and the depth d of CMS-CU in order to achieve the desired precision with the minimum amount of memory. If each counter uses 4 bytes, the memory cost of a CMS-CU is $M = 4wd$ bytes. We compared numerically the memory requirements determined by our approach and by CM05. We do not consider BDLS12 and EF15 in this section as their lower bound on the precision is poor, a configuration method relying on such bounds will thereby return prohibitively large memory requirements.

For target precision values in the range 0.8–0.975, we performed a search for the total number of counters (equal to $w \times d$) in the range $\left[\lceil \frac{2}{\phi} \rceil, N\right]$ (with a step of $\lceil 2/\phi \rceil$) and depth values between 2 and 15 to find the smallest memory which guarantees the target precision. Figure 5 shows the corresponding curves obtained using our approach and CM05. Our approach leads to configuring CMS-CU using a reduced amount of memory. For instance, we observe in Figure 5a a reduction factor of 4.82 for 95% precision target (1.312 MB with CM05 vs. 0.272 MB with Ours) and up to 8.72 for a precision target of 97.5% (2.512 MB with CM05 vs. 0.288 MB with Ours). The gain is of the same order for the other traces. In particular, when the shape parameter of the Zipf distribution is 1.1, the reduction factor is 3.89 and 5.2 for precision targets of 95% and 97.5%, respectively (see Figure 5b). In the case of the Wikipedia trace, the required memory for a precision target of 97.5% is 0.104 MB with our approach as seen in Figure 5c, a reduction factor of 9.61 with respect to the 1 MB advocated by CM05 (reduction factor of 5.58 for a precision target of 95%).

Figure 5 suggests that large values of d are required to minimize the memory while achieving a desired precision target. At the same time, it may be undesirable to select a large number of hash functions due to computational constraints. The heatmap in Figure 6 shows the additional memory needed to achieve a given precision target with respect to the optimal configuration in Figure 5. The lighter the color, the less the additional required memory. Results show that, in practice, a reduction in the value of d is possible with very little memory cost and yet the same precision performance. Finally, we observe that hard computational constraints should lead to explore a smaller set of values for the depth d, i.e., to select a smaller maximal possible value of d. .

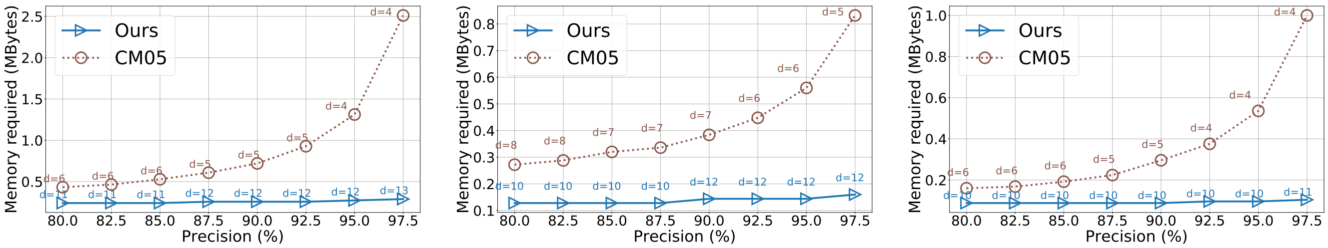

(a) synthetic, $\alpha = 0.8$ (b) synthetic, $\alpha = 1.1$ (c) Wikipedia

Figure 5: Estimated memory requirement for a given precision. (a) and (b) Synthetic traces: $N = 1{,}000{,}000$, $\phi = 5 \cdot 10^{-4}$. (c) Wikipedia trace: 1 million requests, $\phi = 10^{-3}$.

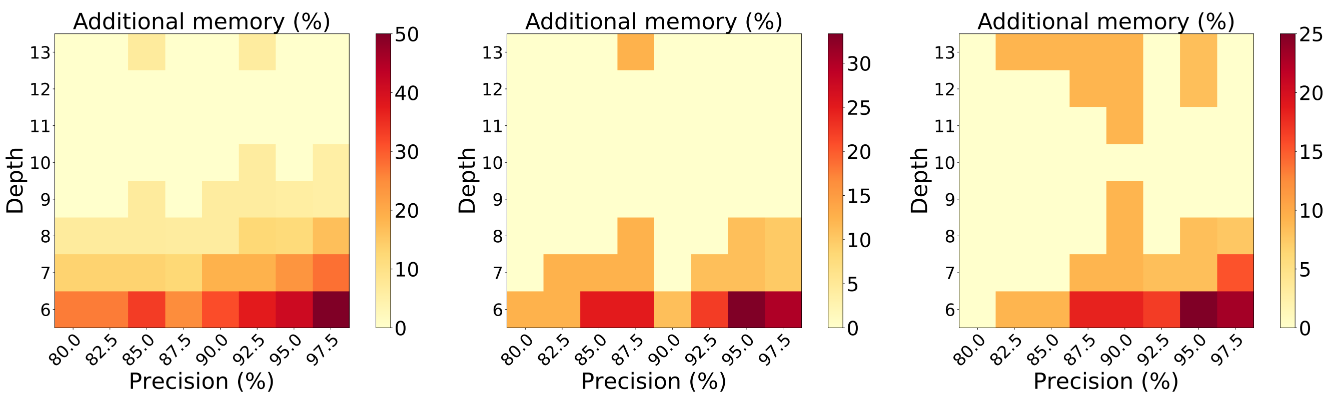

(a) synthetic, $\alpha = 0.8$ (b) synthetic, $\alpha = 1.1$ (c) Wikipedia

Figure 6: Additional memory required in comparison to the optimal setting for different depth values and precision targets (our model's estimations). (a) and (b) Synthetic traces: $N = 1{,}000{,}000$, $\phi = 5 \cdot 10^{-4}$. (c) Wikipedia trace: 1 million requests, $\phi = 10^{-3}$.

5. Conclusion and Perspectives

While it is a common belief that CMS-CU leads to smaller estimation errors for the most popular items [17], our paper is the first to provide quantitative support for such property, thanks to a per-item study of the estimation error. We showed that our analysis significantly improves existing bounds for the most popular items and leads, in comparison to the state of the art, to more accurate estimations for the precision in heavy-hitter detection problems as well as to improved configuration rules, which avoid to oversize the counting data structure.

For less popular items, our bounds are not tighter than existing ones. In the future, we want then to focus on improving the bounds for the tail of the popularity distribution. Moreover, we plan on extending our analysis to other popular sketches based on CMS-CU like the one proposed in [6].

References

[1] R. B. Basat, G. Einziger, R. Friedman, Y. Kassner, Optimal elephant flow detection, in: INFOCOM, IEEE, 2017, pp. 1–9.

[2] H. Liu, Y. Sun, M. S. Kim, Fine-grained ddos detection scheme based on bidirectional count sketch, in: 2011

Proc of 20th International Conference on Computer Communications and Networks (ICCCN), IEEE, 2011, pp. 1–6.

[3] G. Einziger, R. Friedman, B. Manes, TinyLFU: A highly efficient cache admission policy, ACM Trans. Storage 13 (2017) 1–31.

[4] G. Cormode, M. Hadjieleftheriou, Methods for finding frequent items in data streams, The VLDB Journal 19 (2010) 3–20.

[5] G. Cormode, S. Muthukrishnan, An improved data stream summary: the count-min sketch and its applications, Journal of Algorithms 55 (2005) 58–75.

[6] T. Yang, et al., Elastic sketch: Adaptive and fast network-wide measurements, in: Proc. 2018 Conf. ACM Special Interest Group Data Communication, 2018, pp. 561–575.

[7] C.-Y. Hsu, P. Indyk, D. Katabi, A. Vakilian, Learning-based frequency estimation algorithms, in: ICLR, 2019, pp. 1–20.

[8] Y. Zhao, K. Yang, Z. Liu, T. Yang, L. Chen, S. Liu, N. Zheng, R. Wang, H. Wu, Y. Wang, et al., {LightGuardian}: A {Full-Visibility}, lightweight, in-band telemetry system using sketchlets, in: 18th USENIX Symposium on Networked Systems Design and Implementation (NSDI 21), 2021, pp. 991–1010.

[9] K. Yang, Y. Li, Z. Liu, T. Yang, Y. Zhou, J. He, T. Zhao, Z. Jia, Y. Yang, et al., Sketchint: Empowering int with towersketch for per-flow per-switch measurement, in: 2021 IEEE 29th International Conference on Network Protocols (ICNP), IEEE, 2021, pp. 1–12.

[10] C. Estan, G. Varghese, New directions in traffic measurement and accounting: Focusing on the elephants, ignoring the mice, ACM Trans. Comput. Syst. 21 (2003) 270–313.

[11] S. Cohen, Y. Matias, Spectral Bloom filters, in: Proc. ACM SIGMOD Int. Conf. Management Data, 2003, pp. 241–252.

[12] R. Wang, H. Du, Z. Shen, Z. Jia, DAP-Sketch: An accurate and effective network measurement sketch with Deterministic Admission Policy, Computer Networks 194 (2021) 108155.

[13] G. Bianchi, N. d'Heureuse, S. Niccolini, On-demand time-decaying Bloom filters for telemarketer detection, ACM SIGCOMM Comput. Communi. Rev. 41 (2011) 5–12.

[14] A. Goyal, H. Daumé III, G. Cormode, Sketch algorithms for estimating point queries in NLP, in: Proc. 2012 joint Conf. on empirical methods in natural language processing and computational natural language learning, 2012, pp. 1093–1103.

[15] G. Cormode, S. Muthukrishnan, Summarizing and mining skewed data streams, in: Proc. SIAM Int. Conf. Data Mining, SIAM, 2005, pp. 44–55.

[16] F. Ventruto, M. Pulimeno, M. Cafaro, I. Epicoco, On frequency estimation and detection of heavy hitters in data streams, Future Internet 12 (2020) 158.

[17] G. Bianchi, K. Duffy, D. Leith, V. Shneer, Modeling conservative updates in multi-hash approximate count sketches, in: 24th ITC, IEEE, 2012, pp. 1–8.

[18] G. Einziger, R. Friedman, A formal analysis of conservative update based approximate counting, in: ICNC, IEEE, 2015, pp. 255–259.

[19] P. Chen, Y. Wu, T. Yang, J. Jiang, Z. Liu, Precise error estimation for sketch-based flow measurement, in: Proc. 21st ACM Internet Measurement Conf., 2021, pp. 113–121.

[20] A. Broder, M. Mitzenmacher, Network applications of Bloom filters: A survey, Internet mathematics 1 (2004) 485–509.

[21] R. Fagin, Asymptotic miss ratios over independent references, Journal of Computer and System Sciences 14 (1977) 222–250.

[22] R. Motwani, P. Raghavan, Randomized algorithms, Cambridge university press, 1995.

[23] A. Siegel, On universal classes of extremely random constant-time hash functions, SIAM Journal on Computing 33 (2004) 505–543.

[24] Y. Ben Mazziane, S. Alouf, G. Neglia, A formal analysis of the count-min sketch with conservative updates, in: to appear in INFOCOM WORKSHOPS proc., IEEE, 2022, pp. 1–6.

[25] S. Muthukrishnan, Data streams: Algorithms and applications, Now Publishers Inc, 2005.

[26] Y. Zhou, T. Yang, J. Jiang, B. Cui, M. Yu, X. Li, S. Uhlig, Cold filter: A meta-framework for faster and more accurate stream processing, in: Proceedings of the 2018 International Conference on Management of Data, 2018, pp. 741–756.

[27] Y. Zhang, Z. Liu, R. Wang, T. Yang, J. Li, R. Miao, P. Liu, R. Zhang, J. Jiang, Cocosketch: high-performance sketch-based measurement over arbitrary partial key query, in: Proceedings of the 2021 ACM SIGCOMM 2021 Conference, 2021, pp. 207–222.

[28] A. Appleby, `https://sites.google.com/site/murmurhash/`, last accessed on February 2022.

[29] G. Urdaneta, G. Pierre, M. van Steen, Wikipedia workload analysis for decentralized hosting, Elsevier Computer Networks 53 (2009) 1830–1845.

[30] K. Joag-Dev, F. Proschan, Negative association of random variables with applications, The Annals of Statistics (1983) 286–295.

Appendix A. Proof of Proposition 1

From (3) and the fact that the random variables $\{e_i^r(t)\}_{r\in[d]}$ are i.i.d. when using CMS, we have $\Pr\left(e_i(t)/t > x\right) = \left(\Pr\left(e_i^1(t)/t > x\right)\right)^d$. To prove (18) it is then sufficient to show that $\Pr\left(e_i^1(t)/t > x\right) \leq \mathcal{A}_k(x)$ for $k = 0, \dots, w-1$. For a given $k \neq 0$ we consider the event, called $E_{i,k}^r$, of no hash collision in row r between item i and any of the k most popular items (other than i, if $i \leq k$). By first writing the law of total probabilities with respect to the partition $\{E_{i,k}^1, \overline{E_{i,k}^1}\}$, and then using the union bound to write $\Pr\left(\overline{E_{i,k}^1}\right) \leq k/w$ and the Markov inequality to upper bound $\Pr\left(e_i^1(t)/t \geq x \mid E_{i,k}^1\right)$, we obtain

$$
\begin{aligned}
\Pr\left(\frac{e_i^1(t)}{t} \geq x\right) &\leq \Pr\left(\frac{e_i^1(t)}{t} \geq x \mid E_{i,k}^1\right) \cdot 1 + 1 \cdot \Pr\left(\overline{E_{i,k}^1}\right) && \text{(A.1)} \\
&\leq \frac{\mathbb{E}\left[e_i^1(t) \mid E_{i,k}^1\right]}{xt} + \frac{k}{w} && \text{(A.2)} \\
&\leq \frac{1}{x} \sum_{\substack{j>k \\ j\neq i}} p_j \Pr\left(h_1(i) = h_1(j) \mid E_{i,k}^1\right) + \frac{k}{w} && \text{(A.3)} \\
&\leq \frac{1}{x(w-k)} \sum_{\substack{j>k \\ j\neq i}} p_j + \frac{k}{w} = \mathcal{A}_k(x) && \text{(A.4)}
\end{aligned}
$$

where (A.3) follows from (4)-(5) and (A.4) uses $\Pr\left(\overline{E_{i,k}^1}\right) \leq k/w$. By observing that (A.4) holds also for $k=0$, we have completed the proof.

Appendix B. Proof of Lemma 1

We will make use of two quantities to prove Lemma 1.

$$
l_j^r \triangleq \sum_{e \in N_G(h_r(j))} p_e, \qquad g_j \triangleq \min_{r \in [d]} l_j^r \ . \tag{B.1}
$$

For a given realization of G, l_j^r is an upper bound on the growth rate of counter $c_j^r(t)$ and g_j is an upper bound on the growth rate of $\hat{n}_j(t)$. To ease the writing, we use A, B, C, and D_j^r as shorthand for events "$h_r(i) = h_r(j)$", "$\hat{n}_j(s-1) = c_i^r(s-1)$", "$g_j \geq p_i$", and "$l_j^r \geq p_i$", respectively. Starting from (9) we write

$$
\begin{aligned}
\mathbb{E}\left[\delta_{i,j}^r\right] &= p_j \Pr\left(A \cap B\right) \\
&= p_j \Big(\Pr\left(A \cap B \cap C\right) + \Pr\left(A \cap B \cap \overline{C}\right)\Big) \\
&\leq p_j \Big(\Pr\left(A \cap C\right) + \Pr\left(A \cap \overline{C}\right) \Pr\left(B \mid A, \overline{C}\right)\Big) \\
&\leq p_j \Big(\Pr\left(A \cap \left(\cap_{e \in [d], e \neq r} D_j^e\right)\right) + \Pr\left(A\right) \Pr\left(B \mid A, \overline{C}\right)\Big) \\
&\leq p_j \Pr\left(A\right) \Big(\Pr\left(D_j^1\right)^{d-1} + \Pr\left(B \mid A, \overline{C}\right)\Big) \ . && \text{(B.2)}
\end{aligned}
$$

We now move to deriving upper bounds on $\Pr\left(D_j^1\right)$ and $\Pr\left(B \mid A, \overline{C}\right)$. For $j \leq i$, we simply write $\Pr\left(D_j^1\right) \leq 1$. For $j > i$, we follow the steps in (A.1)-(A.2):

$$
\begin{aligned}
\Pr\left(D_j^1\right) = \Pr\left(l_j^1 - p_j \geq p_i - p_j\right) &\leq \Pr\left(l_j^1 - p_j \geq p_i - p_j \mid E_{j,k}^1\right) \cdot 1 + 1 \cdot \Pr\left(\overline{E_{j,k}^1}\right) \\
&\leq \frac{\mathbb{E}\left[l_j^1 - p_j \mid E_{j,k}^1\right]}{p_i - p_j} + \frac{k}{w} \ . && \text{(B.3)}
\end{aligned}
$$

We bound $\mathbb{E}\left[l_j^1 - p_j \mid E_{j,k}^1\right]$ similarly to what was done for $\mathbb{E}\left[e_j^1(t)/t \mid E_{j,k}^1\right]$ in (A.2)-(A.4):

$$\mathbb{E}\left[l_j^1 - p_j \mid E_{j,k}^1\right] = \sum_{i \in I \setminus \{j\}} p_i \cdot \Pr\left(i \in N_G(h_r(j)) \mid E_{j,k}\right) \tag{B.4}$$

$$= \sum_{i>k, i \neq j} p_i \cdot \Pr\left(h_r(i) = h_r(j) \mid E_{j,k}\right) \tag{B.5}$$

$$\leq \sum_{i>k, i \neq j} p_i \cdot \frac{1}{w-k} \tag{B.6}$$

We combine (B.3) and (B.6) to find $\Pr\left(D_j^1\right) \leq \mathcal{A}_k(p_i - p_j)$. We again observe that this holds also for $k = 0$, which implies that

$$\Pr\left(D_j^1\right) \leq \begin{cases} 1, & \forall j \leq i, \\ \min\left(\mathcal{A}(p_i - p_j), 1\right), & \forall j > i, \end{cases} \quad \Leftrightarrow \quad \Pr\left(D_j^1\right)^{d-1} \leq \gamma_{i,j} \tag{B.7}$$

where $\gamma_{i,j}$ is given in (24).

To bound $\Pr\left(B \mid A, \overline{C}\right)$, we start by making a change of variable. Let $z = s - 1$. We define the random variable $y_j(z)$ as,

$$y_j(z) \triangleq \sum_{\substack{e \in I \\ h_{r_0}(j) = h_{r_0}(e)}} n_e(z) : \; r_0 = \arg\min_{r \in [d]} l_j^r \;. \tag{B.8}$$

It follows that $\hat{n}_j(z) \leq y_j(z)$. We now use F, J and K as respective shorthand for events "$n_i(z) > y_j(z)$", "$n_i(z) > m(z)$", and "$y_j(z) < m(z)$", with $m(z) \triangleq (p_i + g_j)z/2$. Under the conditioning on A, the equality $c_i^r(z) = c_j^r(z)$ holds. Equation (2) implies then that $c_i^r(z) \geq \hat{n}_j(z)$. The event $\overline{B}$ conditioned on A boils down to $c_i^r(z) > \hat{n}_j(z)$. Since $c_i^r(z) \geq n_i(z)$ and $y_j(z) \geq \hat{n}_j(z)$, the event F implies the event $\overline{B}$ conditionally on A. We then write

$$\begin{aligned} \Pr\left(B \mid A, \overline{C}\right) = 1 - \Pr\left(\overline{B} \mid A, \overline{C}\right) &\leq 1 - \Pr\left(F \mid A, \overline{C}\right) \\ &\leq 1 - \Pr\left(J\right) \Pr\left(K \mid A, \overline{C}\right) \;. \end{aligned}$$

The last step follows from the fact that $n_i(z)$ and $y_j(z)$ are negatively associated [30]. Following (B.8), for every fixed graph realization of G, $y_j(z)$ is the sum of negatively associated random variables [30] and has an expected value $g_j z$ that is less than $m(z)$ under the conditioning that the fixed graph G verifies $g_j < p_i$ (event $\overline{C}$). Thus using Chernoff bounds on events J and K we get

$$\Pr\left(B \mid A, \overline{C}\right) \leq \beta_{i,j} \mathrm{e}^{-\alpha_{i,j} z} \;. \tag{B.9}$$

Using $\Pr\left(A\right) \leq 1/w$, (B.7), and (B.9) in (B.2) we find (23) concluding the proof.

Appendix C. Proof of Proposition 2

The bounds that are valid with CMS are also valid with CMS-CU, thus by Proposition 1, the CCDF with CMS-CU is less than $\mathcal{A}(x)^d$. We focus on proving the bound by $\mathcal{C}_k(x, i, t)$ for $k = 0, \ldots, w-1$. Recall the events $E_{i,k}^r$ of no hash collision in row r between item i and any of the k most popular items (other than i, if $i \leq k$). We define $F_{i,k}$ as their union for all possible rows $r \in [d]$. By the law of total probabilities with respect to the partition $\{F_{i,k}, \overline{F_{i,k}}\}$, we write, for every k,

$$\begin{aligned}\Pr\left(\frac{e_i(t)}{t} \geq x\right) &\leq \Pr\left(\frac{e_i(t)}{t} \geq x \mid F_{i,k}\right) \cdot 1 + 1 \cdot \Pr\left(\overline{F_{i,k}}\right) \\ &\leq \frac{\mathbb{E}\left[e_i(t) \mid F_{i,k}\right]}{xt} + \left(\frac{k}{w}\right)^d . \end{aligned} \tag{C.1}$$

The last step follows from the Markov inequality and the bound $\Pr\left(\overline{E^r_{i,k}}\right) \leq k/w$. To bound $\mathbb{E}\left[e_i(t) \mid F_{i,k}\right]$, we first observe that conditioning on $F_{i,k}$ implies there exists at least a row r_0 such that the event $E^{r_0}_{i,k}$ is true. Using (3) and (5) we write

$$\mathbb{E}\left[e_i(t) \mid F_{i,k}\right] \leq \mathbb{E}\left[e^{r_0}_i(t) \mid F_{i,k}\right] \leq \sum_{s\in[t]} \sum_{\substack{j>k\\ j\neq i}} \mathbb{E}\left[\delta^{r_0}_{i,j}(s) \mid F_{i,k}\right] \tag{C.2}$$

where we used the fact that $\mathbb{E}\left[\delta^{r_0}_{i,j}(s) \mid F_{i,k}\right] = 0$ for $j \leq k$. For $j > k$, the following holds

$$\mathbb{E}\left[\delta^{r_0}_{i,j}(s) \mid F_{i,k}\right] \leq \frac{\mathbb{E}\left[\delta^{r_0}_{i,j}(s)\right]}{1 - \Pr\left(\overline{F_{i,k}}\right)} \leq \frac{\mathbb{E}\left[\delta^{r_0}_{i,j}(s)\right]}{1 - \left(\frac{k}{w}\right)^d}. \tag{C.3}$$

Combining (C.1)-(C.3) with (23) leads to

$$\Pr\left(\frac{e_i(t)}{t} \geq x\right) \leq \frac{\sum_{j>k, j\neq i} p_j \gamma_{i,j}}{xw\left(1 - \left(\frac{k}{w}\right)^d\right)} + \frac{1}{xwt\left(1 - \left(\frac{k}{w}\right)^d\right)} \sum_{\substack{j>k\\ j\neq i}} \frac{p_j \beta_{i,j}}{1 - \mathrm{e}^{-\alpha_{i,j}}} + \left(\frac{k}{w}\right)^d .$$

The second summation is $B_i(k)$ as defined in (27) and we have found (use (26))

$$\Pr\left(\frac{e_i(t)}{t} \geq x\right) \leq \mathcal{C}_k(x, i, t) .$$

As this bound is valid for $k = 0, \ldots, w-1$, we find (25) which concludes the proof.

Appendix D. Discussion on the bound (15)

When all items are requested at least once (which is true under the IRM model for a large enough stream process), the bound (15) proposed in [18] becomes trivial if $w(N) = o(N)$ and both w and N diverge. This happens because

$$\lim_{N\to+\infty} \mathrm{PFP}(A_k) = 1,\ \forall k \in \mathbb{N} \tag{D.1}$$

as we explain next. A_k is computed recursively using (16), with $A_1 = D_1 = N$ since all items are requested at least once. The first computation of (16) requires $\mathrm{PFP}(N)$. We have

$$
\begin{aligned}
\mathrm{PFP}(N) &= \frac{1}{N}\sum_{i=1}^{N}\left(1-\mathrm{e}^{-\frac{i}{w(N)}}\right)^{d} \\
&= \frac{1}{N}\left(N+\sum_{j=1}^{d}\binom{d}{j}(-1)^{j}\sum_{i=1}^{N}\mathrm{e}^{-\frac{ij}{w(N)}}\right) \\
&= 1+\sum_{j=1}^{d}\binom{d}{j}(-1)^{j}\mathrm{e}^{-\frac{j}{w(N)}}\,\frac{1-\mathrm{e}^{-\frac{jN}{w(N)}}}{N\left(1-\mathrm{e}^{-\frac{j}{w(N)}}\right)} .
\end{aligned}
$$

The following then holds

$$
\lim_{N\to+\infty} N\left(1-\mathrm{e}^{-\frac{j}{w(N)}}\right)=+\infty \quad\Rightarrow\quad \lim_{N\to+\infty}\mathrm{PFP}(N)=1\ .
$$

Applying (16) recursively yields (D.1), confirming that the bound (15) boils down to 1.