Better with Fewer Bits: Improving the Performance of Cardinality Estimation of Large Data Streams

Qingjun Xiao
School of Computer Science and Engineering
Southeast University, Nanjing, China.
Email: csqjxiao@seu.edu.cn

You Zhou, Shigang Chen
Department of Computer and Information Science
and Engineering, University of Florida, USA.
Email: {youzhou, sgchen}@cise.ufl.edu

***Abstract*—Cardinality estimation is the task of determining the number of distinct elements (or the cardinality) in a data stream, under a stringent constraint that the input data stream can be scanned by just a single pass. This is a fundamental problem with many practical applications, such as traffic monitoring of high-speed networks and query optimization of Internet-scale database. To solve the problem, we propose an algorithm named HLL-TailCut+, which implements the estimation standard error $1.0/\sqrt{m}$ using the memory units of three bits each, whose cost is much smaller than the five-bit memory units used by HyperLogLog, the best previously known cardinality estimator. This makes it possible to reduce the memory cost of HyperLogLog by 45%. For example, when the target estimation error is 1.1%, state-of-the-art HyperLogLog needs 5.6 kilobytes memory. By contrast, our new algorithm only needs 3 kilobytes memory consumption for attaining the same accuracy. Additionally, our algorithm is able to support the estimation of very large stream cardinalities, even on the Tera and Peta scale.**

I. Introduction

Cardinality estimation is the task of determining the number of distinct elements in a data stream, which is presented as a sequence of elements and can be examined by only one pass. This problem has attracted significant attention over the past decades, due to its important role in many application domains, e.g., real-time traffic monitoring in high-speed networks [4], [10]–[13], [20] or in software-defined networks [21], query plan optimization in large-scale database [9], in-network query aggregation in wireless sensor networks [17], and file significance evaluation in P2P systems [18].

Practical Importance. In the domain of online traffic monitoring of high-speed networks, the cardinality estimation problem can be used to detect traffic anomalies, such as network IP/port scan and distributed denial-of-service (DDoS) attacks [10], [11], [20]. For instance, if we treat all the packets originated from a same source IP as a data stream, then we can detect whether this source IP is a network scanner by counting the number of distinct destination IP/port addresses in its outward packet stream. A similar estimator can be used to detect whether a server is under DDoS attack, if we treat all the packets towards a common destination IP as a data stream and estimate the number of distinct source addresses in this stream. For other application examples, a server farm may learn the popularity of its hosted contents by tracking the number of distinct users that request for each file, and an institutional gateway may perform cardinality estimation on outbound URL requests to measure the popularity of external web content for caching priority.

According to a recent paper [9], many data analysis systems developed by Google, including Sawzall, Dremel and PowerDrill, need to estimate the cardinalities of very large data sets (e.g., the number of distinct search queries on google) on a daily basis. As pointed out in [9], cardinality estimation over large data sets presents a challenge in terms of computational resources, and memory in particular; for PowerDrill, a non-negligible fraction of queries historically could not be computed since they exceeded the available memory.

Prior Art. Although the cardinality can be easily computed using space *linear* in the cardinality, for many applications, this is impractical as it requires too much memory. Therefore, a large number of algorithms have been developed to produce an approximate estimation of the cardinality based on a summary or "sketch" of the data stream, whose occupied space in memory is merely on a *sublinear* level. Typical sketch-based algorithms include PCSA [8], MultiresolutionBitmap [6] (a generalization of LinearCounting [19]), MinCount [2], [3], LogLog [5], HyperLogLog [7], and just list a few.

We make a quick comparison of existing cardinality estimators in Table I. In the third column, each register may be a partial machine word of a few bits, independently producing a coarse estimation of the cardinality (or say, a machine word may hold multiple registers). To mitigate the high variation of a single register and improve the estimation accuracy, a number m of registers must be used. The second column presents the relationship between the standard error and the value of m, where m refers to the number of registers (or the number of bits for MultiresolutionBitmap, or the number of memory units used by MinCount). The total memory cost of an estimator is m multiplied by the size of a register (or 1 bit for MultiresolutionBitmap, or 32 bits for MinCount).

In the last column, we list the memory needed by each algorithm to control the standard error around 2% of the actual cardinality, which shows the progress in memory saving over the past decades: If we use PCSA as the initial benchmark, the seminal work of LogLog reduces the memory cost by more than half. The followup HyperLogLog (HLL) further cuts the memory cost by over 30%. Therefore, HLL is the state-of-the-

TABLE I
A COMPARISON OF POPULAR CARDINALITY ESTIMATORS.

Algorithm	Std. Err.(σ)	Mem Units	Mem(σ=2%)
MinCount	$1.00/\sqrt{m}$	32-bit keys	10000 bytes
PCSA	$0.78/\sqrt{m}$	32-bit registers	6084 bytes
MultiresBitmap	$\approx 4.4/\sqrt{m}$	1 bit	6050 bytes
LogLog	$1.30/\sqrt{m}$	5-bit registers	2641 bytes
HyperLogLog	$1.04/\sqrt{m}$	5-bit registers	1690 bytes
HLL-TailCut	$1.04/\sqrt{m}$	4-bit registers	1352 bytes
HLL-TailCut+	$1.0/\sqrt{m}$	3-bit registers	938 bytes

art algorithm and has been widely adopted by IT industries, such as Google [9], Ask.com [16], PostgreSQL, file-sharing P2P systems [18], and network security systems for DDoS and scan detection [7], just to list a few.

It may appear that the cardinality estimators in Table I already have small memory overhead (on the scale of KBs), and meanwhile can provide good estimation accuracy of about 2% error. Further reducing their memory cost does not seem to be a critically important issue. However, many applications need a large quantity of counters to work simultaneously. Take network traffic anomaly detection as an example. A core router often receives millions of traffic flows in just a few minutes. In order to monitor all the flow behavior, it has to allocate a cardinality estimator for each flow [10], [11], [20]. For Google's applications, the number of counters that work simultaneously becomes much larger, greater than one billion under extreme cases [9]. Hence, the total memory overhead, which is the per-counter memory cost multiplied by the number of counters, will be a huge value that could easily overwhelm the memory available on devices that maintain these counters. For example, on a high-speed router, the on-chip SRAM available for online anomaly detection is merely on the scale of MBs [10], [11], [20], and on Google servers, the DRAM available for tracking keyword popularity is also limited, typically on the scale of GBs for a commodity server [9]. As a summary, reducing the memory cost of a single counter is an important problem with practical value.

Our Contribution. This paper will present a new cardinality estimator named HLL-TailCut+. As shown at the bottom row of Table I, when comparing with the state-of-the-art HyperLogLog, our algorithm can reduce the memory consumption of a single counter again by 45%. A great contribution is that we reduce the size of each register from 5 bits to 3 bits without degrading the accuracy in cardinality estimation, which represents an extreme in compactness that has not been achieved before. Our technique is called *long tail cutoff* that compresses the information across all registers and meanwhile reduces the variance among the registers, which in turn reduces the standard error in cardinality estimation. Consequently, not only do we have smaller-size registers, but also use fewer registers to attain the same accuracy if compared with the previous algorithms [5], [7], [8]. Moreover, unlike HyperLogLog which has limited operating range within 10^9, our algorithm can support the counting of data streams at Tera or Peta scale. It has no estimation bias on the entire measurement range, even when handling small cardinalities.

II. RELATED WORK

The cardinality estimation problem is to count the number of distinct elements in a stream, wherein each element is allowed to appear more than once. A key challenge is that the stream of elements can be scanned by just one pass to obtain the result, due to the constraint of limited processing time or memory.

Linear-Space Solutions. A naive solution for this problem is to use a hash table to memorize all the elements seen so far, in order to filter the duplicated ones. This solution has the advantage of knowing the *exact* cardinality. But it needs memory linear to the stream cardinality, which in most applications, is far too large to be kept in available memory.

A well-known algorithm that can *approximate* the stream cardinality is LinearCounting (LC) [19]. It distributes all the stream elements uniformly among a bit array, so that each element can be encoded as the index of a bit in the array. Duplicated stream elements will be mapped to the same bit index, and hence are filtered automatically. LC can provide the best accuracy among all the known cardinality estimators, however under a strict condition that there is sufficient memory space roughly linear to the cardinality [16]. Otherwise, its accuracy will degrade severely. Since our interest is to estimate very large cardinality values on Giga or Tera scale, LC is no longer attractive, as it requires too much memory.

Sublinear-Space Solutions. Researchers have developed a whole range of algorithms that requires only sublinear memory space [2], [3], [5]–[8]. A frequently used method for reducing memory cost is sampling. An example is MultiresolutionBitmap [6] that designs a sequence of LC structures, whose sampling probabilities decrease exponentially. Another example is MinCount [2], which records only the k smallest hash values for a stream of data items. For both algorithms, their memory efficiency is worse than LogLog and HyperLogLog, as reported by a comparison study [1].

PCSA (Probabilistic Counting with Stochastic Averaging) also prepares a sequence of sampled subsets, but it reduces their sampling probability *exponentially*, until the probability becomes so small that a sampled subset has no data [8]. For the ease of understanding, the sequence of sampled subsets is depicted in Fig. 1 as a sequence of buckets, whose probability of receiving stream elements reduces by the series $2^{-1}, 2^{-2}, 2^{-3}, \ldots, 2^{-w}$. To record whether each bucket has received any stream elements, PCSA allocates a bit array in memory: If a bucket receives nothing and remains empty, its corresponding bit is zero; Otherwise, the bit is one. The $\times$ mark in Fig. 1 represents that a bit is either zero or one.

By maintaining the state of this bit array upon stream element arrivals, PCSA always knows the index of the leftmost empty bucket, which is denoted in Fig. 1 by the symbol M'. Such a bit array is called a *register*, which can give an independent estimation of stream cardinality as $2^{M'}$. Hence, if a PCSA register is given w bits memory, the range

of its estimated cardinality is as large as 2^w, which is a key advantage of PCSA. Of course, the estimation by a single register will be highly inaccurate. For improving the accuracy, PCSA uses a technique called *stochastic averaging* that allocates multiple registers to produce independent estimations, and returns the average value of their estimations.

The memory efficiency of PCSA still leaves much space for improvement: Its register size must be $\log_2 n_{\max} + O(1)$ bits, where $n_{\max}$ is the upper bound of measured cardinality. In contrast, a follow-up algorithm called LogLog reduces the memory per register to only $\log_2 \log_2 n_{\max} + O(1)$ bits [5]. Such significant memory compression is because, instead of maintaining the state of entire bit array like PCSA, LogLog records only the index of the rightmost non-empty bucket, which is denoted by the symbol M in Fig. 1.

HyperLogLog (HLL) is a variant of LogLog for improving accuracy [7]. Both of them depend on the observation of position M shown in Fig. 1, but they adopt different methods for aggregating the estimation results by a set of registers. LogLog uses geometric averaging, while HLL uses harmonic mean, and its purpose is to mitigate the impact of *outlier registers* with abnormally large estimations, thereby appreciably increasing the quality of estimations. As shown in Table I, the expected error of HLL is $1.04/\sqrt{m}$, which is much smaller than that of LogLog $1.30/\sqrt{m}$. In a word, HyperLogLog is the state-of-the-art algorithm.

After years of development, it appears to be very difficult to further compress the memory cost of a cardinality estimator. However, our HLL-TailCut+ estimator can save memory cost of HLL again by 45%, based on a long tail cutting technique to be proposed in this paper. Our estimator can reduce the size of a register to three bits, which is much smaller than the five-bit register used by HLL, and meanwhile it provides the expected relative error of $1.0/\sqrt{m}$. Therefore, our HLL-TailCut+ algorithm can both reduce the per-register memory cost, and discard large outliers to improve accuracy.

III. Traditional HyperLogLog

In this section, we introduce the traditional HyperLogLog (HLL) algorithm by details, and then identify its inadequacies, which motivate the design of our own algorithms.

A. Basic Idea of HyperLogLog

For the ease of understanding, we firstly explain the estimation procedure of a single HLL register. As shown in Fig. 1, when this register receives a stream of elements, it distributes these elements exponentially among a sequence of buckets, i.e., the probability for the buckets to receive elements reduces exponentially by the series $2^{-1}, 2^{-2}, 2^{-3}, \ldots$.

For implementing this exponential distribution, a hash function h is applied to each stream element e. Let us focus on the binary representation of a hash value $h(e)$. The probability of observing the bit pattern $0^{\rho-1}1$ at its beginning is $1/2^\rho$, where ρ is one plus the number of leading zeros. For instance, if the hash value $h(e)$ has no leading zeros, then $\rho(1\ldots) = 1$, and the probability of observing the bit pattern is $1/2^1$. If

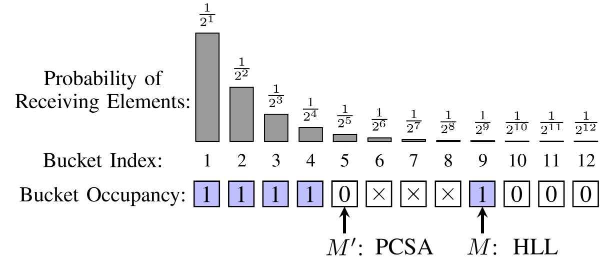

Fig. 1. Observation used by PCSA and HyperLogLog.

there are three leading zeros, then $\rho(0001\ldots) = 4$, and the chance of observing the bit pattern is $1/2^4$. Therefore, we can simply regard the symbol ρ as the index of the bucket a stream element e has been mapped to.

A HLL register will record the largest ρ value for all its input elements, or say, the register will record the position of the rightmost non-empty bucket, which is denoted by M in Fig. 1. Because the probability for this bucket to receive elements is $1/2^M$, intuitively, a good estimation for the number of elements the register receives could be 2^M.

However, the cardinality estimation 2^M by a single register is highly variant. For mitigating the high variance, a technique called *stochastic averaging* is adopted: The input data stream $\mathcal{S}$ is pseudorandomly split into m substreams and then fed into m registers. Each register counts the cardinality of its input substream independently. When needed, their results are aggregated to estimate the cardinality of the data stream $\mathcal{S}$.

B. Detailed Algorithm Procedure

Suppose we have allocated m registers $M_0, M_1, \ldots, M_{m-1}$. The procedure of HyperLogLog can be divided into two parts: an online component that processes each stream element and records critical information into the set of registers, and an offline analysis component that recovers the stream cardinality information from the register set.

Online Component. For an element in stream $\mathcal{S}$, we apply the hash function h to it, and the resultant hash value is denoted by x. For the binary representation of x, let j be its initial p bits, where $p = \log_2 m$ or $m = 2^p$, and let x' be its remaining bits:

$$x = h(e), \qquad j = \langle x_1 x_2 \cdots x_p \rangle, \qquad x' = \langle x_{p+1} x_{p+2} \cdots \rangle.$$

The integer j decides that the register M_j receives this stream element. The integer x' is a hash value that updates M_j:

$$M_j := \max\big(M_j,\, \rho(x')\big), \tag{1}$$

where $:=$ is the assignment operator, and $\max(a, b)$ is a function that returns the greater value of its two parameters. As stated before, $\rho(x')$ is one plus the number of leading zeros in the binary format of x', for instance, $\rho(0001\ldots)=4$. Hence, when the jth substream is nonempty, the register M_j records the index of the *rightmost nonempty bucket* as in Fig. 1.

Offline Analysis Component. Each register M_j in the register set with $0 \le j < m$ can give an estimation 2^{M_j} for the cardinality of its substream. For aggregating the substream cardinalities, HLL uses the normalized harmonic mean:

$$\hat{n} = \alpha_m \cdot m^2 \cdot \big(\textstyle\sum_{0 \le j < m} 2^{-M_j}\big)^{-1}, \tag{2}$$

where α_m is a bias correction constant: $\alpha_{16} = 0.673$, $\alpha_{32} = 0.697$, $\alpha_{64} = 0.709$, $\alpha_m = 0.7213/(1+1.079/m)$ if $m \geq 128$.

C. Shortcomings of HyperLogLog

HyperLogLog is an excellent algorithm that provides the relative standard error $\frac{1.04}{\sqrt{m}}$ at the cost of $5m$ bits memory. Its high accuracy and memory compactness have triggered extensive adoption in IT industries, e.g., Google [9], Ask.com [16] and PostgreSQL. However, this algorithm still possesses two inadequacies which open doors to further improvements.

Threat of Outliers. As mentioned before, the observation used by HyperLogLog, which is the value of each register M_j, is highly variant. To give an impression of the high variance, we illustrate in Fig. 2 the probabilistic distribution for a register to carry an arbitrary value k. The plot (a) is drawn in normal scale, and the plot (b) is drawn by log scale for Y-axis. The mathematical formula of this probabilistic distribution will be described later in Eq. (3). Here, the two plots show that it is a *right-skewed distribution* with a long tail stretching out to the right side of the peak. Note that this property of M_j distribution has no relation with the input stream data. It originates from the uniform distribution of hash function h.

The registers whose value strongly deviates from the peak are called *outliers*, which are most likely to exist on the right-side long tail of the distribution as illustrated in Fig. 2(b). In order to mitigate the impact of the outliers existing on the right tail that have abnormally large register values, HLL adopts harmonic average to aggregate the estimation results of a register set. Our intuition is to completely remove the impact of large outliers, by cutting off the right-side long tail on such a histogram, which contains plenty of outliers instead of useful information. It may appear that the outlier rejection can be easily implemented by discarding the registers whose values are much larger than the average. But the difficulty is how to achieve the tail cutoff when the size of each register is reduced to less than five bits for space saving.

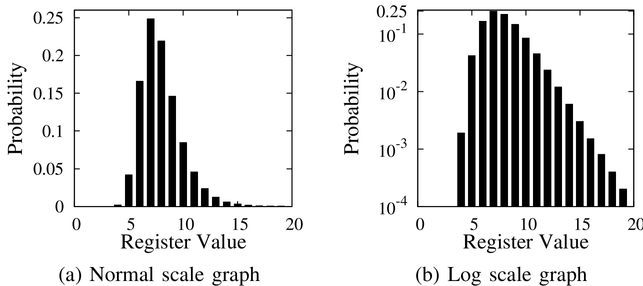

(a) Normal scale graph (b) Log scale graph

Fig. 2. Probability of register values, when $m = 512$ and load factor $\frac{n}{m} = 100$.

Inefficient Register Encoding. The second inadequacy of HyperLogLog is its inefficient encoding of each register state. The size of a HLL register is five bits long, so that the cardinality estimation by a single register can be up to $2^{2^5} \approx 4 \times 10^9$. A recent paper proposed to expand the register size to six bits, in order to support the counting of big data on Tera or Peta scale [9]. On the contrary, we discover that it is possible to reduce the register size to four bits or even to three bits, and meanwhile support the same large operating range.

Our inspiration comes from the Fig. 2(b), where only for the sixteen highest bars between 4 and 19, their probabilities are greater than 0.01%. It implies that, when the number of registers m is on the scale of thousands, the spread of a register set (i.e., the largest register value minus the smallest register value) is less than sixteen in most cases. From the perspective of information theory, it is redundant to use five bits to encode each register, and four bits may be sufficient in most cases.

Moreover, in Fig. 2(a), only for the eight highest bars between 5 and 12, their probabilities are greater than 2%. These eight bars are the most informative part of a histogram, and others are more prone to contain outliers, which implies the possibility of abandoning the right tail for outlier rejection and encoding each register by only three bits memory.

Conclusion. Our basic idea is that the memory per register may be compressed to four bits or even to three bits with no significant loss of useful information. Due to the smaller register size, we can allocate a larger number of registers from the same memory budget, which drives down the estimation error. However, what we have proposed is *lossy compression of registers*, and the challenge is how to avoid its side effect.

IV. MLE-based HyperLogLog

Before introducing our algorithm, we replace the estimation equation of HyperLogLog in (2) with an alternative formula in (5), which is based on MLE (maximum likelihood estimation). The analysis in this section is the theoretical foundation of our own algorithm. Moreover, **HyperLogLog has another inadequacy, which is strongly biased when handling small cardinalities** (see Fig. 6(a) in the simulation section). Our MLE-based substitute can solve this problem and provide unbiased estimations in the entire measurement range.

A. Maximum Likelihood Estimator

In this subsection, we will present a maximum likelihood estimator for the number of distinct elements in a data stream.

We analyze the probabilistic distribution for a HLL register to carry an arbitrary k value, which has been illustrated in Fig. 2, and we have the following theorem.

Theorem 1 (Probability of Register Value)**.** *The probability for a register (for instance, the jth register M_j) to demonstrate a particular value k is approximately*

$$Pr\{M_j = k\} \approx \begin{cases} \left(1 - \frac{1}{m}\right)^n & \text{if } k = 0, \\ (1 - \frac{1}{m2^k})^n - (1 - \frac{1}{m2^{k-1}})^n & \text{if } k \geq 1. \end{cases} \quad (3)$$

Proof. Check Appendix A of the extended version [14]. □

For an arbitrary non-negative k value, let N_k be the number of registers, among the m registers, which carry the k value. If observing exactly N_k registers carrying a particular k value, the probability of this observation is $Pr\{M_j = k\}^{N_k}$, assuming these registers are mutually independent. Then, the combined probability of making the observations

$N_0, N_1, \ldots, N_\infty$ for all the k values from zero to infinity is as follows, under the condition that the stream cardinality is n.

$$Pr\{N_0, N_1, \ldots, N_\infty \mid n\} = \frac{m!}{N_1!N_2!\ldots N_\infty!} \prod_{k=0}^{\infty} Pr\{M_j = k\}^{N_k}$$

However, it is impossible to measure the number of registers N_k for an arbitrarily large k value up to infinity, because each register are given limited memory space (typically 5 bits). We use a symbol $\mathcal{K}$ to characterize a register's up-bound capacity of recording k value. For example, if each register is given 5 bits, then it can record a limited range of k values starting from 0 up to $2^5 - 1 = 31$, and in this case, $\mathcal{K} = 32$. If each register is of 4 bits, then the threshold $\mathcal{K} = 2^4 = 16$. Considering this upper limit of recording k values, the above probability function needs to be modified, assuming only the availability of the observations $N_0, N_1, \ldots, N_{\mathcal{K}-1}$.

$$Pr\{N_0, N_1, \ldots, N_{\mathcal{K}-1} \mid n\} \approx \frac{m!}{N_1!N_2!\ldots N_{\mathcal{K}-1}!} \prod_{k=0}^{\mathcal{K}-1} Pr\{M_j = k\}^{N_k}$$

This probability function is also called the *likelihood* of unknown parameter n, when given the observations about the number of registers carrying each value: $N_0, N_1, \ldots, N_\infty$.

$$\mathcal{L}(n \mid N_0, N_1, \ldots, N_{\mathcal{K}-1}) \approx \frac{m!}{N_1!N_2!\ldots N_{\mathcal{K}-1}!} \prod_{k=0}^{\mathcal{K}-1} Pr\{M_j{=}k\}^{N_k} \quad (4)$$

Applying the well-known *maximum likelihood estimation*, we can find the best n value that maximizes this log-likelihood function, and we use the symbol $\hat{n}$ to denote this optimized estimation of the stream cardinality n.

$$\hat{n} = \arg\max_n \log \mathcal{L}(n \mid N_0, N_1, \ldots, N_{\mathcal{K}-1}) \quad (5)$$

B. Gradient Ascent Solution for MLE

In this subsection, we present our solution to the MLE optimization problem in (5). Although it is viable to solve this problem symbolically by finding the closed-form root to the equation $\frac{\partial}{\partial n} \log \mathcal{L}(n \mid N_0, N_1, \ldots, N_{\mathcal{K}-1}) = 0$, this solution will be complex and have low flexibility (We will demonstrate this point in the next section, when the symbol $\mathcal{K}$ is configured to some other value smaller than 2^5). Therefore, we choose to solve this optimization problem numerically.

We use the following iterative optimization procedure to obtain an optimized estimation of stream cardinality n:

$$\hat{n}^{(i+1)} = \hat{n}^{(i)} + \eta \cdot \frac{\partial}{\partial \hat{n}^{(i)}} \log \mathcal{L}(\hat{n}^{(i)} \mid N_0, N_1, \ldots, N_{\mathcal{K}-1}), \quad (6)$$

where $\frac{\partial}{\partial n} \log \mathcal{L}(n \mid N_0, N_1, \ldots, N_{\mathcal{K}-1})$ is the gradient of log-likelihood function, whose mathematical expression is given in Appendix B, $\hat{n}^{(i)}$ is the current estimation of n, $\hat{n}^{(i+1)}$ is the next-round estimation, η is the optimization step size assigned to $2^{\mathcal{B}} m$, and $\mathcal{B}$ is the smallest value among all registers.

For the above iterative optimization method, its computational cost is only to evaluate the gradient of log-likelihood function for ten or twenty rounds. Moreover, we can speed up its convergence, if we generate a good initial guess using the closed-form cardinality estimator in (2) by HyperLogLog.

V. HLL-TailCut Algorithm

In this section, we reduce the size of each HyperLogLog register from five bits to four bits, and we call the new algorithm HLL-TailCut, because it essentially applies the long tail cutoff technique to the histogram in Fig. 2. HLL-TailCut reduces the memory cost of HLL by 20%. Unlike HLL, it can **support the counting of Tera- or Peta-scale data streams**.

A. Base Register and Offset Registers

Our basic idea is to use a shared base register for storing the smallest value among the set of HyperLogLog registers, so that the m registers only need to store their offsets relative to the base register. Intuitively, the offset stays in a much smaller range than 2^5 (see Fig. 2) and can be encoded by less than five bits. In following, we explain how to maintain the base register and the m offset registers, upon the arrival of stream elements.

Let $\mathcal{B}$ be the shared base register that records the smallest value of the HLL register set. Due to the base register, Eq. (1) that updates each offset register $\tilde{M}_j$ should be changed to

$$\tilde{M}_j := \max\big(\tilde{M}_j,\ \rho(x') - \mathcal{B}\big), \quad (7)$$

where $\tilde{M}_j$ has an upper tilde indicating it is the jth offset register that records the offset of M_j relative to base register $\mathcal{B}$, $\rho(x')$ is the index of the bucket a stream element is mapped to, and $\rho(x') - \mathcal{B}$ is the offset of the bucket index relative to $\mathcal{B}$.

Handle Overflow of Offset Register. We define each offset register to be four bits long. Thus, an offset register's recording capacity $\mathcal{K}$ is $2^4 = 16$, implying that the recorded offset value must be smaller than $\mathcal{K} = 16$. However, occasionally, the offset values $\rho(x') - \mathcal{B}$ of some stream elements are at least $\mathcal{K}$. We use the term "overflow" to refer to the attempts of updating the offset register $\tilde{M}_j$ to the $\mathcal{K}$ value or above.

In order to handle the overflow event, we scan the m offset registers to find the smallest offset value, which is denoted as $\Delta\mathcal{B}$. If $\Delta\mathcal{B}$ is non-zero, it implies that the shared base register $\mathcal{B}$ can be increased by this amount to reduce the offset value stored in each offset register. We call this operation "update the base register": Whenever $\mathcal{B}$ is increased by $\Delta\mathcal{B}$, each offset register $\tilde{M}_j$ needs to be decreased by $\Delta\mathcal{B}$, as they record the offsets to $\mathcal{B}$. Thanks to this base register updating operation, we can easily count the data streams on Tera or Peta scale.

After the increase of the base register by $\Delta\mathcal{B}$, the new offset value $\rho(x') - \mathcal{B}$ in (7) may become smaller than $\mathcal{K}$. If that is true, then the overflow event disappears. Otherwise, the overflow problem can not be resolved, and the jth offset register has to be truncated by the cutoff bound $\mathcal{K}$ as follows.

$$\tilde{M}_j := \max\big(\tilde{M}_j,\ \min(\rho(x') - \mathcal{B},\ \ \mathcal{K} - 1)\big) \quad (8)$$

B. Pseudocode of HLL-TailCut

In this subsection, we describe the procedure of the HLL-TailCut algorithm (also abbreviated as HLL-TC), which can be divided into two parts: an online component that updates the base register $\mathcal{B}$ and offset registers $\tilde{M}_j$, $0 \le j < m$, upon the arrival of stream elements, and an offline component that estimates the stream cardinality n using these registers.

We present the pseudocode of the online component in Algorithm 1. We use the term "truncated register" to refer to the sum of base register $\mathcal{B}$ and offset register $\tilde{M}_j$. The Algorithm 1 essentially maintains a set of truncated registers $\mathcal{B}+\tilde{M}_j$, $0 \le j < m$. Different from the HLL register M_j maintained by (1), these truncated registers chop off the long tail of a histogram (like in Fig. 2) by the bound $\mathcal{B}+\mathcal{K}$, and the chopped part is stacked above the $(\mathcal{B}+\mathcal{K}-1)$th bar, due to the line 9. Thus, the resultant histogram will exhibit an edge peak distribution with a spike close to the tail truncation point.

Algorithm 1: Online Component of HLL-TailCut

```
1 initialize B and M~_j to zero, for each j ∈ [0, m)
2 foreach element e in data stream S do
3     x := h(e),   j := ⟨x_1 x_2 ··· x_b⟩,   x' := ⟨x_{b+1} x_{b+2} ···⟩
4     if ρ(x') − B ≥ K then                 // detect overflow
5         ΔB := min_{0≤j<m} M~_j
6         if ΔB > 0 then
7             B := B + ΔB        // update base register
8             foreach j ∈ [0, m) do M~_j := M~_j − ΔB
9     M~_j := max(M~_j,  min(ρ(x') − B,  K − 1))
```

Since the online component above no longer maintains the HLL register M_j, we need to modify the offline estimation equation in (2), using the newly designed base register $\mathcal{B}$ and offset registers $\tilde{M}_j$. A straightforward solution is to replace M_j by the jth truncated register $\mathcal{B}+\tilde{M}_j$.

$$\hat{n} = \alpha_m \cdot m^2 \cdot \left(\sum_{0\le j<m} 2^{-(\mathcal{B}+\tilde{M}_j)}\right)^{-1} \tag{9}$$

C. Performance Evaluation of HLL-TailCut

In this subsection, we evaluate how the estimated result by (9) is affected, when we truncate the right-side tail of a HLL histogram by $\mathcal{B}+\mathcal{K}$. In Appendix C of extended version [14], we prove HLL-TC in (9) can generate unbiased estimations, if each offset register $\tilde{M}_j$ is given four-bits memory. In Appendix C, we also prove that, if multiple HLL-TC estimators are deployed at different locations, these estimators **can be composed to estimate the union of data streams**.

In following, we use experiments to verify that the tail cutoff across the boundary $\mathcal{B}+16$ has negligibly small impact to the cardinality estimation by (9). We plot the evaluation results in Fig. 3. The subfigure (a) illustrates the estimation bias $E(\hat{n}-n)/n$, where n is the actual cardinality and $\hat{n}$ is the estimated value. The subfigure (b) depicts the relative standard deviation of estimated results $\sqrt{Var(\hat{n})}/E(\hat{n})$. We illustrate both the results of LinearCounting and HLL-TC, which are configured with the same number of memory units: LinearCounting is given m bits, and HLL-TC is given m offset registers. Plot (a) shows that the tail cutoff won't cause severe bias to HLL-TC, when the cutoff bound is $\mathcal{B}+16$.

VI. HLL-TailCut+ Algorithm

In this section, we reduce the size of each offset register to three bits, and save the memory cost by over 40% than HLL.

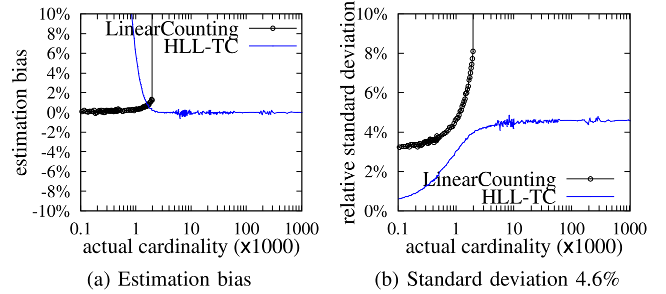

(a) Estimation bias (b) Standard deviation 4.6%

Fig. 3. Performance of HLL-TailCut allocated with 512 offset registers, each of which is given four bits memory.

A. Bias Problem of the Naive HLL-TailCut

When the offset register size reduces to three bits, we assign the cutoff bound $\mathcal{K}$ to $2^3 = 8$, and then we can reuse the online component in Algorithm 1 to maintain the base register $\mathcal{B}$ and each offset register $\tilde{M}_j$, upon the arrival of stream elements. However, the offline analysis component in (9) used by the naive HLL-TailCut has a serious "estimation bias" problem, which will be identified and explained as follows.

HLL-TC adopts an estimation equation in (9) similar to HyperLogLog. For this solution, we illustrate its experimental results in Fig. 4. The subfigure (a) shows that HLL-TC with cutoff bound $\mathcal{K}=8$ produces the estimation bias of -5.2%. This is because, when the offset register is three bits and $\mathcal{K}$ reduces to eight, the percentage of registers truncated by (8), called overflow probability, will greatly increases to about 5%. Thus, a non-negligible fraction of offset registers are truncated.

To make things worse, in Fig. 4(a), the bias of HLL-TC exaggerates to -5.2% by a non-linear curve, implying that we cannot compensate such bias simply by applying a constant corrector to the biased estimation result.

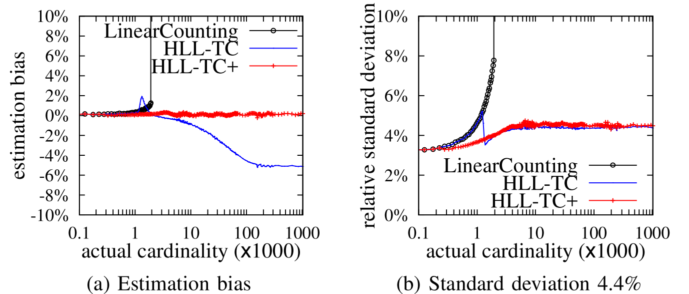

(a) Estimation bias (b) Standard deviation 4.4%

Fig. 4. Performance of HLL-TailCut configured with 512 offset registers, each of which is given three bits.

Interestingly, the standard deviation of HLL-TC decreases from 4.6% shown in Fig. 3(b) to 4.4% shown in Fig. 4(b). This is because more outliers in the long tail are discarded, as the tail cutoff bound changes from 16 to 8. More aggressive outlier rejection brings a small degree of accuracy gain.

B. Probabilistic Model of Truncated Register

To address the negative bias problem of HLL-TC, we will propose a HLL-TailCut+ algorithm, which modifies the MLE-based HyperLogLog algorithm discussed in Section IV. We have already depicted its performance in Fig. 4. Plot (a) shows

that it can provide unbiased estimations in the entire range, and plot (b) shows that it has comparable accuracy with HLL-TC.

Before introducing this new algorithm, in this subsection, we analyze the probability of a truncated register $\mathcal{B} + \tilde{M}_j$ to exhibit an arbitrary k value. The key difficulty for this analysis is that, during the execution time of the online component, the shared register $\mathcal{B}$ is not a fixed value, but gradually increases as the register set receives more and more stream elements.

We begin by defining several notations. Let b be the final value which the base register $\mathcal{B}$ has been updated to. Note that b is typically a small value even for a very large stream at the scale of 10^9. For example, let the number of registers m be 1024. When the stream size n equals 10^2m, the b value alternates between 4 and 5, as in Fig. 7(a). When n increases to 10^6m (about 10^9), the b value grows only to 17 or 18.

As the base register $\mathcal{B}$ undergoes the step-by-step increase in the range of $[0, b]$, the register set consisting of m offset registers will receive different numbers of stream elements.

- When $\mathcal{B}$ is equal to 0, we assume that the register set receives n_0 distinct stream elements.
- When $\mathcal{B}$ is equal to 1, the register set receives n_1 stream elements that are distinct from the previous n_0 elements.
- ...
- When $\mathcal{B}$ is equal to b, the register set receives n_b stream elements that are distinct from the elements $n_0, n_1, \ldots,$ n_{b-1} received when $\mathcal{B}$ is equal to the previous values.

The purpose of our problem is to estimate the total cardinality n of the data stream, which is equal to $n_0 + n_1 + \ldots + n_b$.

Let $\tilde{M}_j^{(0)}, \tilde{M}_j^{(1)}, \ldots, \tilde{M}_j^{(b)}$ be the values of the jth offset register, when the base register $\mathcal{B}$ is fixed to $0, 1, \ldots, b$ and the register set *independently* receives $n_0, n_1, \ldots, n_b$ distinct elements, respectively. For example, $\tilde{M}_j^{(1)}$ is the value of jth offset register, when the base register $\mathcal{B}$ is fixed to 1 and the register set receives n_1 unique elements that are totally different from the n_0 elements received when $\mathcal{B}$ is still zero.

After the register set receives all the $n = n_0 + n_1 + \ldots + n_b$ stream elements, the jth truncated register $\mathcal{B} + \tilde{M}_j$ becomes

$$\mathcal{B} + \tilde{M}_j = \max(\mathcal{B} + \tilde{M}_j^{(0)},\ \mathcal{B} + \tilde{M}_j^{(1)},\ \ldots,\ \mathcal{B} + \tilde{M}_j^{(b)}).$$

Because $\mathcal{B} + \tilde{M}_j^{(0)}, \mathcal{B} + \tilde{M}_j^{(1)}, \ldots, \mathcal{B} + \tilde{M}_j^{(b)}$ are independent, the cumulative probability for $\mathcal{B} + \tilde{M}_j$ (i.e., the probability for the jth truncated register to carry a value of at most k) is

$$Pr\{\mathcal{B} + \tilde{M}_j \le k \mid n_0, n_1, \ldots, n_b\} = \prod_{0 \le i \le b} Pr\{\mathcal{B} + \tilde{M}_j^{(i)} \le k \mid n_i\}. \quad (10)$$

Here, it needs the cumulative distributions of the jth truncated register $Pr\{\mathcal{B} + \tilde{M}_j^{(i)} \le k \mid n_i\}$, when the base register is fixed to a value i ranging from 0 to b. If the base register is equal to b value, the cumulative distributions $Pr\{\mathcal{B} + \tilde{M}_j^{(b)} \le k \mid n_b\}$ is given in the following theorem. When the base register is equal to other values 0, 1, ..., or $b - 1$, we can easily obtain their corresponding cumulative probability, if we replace the symbol b in (11) by 0, 1, ..., or $b - 1$, respectively.

Theorem 2 (Cumulative Distribution of Truncated Register $\mathcal{B} + \tilde{M}_j^{(b)}$ with Fixed Base Register)**.** *When the base register $\mathcal{B}$ is fixed to a value b and the register set receives n_b distinct elements, the probability for the truncated register $\mathcal{B} + \tilde{M}_j^{(b)}$ to exhibit a value of no more than k is as follows.*

$$Pr\{\mathcal{B} + \tilde{M}_j^{(b)} \le k \mid n_b\} \approx \begin{cases} \left(1 - \frac{1}{m2^k}\right)^{n_b} & \text{if } 0 \le k \le b + \mathcal{K} - 2 \\ 1 & \text{if } k \ge b + \mathcal{K} - 1 \end{cases} \quad (11)$$

Proof. Directly derived from Theorem 1. □

By applying (11) to (10), we can obtain the cumulative probability of the jth truncated register $Pr\{\mathcal{B} + \tilde{M}_j \le k \mid n_0, n_1, \ldots, n_b\}$. We refrain from expanding this formula, which otherwise will become too complicated. Then, with the cumulative probability in (10), we can derive the probability density function for the jth truncated register $\mathcal{B} + \tilde{M}_j$.

$$Pr\{\mathcal{B} + \tilde{M}_j = k \mid n_0, n_1, \ldots, n_b\} = \begin{cases} 0 & \text{if } k < b \\ Pr\{\mathcal{B} + \tilde{M}_j \le k \mid n_0, n_1, \ldots, n_b\} & \text{if } k = b \\ Pr\{\mathcal{B} + \tilde{M}_j \le k \mid n_0, n_1, \ldots, n_b\} - \\ Pr\{\mathcal{B} + \tilde{M}_j \le k{-}1 \mid n_0, n_1, \ldots, n_b\} & \text{if } k > b \end{cases} \quad (12)$$

Here, the probability for the truncated register to take a value less than b is zero, because the base register $\mathcal{B}$ increases to b after receiving all the n elements, which makes it impossible for the truncated register $\mathcal{B} + \tilde{M}_j$ to be smaller than b.

C. Maximum Likelihood Estimator

As the probability for a truncated register $\mathcal{B} + \tilde{M}_j$ to carry an arbitrary k value is available in (12), the only problem that remains is how we use this parameterized probabilistic model with b unknown variables $n_0, n_1, \ldots, n_b$, to generate an unbiased estimation of the total stream cardinality n.

We address the problem by estimating the b unknown parameters one by one. When the base register $\mathcal{B}$ is about to increase from zero to one, we estimate n_0, the number of distinct elements received. To accomplish this task, since the base $\mathcal{B}$ is still zero, we can use directly the maximum likelihood estimator in Section IV. Note that when the estimation of n_0 is smaller than m, we will use instead the estimated result by LinearCounting [19] for better accuracy, as inspired by the work [16] that argues LinearCounting is more accurate than HyperLogLog if given enough memory space.

Then, following the principle of mathematical induction, we assume that the stream cardinalities $n_0, n_1, \ldots, n_{b-1}$ all have been estimated as $\hat{n_0}, \hat{n_1}, \ldots, \hat{n_{b-1}}$, at the time that the base register $\mathcal{B}$ is about to update to $1, 2, \ldots, b$, respectively. Based on them, we will further estimate the next unknown variable n_b. The likelihood function of n_b is as follows.

$$\mathcal{L}(n_b \mid N_0, N_1, \ldots, N_{b+\mathcal{K}-1}) = \frac{m!}{N_0! N_1! \ldots N_{b+\mathcal{K}-1}!} \cdot \prod_{k=0}^{b+\mathcal{K}-1} Pr\{\mathcal{B} + \tilde{M}_j = k \mid \hat{n_0}, \hat{n_1}, \ldots, \hat{n_{b-1}}, n_b\}^{N_k} \quad (13)$$

The probability density function $Pr\{\mathcal{B}+\tilde{M}_j = k \mid n_0, n_1, \ldots, n_b\}$ of truncated register $\mathcal{B}+\tilde{M}_j$ is in (12). We replace the true values of $n_0, n_1, \ldots, n_{b-1}$ by their estimated values in (13).

By maximizing the log-likelihood function of n_b, we obtain an optimized estimation of n_b.

$$\hat{n_b} = \arg\max_{n_b} \log \mathcal{L}(n_b \mid N_0, N_1, \ldots, N_{b+\mathcal{K}-1}) \quad (14)$$

To solve this maximum likelihood problem efficiently, we use a steepest-ascent optimization method, which is described in the Appendix D of the extended version [14].

Because the stream cardinalities $n_0, n_1, \ldots, n_b$ all have been estimated, we can obtain an estimation of the total stream cardinality as $\hat{n} = \hat{n_0} + \hat{n_1} + \ldots + \hat{n_b}$. We call this algorithm HLL-TailCut+ (abbreviated as HLL-TC+). Unlike HLL-TC, this algorithm has no bias problem as illustrated in Fig. 4(a).

D. Analysis of Memory Cost

The memory cost of HLL-TC+ is the number of offset registers m multiplied by three bits, and its relative standard error is roughly $\frac{1.0}{\sqrt{m}}$. We obtain this relative error $\frac{1.0}{\sqrt{m}}$ by applying HLL-TC+ to a fixed cardinality (e.g., ten million) for ten thousand times, and then calculating the relative standard deviation of estimated results. Later in Table II, we will use more extensive experiments to verify this relative error equation also applies for other m values and other n values.

Since the standard error of HLL-TC+ is $\frac{1.0}{\sqrt{m}}$ and that of HLL is $\frac{1.04}{\sqrt{m}}$, we can show HLL-TC+ only needs 55% memory of HLL to attain the same accuracy. Let $m_{\text{HLL-TC+}}$ (or m_{HLL}) be the number of registers used by HLL-TC+ (or HLL). Then, we have $\frac{1.0}{\sqrt{m_{\text{HLL-TC+}}}} = \frac{1.04}{\sqrt{m_{\text{HLL}}}}$, to attain the same accuracy. Since the register size of HLL is five bits and that of HLL-TC+ is only three bits, the memory cost of HLL-TC+ divided by that of HLL is $\frac{\text{Memory}_{\text{HLL-TC+}}}{\text{Memory}_{\text{HLL}}} = \frac{3 \text{ bits} \cdot m_{\text{HLL-TC+}}}{5 \text{ bits} \cdot m_{\text{HLL}}} \approx \frac{3}{5} \cdot \left(\frac{1.0}{1.04}\right)^2 \approx 55\%$.

VII. Experiments

In this section, we evaluate the performance of our proposed HLL-TC and HLL-TC+ algorithms, and compare them with state-of-the-art algorithms, including HyperLogLog (HLL) [7] and HyperLogLog+ (HLL+) [9]. Note that we have shared online the source code of all these four algorithms [15].

Experiment Setup. For each cardinality estimator, we will evaluate two performance metrics: the average estimation bias and the average estimation error when given a same amount of memory. We will evaluate the performance of the cardinality estimators under three different scenarios. First, we assume very limited memory budget, no more than a few hundreds bytes per stream, to support cardinality measurements with coarse accuracy ranging from 4% to 10%. Second, we assume the available memory is several kilobytes per stream, which enables highly accurate estimations with the expected errors lower than 2% or even 1%. Third, we would like to verify whether our HLL-TailCut+ estimator can support the measurement of extra large streams whose cardinalities exceed 4×10^9. This bound is important since a five-bit HLL register can only count cardinalities up to $2^{2^5} \approx 4 \times 10^9$.

Coarse-Accuracy Estimation. We consider the coarse accuracy $\sigma \approx 4.4\%$. Then, the number of registers m should be $(1.0/0.044)^2 \approx 512$ for HLL-TC+, occupying $512 \times 3 = 1.54$k bits memory. We give the same amount of memory to the other three algorithms, and depict their performance in Fig. 5.

Fig. 5(a) shows that all four algorithms are approximately unbiased. Fig. 5(b) shows that the estimation error of HLL is slightly smaller than the error of HLL+. This is because HLL defines the register size to be five bits to support the counting of data on Giga scale, while HLL+ enlarges the register size to six bits, to extend the operating range to Tera or Peta scale. Hence, when given the same memory budget, HLL+ can allocate a smaller number of registers than HLL. Fig. 5(b) also shows that our HLL-TC and HLL-TC+ algorithms can provide smaller estimation error. This is because HLL-TC and HLL-TC+ have compressed the register size to four bits and three bits, respectively. Given the same amount of memory, they can allocate more registers to achieve higher accuracy.

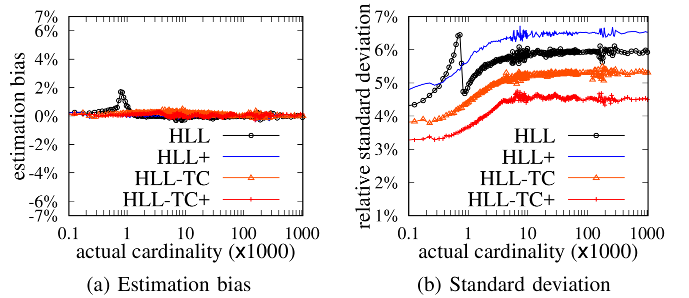

(a) Estimation bias (b) Standard deviation

Fig. 5. Compare cardinality estimators with the same 1.54k bits memory.

Fine-Accuracy Estimation. We consider the estimation error $\sigma \approx 1.1\%$. To achieve such fine accuracy, HLL-TC+ needs about $(1.0/0.011)^2 \approx 8192$ registers, which occupies 24.58k bits memory. We give the same amount of memory to the other three algorithms, and evaluate their performance. Fig. 6(b) shows that HLL-TC+ provide the best accuracy among the four algorithms, and its expected error is $1.0/\sqrt{m} \approx 1.1\%$.

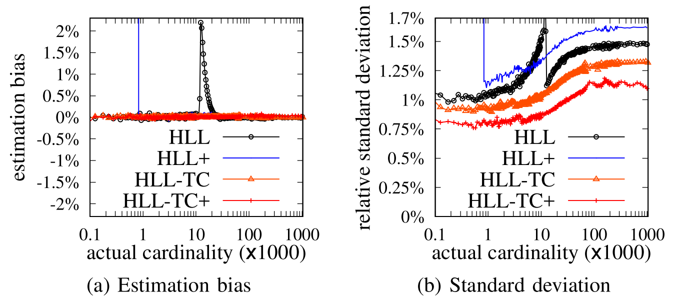

(a) Estimation bias (b) Standard deviation

Fig. 6. Compare cardinality estimators with the same 24.58k bits memory.

Fig. 6(a) shows that HLL has a high spike that is strongly biased. This is because, in the small region around $2.5m = 2.5 \cdot 8192 \cdot \frac{3 \text{ bits}}{5 \text{ bits}} \approx 12288$, HLL makes a switch between LinearCounting and its raw estimation equation in (2). This bias problem has also been elaborated by previous work [9]. As shown in Fig. 6(a), this bias problem has been solved by HLL+, HLL-TC and HLL-TC+, however using different methods. HLL+ corrects the bias in a brute-force way [9]: It

empirically calculates the bias of 200 hundred reference values in the small region from $2m$ to $5m$, and then interpolates between the 200 reference points to determine the correction to apply for any given raw estimation value by (2). In contrast, our HLL-TC addresses this problem elegantly, by substituting the equation (2) with a MLE estimator in (5), within the small region from $2m$ to $5m$. Our HLL-TC+ also does not have the bias problem, because it uses the MLE estimator in (14).

Extra Large Measurement Range. The previous experiments only show the evaluation results for cardinalities up to one million. In following, we will verify that our HLL-TC+ can measure extra large streams that have over four billions distinct elements. Unlike HyperLogLog+ which increases the register size to six bits to support such large streams, we only need an array of three-bits offset registers (whose number is m) plus a single base register which is at least six bits long.

We list in Table II the average estimation bias and error of our HLL-TC+ algorithm, when it is given different numbers of registers m, such as 2^{10}, 2^{12} and 2^{13}. We only show the experimental results of a single stream cardinality value 16×10^9, since it takes days to process such a large data stream for ten thousands times. In this table, the second column lists the average estimation bias of HLL-TC+, which is negligibly small as compared with its standard deviation shown in the third column. This implies that our algorithm can unbiasedly estimate extra large streams beyond the bound of four billions.

TABLE II
APPLY HLL-TC+ TO DATA STREAMS WITH 16×10^9 DISTINCT ELEMENTS.

Register Number (m)	Avg Bias	Std Deviation	Error Eqn
1024	-0.02%	3.13%	$1.00/\sqrt{m}$
4096	-0.06%	1.56%	$1.00/\sqrt{m}$
8192	-0.01%	1.11%	$1.00/\sqrt{m}$

The last column of Table II rewrites the standard deviation of estimated results (shown in the third column) into the form of a constant divided by $\sqrt{m}$. It shows that the standard deviation of our algorithm can be accurately approximated by $1.0/\sqrt{m}$. According to our previous analysis in Section VI-D, if the expected relative error of HLL-TC+ is $1.0/\sqrt{m}$, then it can save 45% memory cost than traditional HyperLogLog.

ACKNOWLEDGEMENT

This work is supported in part by National Natural Science Foundation of China under Grants 61502098, 61632008 and 61320106007, by Jiangsu Provincial Natural Science Foundation of China under Grant BK20150629, by National Science Foundation of United States under Grant STC-1562485, by a grant from Florida Cybersecurity Center, by Jiangsu Provincial Key Lab of Network and Information Security under grant BM2003201, by Key Lab of Computer Network and Information Integration of Ministry of Education of China under Grant 93K-9, and by Collaborative Innovation Center of Novel Software Technology and Industrialization.

VIII. CONCLUSION

This paper studies a fundamental problem called cardinality estimation, in the domain of one-pass processing of streaming data. We present a new solution named HLL-TailCut+, which is able to reduce memory consumption by 45% than the state-of-the-art HyperLogLog. This remarkable improvement originates from a technique we proposed that truncates the right-side long tail of the register distribution of HyperLogLog. This technique brings two key benefits — improve estimation accuracy by rejecting outliers in the long tails, and compress the register size by recording only eight highest bars in the histogram of HLL. Therefore, our algorithm can provide the standard error $\frac{1.0}{\sqrt{m}}$ using only three-bit memory per register. Moreover, this HLL-TailCut+, based on maximum likelihood estimation, can provide approximately unbiased estimations in the entire range of cardinality, even at the point where it switches to LinearCounting for handling small streams.

REFERENCES

[1] K. Aouiche and D. Lemire. A comparison of five probabilistic view-size estimation techniques in OLAP. *Proc. of DOLAP*, 2007.

[2] Z. Bar-yossef, T. S. Jayram, R. Kumar, D. Sivakumar, L. Trevisan, and Luca. Counting distinct elements in a data stream. *Proc. of RANDOM: International Workshop on Randomization and Approximation*, 2002.

[3] K. Beyer, P. J. Haas, B. Reinwald, Y. Sismanis, and R. Gemulla. On synopses for distinct-value estimation under multiset operations. *Proc. of ACM SIGMOD*, 2007.

[4] J. Cao, Y. Jin, A. Chen, T. Bu, and Z. Zhang. Identifying high cardinality internet hosts. *Proc. of IEEE INFOCOM*, April 2009.

[5] M. Durand and P. Flajolet. Loglog counting of large cardinalities. *Proc. of ESA: European Symposia on Algorithms*, 2003.

[6] C. Estan, G. Varghese, and M. Fisk. Bitmap algorithms for counting active flows on high-speed links. *IEEE/ACM Transactions on Networking*, 14(5), Oct. 2006.

[7] P. Flajolet, E. Fusy, O. Gandouet, and et al. HyperLogLog: The analysis of a near-optimal cardinality estimation algorithm. *Proc. of AOFA*, 2007.

[8] P. Flajolet and G. N. Martin. Probabilistic counting algorithms for database applications. *J. Comput. Syst. Sci.*, 31(2), 1985.

[9] S. Heule, M. Nunkesser, and A. Hall. HyperLogLog in practice: Algorithmic engineering of a state-of-the-art cardinality estimation algorithm. *Proc. of EDBT*, 2013.

[10] A. Kumar, M. Sung, J. J. Xu, and J. Wang. Data streaming algorithms for efficient and accurate estimation of flow size distribution. *Proc. of ACM SIGMETRICS*, 2004.

[11] A. Lall, V. Sekar, M. Ogihara, J. Xu, and H. Zhang. Data streaming algorithms for estimating entropy of network traffic. *Proc. of ACM SIGMETRICS Perform. Eval. Rev.*, 34(1), June 2006.

[12] P. Lieven and B. Scheuermann. High-speed per-flow traffic measurement with probabilistic multiplicity counting. *Proc. of INFOCOM*, 2010.

[13] Y. Lu and B. Prabhakar. Robust counting via counter braids: An error-resilient network measurement architecture. *Proc. of INFOCOM*, 2009.

[14] ——. Better with fewer bits: Improving the performance of cardinality estimation of large data streams (extended version online). [url:] https://www.dropbox.com/s/60p90gavz72spwl/infocom17-cll.pdf, 2016.

[15] ——. Source code and tester code of HLL-TailCut+ algorithm. [url:] https://www.dropbox.com/s/l0eaexhzvi34x9u/HLLPlus.zip, 2016.

[16] A. Metwally, D. Agrawal, and A. E. Abbadi. Why go logarithmic if we can go linear?: Towards effective distinct counting of search traffic. *Proc. of EDBT*, 2008.

[17] S. Nath, P. B. Gibbons, S. Seshan, and Z. Anderson. Synopsis diffusion for robust aggregation in sensor networks. *ACM Transactions on Sensor Networks*, 4(2):7:1–7:40, Apr. 2008.

[18] N. Ntarmos, P. Triantafillou, and G. Weikum. Counting at large: Efficient cardinality estimation in internet-scale data networks. *ICDE*, 2006.

[19] K.-Y. Whang, B. T. Vander-Zanden, and H. M. Taylor. A linear-time probabilistic counting algorithm for database applications. *ACM Transactions on Database Systems*, 15(2), June 1990.

[20] M. Yoon, T. Li, S. Chen, and J. Peir. Fit a spread estimator in small memory. *Proc. of IEEE INFOCOM (Review Scores: 5/5/5)*, 2009.

[21] M. Yu, L. Jose, and R. Miao. Software defined traffic measurement with OpenSketch. *Proc. of USENIX NSDI*, 2013.