Utilizing Dynamic Properties of Sharing Bits and Registers to Estimate User Cardinalities over Time

Pinghui Wang^{1,2}, Peng Jia¹, Xiangliang Zhang³, Jing Tao^{1,2,4}, Xiaohong Guan^{2,1,5}, Don Towsley⁶

¹MOE Key Laboratory for Intelligent Networks and Network Security, Xi'an Jiaotong University, China

²Shenzhen Research Institute of Xi'an Jiaotong University, Shenzhen, China

³King Abdullah University of Science and Technology, Thuwal, SA

⁴Zhejiang Research Institute of Xi'an Jiaotong University, Hangzhou, China

⁵Department of Automation and NLIST Lab, Tsinghua University, Beijing, China

⁶School of Computer Science, University of Massachusetts Amherst, MA, USA

Email: {phwang, jtao, xhguan, pengjia}@sei.xjtu.edu.cn, xiangliang.zhang@kaust.edu.sa, towsley@cs.umass.edu

Abstract—Online monitoring user cardinalities (or degrees) in graph streams is fundamental for many applications. For example in a bipartite graph representing user-website visiting activities, user cardinalities (the number of distinct visited websites) are monitored to report network anomalies. These real-world graph streams may contain user-item duplicates and have a huge number of distinct user-item pairs, therefore, it is infeasible to exactly compute user cardinalities when memory and computation resources are limited. Existing methods are designed to approximately estimate user cardinalities, whose accuracy highly depends on parameters that are not easy to set. Moreover, these methods cannot provide anytime-available estimation, as the user cardinalities are computed at the end of the data stream. Real-time applications such as anomaly detection require that user cardinalities are estimated on the fly. To address these problems, we develop novel bit and register sharing algorithms, which use a bit array and a register array to build a compact sketch of all users' connected items respectively. Compared with previous bit and register sharing methods, our algorithms exploit the dynamic properties of the bit and register arrays (e.g., the fraction of zero bits in the bit array at each time) to significantly improve the estimation accuracy, and have low time complexity ($O(1)$) to update the estimations each time they observe a new user-item pair. In addition, our algorithms are simple and easy to use, without requirements to tune any parameter. We evaluate the performance of our methods on real-world datasets. The experimental results demonstrate that our methods are several times more accurate and faster than state-of-the-art methods using the same amount of memory.

I. INTRODUCTION

Many real-world networks are given in the form of graph streams. Calling network is such an example with nodes representing users and an edge representing a call from one user to another. When web surfing activities are modeled as a bipartite graph stream where users and items refer to network hosts and websites respectively, an edge represents a visit by a user to a website. Monitoring the cardinalities (or degrees) of users in these networks is fundamental for many applications such as network anomaly detection [17], [54], [36], [51], where a user's cardinality is defined to be

the number of *distinct* users/items that the user connects to in the regular/bipartite graph stream of interest. Due to the large-size and high-speed nature of these graph streams, it is infeasible to collect the entire graph especially when the computation and memory resources are limited. For example, network routers have fast but very small memories, which leads their traffic monitoring modules incapable to exactly compute the cardinalities of network users. Therefore, it is important to develop fast and memory efficient algorithms to approximately compute user cardinalities over time.

Compared with using a counter to record a user's frequency (i.e., the number of times the user occurred) over time, one needs to build a hash table of distinct occurred edges to handle edge duplicates in graph streams when computing user cardinalities. Therefore, computing user cardinalities is more complex and difficult than computing user frequencies for large data streams, and frequency estimation methods such as Count-Min sketch [10] fails to approximate user cardinalities. To address this challenge, a variety of cardinality estimation methods such as Linear-Time Probabilistic Counting (LPC) [46] and HyperLogLog (HLL) [19] are developed to approximately compute cardinalities. An LPC/HLL sketch consists of m bits/registers, where m is a parameter affecting the estimation accuracy. Since user cardinalities are not known in advance and change over time, one needs to set m large (e.g., thousand) to achieve reasonable accuracy for each user, whose cardinality may vary over a large range. However, this method wastes considerable memory because a large value of m is not necessary for most users, which have small cardinalities. To solve this problem, [54], [50], [43], [47] develop different virtual sketch methods to compress each user's LPC/HLL sketch into a large bit/register array shared by all users. These virtual sketch methods build each user's virtual LPC/HLL sketch using m bits/registers randomly selected from the large bit/register array. This significantly reduces memory usage because each bit/register may be used by more than one user. However, bits/registers in a user's virtual LPC/HLL sketch may be contaminated by other users.

*Peng Jia, Jing Tao and Xiaohong Guan are corresponding authors.

We refer to these bits/registers as “noisy” bits/registers. In practice, most users have small cardinalities and their virtual LPC/HLL sketches tend to contain many “noisy” bits/registers, which results in large estimation errors. Another limitation of existing methods is that they are unable to report user cardinalities on the fly, because they are customized to estimate user cardinalities after all the data has been observed. For real-time applications like on-line anomaly detection, it is important to track user cardinalities in real-time. For example, network monitoring systems are required to detect abnormal IP addresses such as super spreaders (i.e., IP addresses with cardinalities larger than a specified threshold) on the fly. Moreover, online monitoring of IP address cardinalities over time also facilitates online detection of stealthy attacks launched from a subclass of IP addresses.

To address the above challenges, we develop two novel streaming algorithms FreeBS and FreeRS to accurately estimate user cardinalities over time. We summarize our main contributions as:

- Compared with previous bit and register sharing methods, our algorithms FreeBS and FreeRS exploit the dynamic properties of the bit and register arrays over time (e.g., the fraction of zero bits in the bit array at each time) to significantly improve the estimation accuracy. To be more specific, FreeBS/FreeRS allows the number of bits/registers used by a user to dynamically increase as its cardinality increases over time and each user can use all shared bits/registers, which results in more accurate user cardinality estimations.
- Our algorithms report user cardinality estimations on the fly and allow to track user cardinalities in real-time. The time complexity is reduced from $O(m)$ in state-of-the-art methods CSE [50] and vHLL [47] to $O(1)$ for updating user cardinality estimations each time they observe a new user-item pair.
- We evaluate the performance of our methods on real-world datasets. Experimental results demonstrate that our methods are orders of magnitude faster and up to 10,000 times more accurate than state-of-the-art methods using the same amount of memory.

The rest of this paper is organized as follows. The problem formulation is presented in Section II. Section III introduces preliminaries. Section IV presents our algorithms FreeBS and FreeRS. The performance evaluation and testing results are presented in Section V. Section VI summarizes related work. Concluding remarks then follow.

II. PROBLEM FORMULATION

To formally define our problem, we first introduce some notation. Let $\Gamma = e^{(1)}e^{(2)} \dots$ be the graph stream of interest consisting of a sequence of edges. Note that an edge in Γ may appear more than once. In this paper, we focus on bipartite graph streams consisting of edges between users and items. Our methods however easily extend to regular graphs. Let S and D denote the user and item sets, respectively. For $t = 1, 2, \dots$, let $e^{(t)} = (s^{(t)}, d^{(t)})$ denote the t^{th} edge of Γ , where $s^{(t)} \in S$ and $d^{(t)} \in D$ are the user and the item of $e^{(t)}$ respectively. Let $N_s^{(t)}$ denote the set of distinct items

that user s connects to before and including time t . Define $n_s^{(t)} = |N_s^{(t)}|$ to be the cardinality of user s at time t . Then, $n^{(t)} = \sum_{s \in S} |N_s^{(t)}|$ is the sum of all user cardinalities at time t . In this paper, we develop fast and accurate methods for estimating user cardinalities at times $t = 1, 2, \dots$ using a limited amount of memory. When no confusion arises, we omit the superscript (t) to ease exposition.

III. PRELIMINARIES

A. Estimating a Single User's Cardinality

1) *Linear-Time Probabilistic Counting*: For a user $s \in S$, Linear-Time Probabilistic Counting (LPC) [46] builds a sketch B_s to store the set of items that s connects to, i.e., $N_s^{(t)}$. Formally, B_s is defined as a bit array consisting of m bits, which are initialized to zero. Let $h(d)$ be a uniform hash function with range $\{1, \dots, m\}$. When user-item pair (s, d) arrives, the $h(d)^{\text{th}}$ bit in B_s is set to one, i.e., $B_s[h(d)] = 1$. For any bit $B_s[i]$, $1 \leq i \leq m$, the probability that it remains zero at time t is $P(B_s[i] = 0) = (1 - \frac{1}{m})^{n_s^{(t)}}$. Denote by $U_s^{(t)}$ the number of zero bits in B_s at time t . Then, the expectation of $U_s^{(t)}$ is computed as $\mathbb{E}(U_s^{(t)}) = \sum_{i=1}^m P(B_s[i] = 0) \approx m e^{-\frac{n_s^{(t)}}{m}}$. Based on the above equation, when $U_s^{(t)} > 0$, Whang et al. [46] estimate $n_s^{(t)}$ as

$$\hat{n}_s^{(t, \text{LPC})} = -m \ln \frac{U_s^{(t)}}{m}.$$

The range of $\hat{n}_s^{(t, \text{LPC})}$ is $[0, m \ln m]$, and its expectation and variance are computed as

$$\mathbb{E}(\hat{n}_s^{(t, \text{LPC})}) \approx n_s^{(t)} + \frac{1}{2} \left(e^{\frac{n_s^{(t)}}{m}} - \frac{n_s^{(t)}}{m} - 1 \right),$$

$$\text{Var}(\hat{n}_s^{(t, \text{LPC})}) \approx m \left(e^{\frac{n_s^{(t)}}{m}} - \frac{n_s^{(t)}}{m} - 1 \right).$$

2) *HyperLogLog*: To estimate the cardinality of user s , HyperLogLog (HLL) [19] is developed based on the Flajolet-Martin (FM) sketch [20] R_s consisting of m registers $R_s[1], \dots, R_s[m]$. All m registers are initialized to 0. For $1 \leq i \leq m$, let $R_s^{(t)}[i]$ be the value of $R_s[i]$ at time t . When a user-item pair (s, d) arrives, HLL maps the item into a pair of random variables $h(d)$ and $\rho(d)$, where $h(d)$ is an integer uniformly selected from $\{1, \dots, m\}$ at random and $\rho(d)$ is drawn from a *Geometric*(1/2) distribution, $P(\rho(d) = k) = \frac{1}{2^k}$ for $k = 1, 2, \dots$ ¹ Then, register $R_s[h(d)]$ is updated as

$$R_s^{(t)}[h(d)] \leftarrow \max\{R_s^{(t-1)}[h(d)], \rho(d)\}.$$

At time t , Flajolet et al. [19] estimate $n_s^{(t)}$ as

$$\hat{n}_s^{(t, \text{HLL})} = \frac{\alpha_m m^2}{\sum_{i=1}^m 2^{-R^{(t)}[i]}},$$

¹Functions $h(d)$ and $\rho(d)$ are usually implemented as: Let $\Phi(d) = \langle x_1 x_2 \dots \rangle$ be the binary format of the output of a hash function $\Phi(d)$, and $b = \lceil \log_2 m \rceil$. Then, $h(d)$ is defined as $h(d) = (x_1 x_2 \dots x_b \bmod m) + 1$ and $\rho(d)$ is defined as the number of leading zeros in $\langle x_{b+1} x_{b+2} \dots \rangle$ plus one.

Figure 1. Overview of bit sharing method CSE and register sharing method vHLL. Virtual CSE/vHLL sketches of users may contain “noisy” bits/registers (e.g., the bit and register in red and bold in the figure).

where α_m is the following constant to correct for bias

$$\alpha_m = \left(m \int_0^\infty \left(\log_2 \frac{2+x}{1+x} \right)^m dx \right)^{-1}.$$

The above formula for α_m is complicated. In practice, α_m is computed numerically, e.g., $\alpha_{16} \approx 0.673$, $\alpha_{32} \approx 0.697$, $\alpha_{64} \approx 0.709$, and $\alpha_m \approx 0.7213/(1 + 1.079/m)$ for $m \geq 128$. The error of $\hat{n}_s^{(t, \text{HLL})}$ is analyzed as

$$\lim_{n_s^{(t)} \rightarrow \infty} \frac{\mathbb{E}(\hat{n}_s^{(t, \text{HLL})})}{n_s^{(t)}} = 1 + \delta_1(n_s^{(t)}) + o(1),$$

$$\lim_{n_s^{(t)} \rightarrow \infty} \frac{\sqrt{\text{Var}(\hat{n}_s^{(t, \text{HLL})})}}{n_s^{(t)}} = \frac{\beta_m}{\sqrt{m}} + \delta_2(n_s^{(t)}) + o(1).$$

where δ_1 and δ_2 represent oscillating functions of a tiny amplitude (e.g., $|\delta_1(n_s^{(t)})| < 5 \times 10^{-5}$ and $|\delta_2(n_s^{(t)})| < 5 \times 10^{-4}$ as soon as $m \geq 16$) which can be safely neglected for all practical purposes, and β_m is a constant for a specific m , e.g., $\beta_{16} \approx 1.106$, $\beta_{32} \approx 1.070$, $\beta_{64} \approx 1.054$, $\beta_{128} \approx 1.046$, and $\beta_\infty \approx 1.039$.

Since $\hat{n}_s^{(t, \text{HLL})}$ is severely biased for small cardinalities, HLL treats the HLL sketch R_s as an LPC sketch (i.e., a bitmap of m bits) when $\frac{\alpha_m m^2}{\sum_{i=1}^m 2^{-R_s^{(t)}[i]}} < 2.5m$. In this case, $n_s^{(t)}$ is estimated as $-m \ln \frac{\hat{U}_s^{(t)}}{m}$, where $\hat{U}_s^{(t)}$ is the number of registers among $R_s[1], \dots, R_s[m]$ that equal 0 at time t . Therefore, we easily find that LPC outperforms HLL for small cardinalities under the same memory usage.

3) *Discussions*: To compute all user cardinalities, one can use an LPC or HLL sketch to estimate each user cardinality. Clearly, using a small m , LPC and HLL will exhibit large errors for users with large cardinalities. Most user cardinalities are small and assigning an LPC/HLL sketch with large m to each user in order to accurately estimate large user cardinalities is wasteful as LPC and HLL do not require to set a large m to achieve reasonable estimation accuracy for users with small cardinalities. In practice, the user cardinalities are

not known in advance and vary over time. Therefore, it is difficult to set an optimal value of m when using LPC and HLL to estimate all user cardinalities. In the next subsection, we introduce state-of-the-art methods to address this problem, and also discuss their shortcomings.

B. Estimating All User Cardinalities

1) *CSE: Compressing LPC Sketches of All Users into a Shared Bit Array*: As shown in Figure 1 (a), CSE [50] consists of a large bit array A and m independent hash functions $f_1(s), \dots, f_m(s)$, each mapping users to $\{1, \dots, M\}$, where M is the length of the one dimensional bit array A . Similar to LPC, CSE builds a virtual LPC sketch for each user and embeds LPC sketches of all users in A . For user s , its virtual LPC sketch $\hat{B}_s$ consists of m bits selected randomly from A by the group of hash functions $f_1(s), \dots, f_m(s)$, that is

$$\hat{B}_s = (A[f_1(s)], \dots, A[f_m(s)]).$$

Each bit in A is initially set to zero. When a user-item pair (s, d) arrives, CSE sets the $h(d)$ th bit in $\hat{B}_s$ to one. Similar to LPC, $h(d)$ is a uniform hash function with range $\{1, \dots, m\}$. Since the $h(d)$ th element in $\hat{B}_s$ is bit $A[f_{h(d)}(s)]$, CSE only needs to set bit $A[f_{h(d)}(s)]$, i.e., $A[f_{h(d)}(s)] \leftarrow 1$. Let $\hat{U}_s^{(t)}$ be the number of zero bits in $\hat{B}_s$ and $U^{(t)}$ be the number of zero bits in A at time t . A user’s virtual LPC sketch can be viewed as a regular LPC sketch containing “noisy” bits (e.g., the bit in red and bold in Figure 1 (a)), that are wrongly set from zero to one by items of other users. To remove the estimation error introduced by “noisy” bits, Yoon et al. [50] estimate $n_s^{(t)}$ as

$$\hat{n}_s^{(t, \text{CSE})} = -m \ln \frac{\hat{U}_s^{(t)}}{m} + m \ln \frac{U^{(t)}}{M}.$$

On the right-hand side of the above equation, the first term is the same as the regular LPC, and the second term corrects the error introduced by “noisy” bits. The bias and variance of $\hat{n}_s^{(t, \text{CSE})}$ are given by eqs. (23) and (24) in the original paper [50].

2) *vHLL: Compressing HLL Sketches of All Users into a Shared Bit Array*: Xiao et al. [47] develop a register sharing method, vHLL, which extends the HLL method to estimate all user cardinalities. vHLL consists of a list of M registers $R[1], \dots, R[M]$, which are initialized to zero. To maintain the virtual HLL sketch $\hat{R}_s$ of a user s , vHLL uses m independent hash functions $f_1(s), \dots, f_m(s)$ to randomly select m registers from all registers $R[1], \dots, R[M]$, where each function $f_1(s), \dots, f_m(s)$ maps users to $\{1, \dots, M\}$. Formally, $\hat{R}_s$ is defined as

$$\hat{R}_s = (R[f_1(s)], \dots, R[f_m(s)]).$$

For $1 \leq i \leq M$, let $R^{(t)}[i]$ be the value of $R[i]$ at time t . When a user-item pair (s, d) arrives, it maps the item to a pair of two random variables $h(d)$ and $\rho(d)$, where $h(d)$ is an integer uniformly selected from $\{1, \dots, m\}$ at random, and $\rho(d)$ is a random integer drawn from a *Geometric*(1/2) distribution, which is similar to HLL. We can easily find that the $h(d)^{\text{th}}$ element in the virtual HLL sketch of user s is $R[f_{h(d)}(s)]$, therefore, vHLL only needs to update register $R[f_{h(d)}(s)]$ as

$$R^{(t)}[f_{h(d)}(s)] \leftarrow \max\{R^{(t-1)}[f_{h(d)}(s)], \rho(d)\}.$$

A user's virtual HLL sketch can be viewed as a regular HLL containing "noisy" registers (e.g., the register in red and bold in Figure 1 (b)), which are wrongly set by items of other users. To remove the estimation error introduced by "noisy" registers, Xiao et al. [47] estimate $n_s^{(t)}$ as

$$\hat{n}_s^{(t, \text{vHLL})} = \frac{M}{M-m} \left(\frac{\alpha_m m^2}{\sum_{i=1}^m 2^{-R^{(t)}[f_i(s)]}} - \frac{m \alpha_M M}{\sum_{i=1}^M 2^{-R^{(t)}[i]}} \right),$$

where α_m is the same as that of HLL. For the two terms between the parentheses on the right-hand side of the above equation, the first term is the same as the regular HLL, and the second term corrects the error introduced by "noisy" registers. Similar to the regular HLL, the first term is replaced by $-m \ln \frac{\hat{U}_s^{(t)}}{m}$ when $\frac{\alpha_m m^2}{\sum_{i=1}^m 2^{-R^{(t)}[f_i(s)]}} < 2.5m$

where $\hat{U}_s^{(t)}$ is the number of registers among $R[f_1(s)], \dots, R[f_m(s)]$ that equal 0 at time t . The expectation of $\hat{n}_s^{(t, \text{vHLL})}$ approximately equals $n_s^{(t)}$, that is, $\mathbb{E}(\hat{n}_s^{(t, \text{vHLL})}) \approx n_s^{(t)}$. The variance of $\hat{n}_s^{(t, \text{vHLL})}$ is approximately computed as $\text{Var}(\hat{n}_s^{(t, \text{vHLL})}) \approx \frac{M^2}{(M-m)^2} \left(\frac{1.04^2}{m} (n_s^{(t)} + (n^{(t)} - n_s^{(t)}) \frac{m}{M})^2 + (n^{(t)} - n_s^{(t)}) \frac{m}{M} (1 - \frac{m}{M}) + \frac{(1.04 n^{(t)} m)^2}{M^3} \right)$, where $n^{(t)} = \sum_{s \in S} n_s^{(t)}$ counts the total number of distinct user-item pairs occurred before and including time t .

C. Unsolved Challenges

Challenge 1: It is difficult to set parameter m for both CSE and vHLL. The estimation accuracy of CSE and vHLL highly depends on the value of m , as we can see in above discussions. Increasing m introduces more "unused" bits in virtual LPC sketches of occurred users, which can become contaminated with noise. Here "unused" bits refer to the bits in a user's virtual LPC sketch that no user-item pairs of

user are hashed into. However, decreasing m introduces large estimation errors for users with large cardinalities. Similarly, vHLL also confronts the same challenge in determining an optimal m . Later our experimental results will also verify that errors increase with m for users with small cardinalities under CSE and vHLL.

Challenge 2: It is computationally intensive to estimate user cardinalities for all values of t . At each time, both CSE and vHLL require time complexity $O(m)$ to compute the cardinality of a single user. When applied to compute cardinalities for all users in S at all times, CSE and vHLL have to be repeatedly called and will incur high computational cost, which prohibits their application to high speed streams in an on-line manner.

IV. OUR METHODS

In this section, we present our streaming algorithms FreeBS and FreeRS for estimating user cardinalities over time. FreeBS and FreeRS are designed based on two novel bit sharing and register sharing techniques, respectively. The basic idea behind our methods can be summarized as: Unlike vHLL/CSE mapping each user's items into $m \ll M$ bits/registers, FreeBS/FreeRS randomly maps user-item pairs into all M bits/registers in the bit/register array. Thus, users with larger cardinalities (i.e., connecting to a larger number of items) tend to use more bits/registers. For each user-item pair $e^{(t)} = (s^{(t)}, d^{(t)})$ occurred at time t , we discard it when updating $e^{(t)}$ does not change any bit/register in the bit/register array shared by all users. Otherwise, $e^{(t)}$ is a new user-item pair that does not occur before time t , and we increase the cardinality estimation of user $s^{(t)}$ by $\frac{1}{q^{(t)}}$, where $q^{(t)}$ is defined as the probability that a new user-item pair changes any bit/register in the bit/register array at time t . To some extent, the above procedure can be viewed as a sampling method such that each new user-item pair arriving at time t is sampled with probability $q^{(t)}$ and each user's cardinality is estimated using the Horvitz-Thompson estimator [25].

A. FreeBS: Parameter-Free Bit Sharing

Data Structure. The pseudo-code for FreeBS is shown as Algorithm 1. FreeBS consists of a one-dimensional bit array B of length M , where each bit $B[j]$, $1 \leq j \leq M$, is initialized to zero. In addition, FreeBS uses a hash function $h^*(e)$ to uniformly and independently map each user-item pair $e = (s, d)$ to an integer in $\{1, \dots, M\}$ at random, i.e., $P(h^*(e) = i) = \frac{1}{M}$, and $P(h^*(e) = i \wedge h^*(e') = i' | e \neq e') = \frac{1}{M^2}$, $i, i' \in \{1, \dots, M\}$. Note that $h^*(e)$ differs from the hash function $h(s)$ used by CSE that maps **user** s to an integer in $\{1, \dots, M\}$ at random.

Update Procedure. When a user-item pair $e = (s, d)$ arrives at time t , FreeBS first computes a random variable $h^*(e)$, and then sets $B[h^*(e)]$ to one, i.e., $B[h^*(e)] \leftarrow 1$. Let $B_0^{(t)}$ denote the set of the indices corresponding to zero bits in B at time t . Formally, $B_0^{(t)}$ is defined as $B_0^{(t)} = \{i : B^{(t)}[i] = 0, 1 \leq i \leq M\}$. Let $\hat{n}_s^{(t, \text{FreeBS})}$ denote the cardinality estimate for user s at time t . We initialize $\hat{n}_s^{(0, \text{FreeBS})}$ to 0. Next, we describe

Algorithm 1: The pseudo-code for FreeBS.

```

 $B[1, \dots, M] \leftarrow [0, \dots, 0];$ 
 $\hat{n}_s^{(\text{FreeBS})} \leftarrow 0, s \in S;$ 
 $m_0 \leftarrow M;$ 
foreach  $e = (s, d)$  in  $\Gamma$  do
    if  $B[h^*(e)] == 0$  then
         $B[h^*(e)] \leftarrow 1;$ 
         $\hat{n}_s^{(\text{FreeBS})} \leftarrow \hat{n}_s^{(\text{FreeBS})} + \frac{M}{m_0};$ 
         $m_0 \leftarrow m_0 - 1;$ 
    end
end

```

how $\hat{n}_s^{(t, \text{FreeBS})}$ is computed on-line. Let $m_0^{(t)} = |B_0^{(t)}|$ denote the number of zero bits in B at time t . Let $q_B^{(t)}$ denote the probability of e changing a bit in B from 0 to 1 at time t . Formally, $q_B^{(t)}$ is defined as

$$q_B^{(t)} = \sum_{i \in B_0^{(t-1)}} P(h^*(e) = i) = \frac{|B_0^{(t-1)}|}{M} = \frac{m_0^{(t-1)}}{M}.$$

Let $\mathbf{1}(\mathbb{P})$ denote the indicator function that equals 1 when predicate $\mathbb{P}$ is true and 0 otherwise. Besides setting $B[h^*(e)]$ to 1 at time t with the arrival of user-item pair $e = (s, d)$, we also update the cardinality estimate of user s as

$$\hat{n}_s^{(t, \text{FreeBS})} \leftarrow \hat{n}_s^{(t-1, \text{FreeBS})} + \frac{\mathbf{1}(B^{(t-1)}[h^*(e)] = 0)}{q_B^{(t)}}.$$

For any other user $s' \in S \setminus \{s\}$, we keep its cardinality estimate unchanged, i.e., $\hat{n}_{s'}^{(t, \text{FreeBS})} \leftarrow \hat{n}_{s'}^{(t-1, \text{FreeBS})}$. We easily find that $q_B^{(t)}$ can be fast computed incrementally. That is, we initialize $q_B^{(1)}$ to 1, and incrementally compute $q_B^{(t+1)}$ as

$$q_B^{(t+1)} \leftarrow q_B^{(t)} - \frac{\mathbf{1}(B^{(t-1)}[h^*(e)] = 0)}{M}, \quad t \geq 1.$$

Hence, the time complexity of FreeBS for processing each user-item pair is $O(1)$.

Error Analysis. Let $T_s^{(t)}$ denote the set of the first occurrence times of user-item pairs associated with user s in stream Γ . Formally, we define $T_s^{(t)}$ as

$$T_s^{(t)} = \{i : s^{(i)} = s \wedge e^{(j)} \neq e^{(i)}, 0 < j < i \leq t\}.$$

Theorem 1: The expectation and variance of $\hat{n}_s^{(t, \text{FreeBS})}$ are

$$\mathbb{E}(\hat{n}_s^{(t, \text{FreeBS})}) = n_s^{(t)},$$

$$\text{Var}(\hat{n}_s^{(t, \text{FreeBS})}) = \sum_{i \in T_s^{(t)}} \mathbb{E}\left(\frac{1}{q_B^{(i)}}\right) - n_s^{(t)} \leq n_s^{(t)} \left(\mathbb{E}\left(\frac{1}{q_B^{(t)}}\right) - 1 \right),$$

where

$$\begin{aligned} \mathbb{E}\left(\frac{1}{q_B^{(i)}}\right) &= \sum_{j=1}^M \frac{M \binom{n^{(i)}}{j} \sum_{k=0}^{j-1} (-1)^k \binom{j}{k} \left(\frac{j-k}{M}\right)^{n^{(i)}}}{M-j} \\ &\approx e^{\frac{n^{(i)}}{M}} \left(1 + \frac{1}{M} \left(e^{\frac{n^{(i)}}{M}} - \frac{n^{(i)}}{M} - 1 \right) \right). \end{aligned}$$

Proof. Let δ_e denote an indicator that equals 1 when updating a user-item pair e incurs a value change of $B[h^*(e)]$ (i.e., $B[h^*(e)]$ changes from 0 to 1), and 0 otherwise. We easily have

$$\hat{n}_s^{(t, \text{FreeBS})} = \sum_{i \in T_s^{(t)}} \frac{\delta_{e^{(i)}}}{q_B^{(i)}}.$$

For each $\delta_{e^{(i)}}$, we have

$$\mathbb{E}(\delta_{e^{(i)}} | q_B^{(i)}) = q_B^{(i)}, \quad 1 \leq i \leq t,$$

$$\text{Var}(\delta_{e^{(i)}} | q_B^{(i)}) = \mathbb{E}(\delta_{e^{(i)}}^2 | q_B^{(i)}) - (\mathbb{E}(\delta_{e^{(i)}} | q_B^{(i)}))^2 = q_B^{(i)} - (q_B^{(i)})^2.$$

Define $Q_B^{(t)} = \{q_B^{(1)}, \dots, q_B^{(t)}\}$. Then, we have

$$\begin{aligned} \mathbb{E}(\hat{n}_s^{(t, \text{FreeBS})} | Q_B^{(t)}) &= \mathbb{E}\left(\sum_{i \in T_s^{(t)}} \frac{\delta_{e^{(i)}}}{q_B^{(i)}} \middle| Q_B^{(t)}\right) \\ &= \sum_{i \in T_s^{(t)}} \frac{\mathbb{E}(\delta_{e^{(i)}} | Q_B^{(t)})}{q_B^{(i)}} \\ &= \sum_{i \in T_s^{(t)}} \frac{P(\delta_{e^{(i)}} = 1 | Q_B^{(t)})}{q_B^{(i)}} \\ &= n_s^{(t)}. \end{aligned}$$

Given $Q_B^{(t)}$, random variables $\delta_{e^{(i)}}$, $i \in T_s^{(t)}$, are independent of each other. Then, we have

$$\mathbb{E}(\hat{n}_s^{(t, \text{FreeBS})}) = \mathbb{E}(\mathbb{E}(\hat{n}_s^{(t, \text{FreeBS})} | Q_B^{(t)})) = \mathbb{E}(n_s^{(t)}) = n_s^{(t)}.$$

The variance of $\hat{n}_s^{(t, \text{FreeBS})}$ is computed as

$$\begin{aligned} \text{Var}(\hat{n}_s^{(t, \text{FreeBS})} | Q_B^{(t)}) &= \text{Var}\left(\sum_{i \in T_s^{(t)}} \frac{\delta_{e^{(i)}}}{q_B^{(i)}} \middle| Q_B^{(t)}\right) \\ &= \sum_{i \in T_s^{(t)}} \frac{\text{Var}(\delta_{e^{(i)}} | Q_B^{(t)})}{(q_B^{(i)})^2} \\ &= \sum_{i \in T_s^{(t)}} \frac{q_B^{(i)} - (q_B^{(i)})^2}{(q_B^{(i)})^2} \\ &= \sum_{i \in T_s^{(t)}} \frac{1}{q_B^{(i)}} - n_s^{(t)}. \end{aligned}$$

Since $\text{Var}(\mathbb{E}(\hat{n}_s^{(t, \text{FreeBS})} | Q_B^{(t)})) = 0$, using the equation $\text{Var}(X) = \text{Var}(\mathbb{E}(X|Y)) + \mathbb{E}(\text{Var}(X|Y))$, we have

$$\begin{aligned} \text{Var}(\hat{n}_s^{(t, \text{FreeBS})}) &= \mathbb{E}(\text{Var}(\hat{n}_s^{(t, \text{FreeBS})} | Q_B^{(t)})) \\ &= \mathbb{E}\left(\sum_{i \in T_s^{(t)}} \frac{1}{q_B^{(i)}}\right) - n_s^{(t)} \\ &= \sum_{i \in T_s^{(t)}} \mathbb{E}\left(\frac{1}{q_B^{(i)}}\right) - n_s^{(t)}. \end{aligned}$$

In what follows we derive the formula for $\mathbb{E}(\frac{1}{q_B^{(i)}})$. For j specific distinct bits in B , there exist $j! \tau(n^{(i)}, j)$ ways to map $n^{(i)}$ distinct user-item pairs occurred in stream Γ before and

including time i into these bits given that each bit has at least one user-item pair, where $\tau(n^{(i)}, j)$, the Stirling number of the second kind [2], is computed as

$$\tau(n^{(i)}, j) = \sum_{k=0}^{j-1} (-1)^k \binom{j}{k} (j-k)^{n^{(i)}}, \quad 0 < j \leq n^{(i)}.$$

In addition, there exist $\binom{M}{j}$ ways to select j distinct bits from B , therefore we have

$$P(m_0^{(i)} = M - j | n^{(i)}) = \frac{\binom{M}{j} j! \tau(n^{(i)}, j)}{M^{n^{(i)}}}.$$

Then, we have

$$\mathbb{E}\left(\frac{1}{q_B^{(i)}}\right) = \sum_{j=1}^M \frac{M \binom{n^{(i)}}{j} \sum_{k=0}^{j-1} (-1)^k \binom{j}{k} \left(\frac{j-k}{M}\right)^{n^{(i)}}}{M - j}.$$

Next, we introduce a method to approximately compute $\mathbb{E}(\frac{1}{q_B^{(i)}})$. We expand the function $\frac{1}{q_B^{(i)}}$ by its Taylor series around $\mathbb{E}(q_B^{(i)})$ as

$$\begin{aligned} \mathbb{E}\left(\frac{1}{q_B^{(i)}}\right) &\approx \mathbb{E}\left(\frac{1}{\mathbb{E}(q_B^{(i)})} - \frac{q_B^{(i)} - \mathbb{E}(q_B^{(i)})}{(\mathbb{E}(q_B^{(i)}))^2} + \frac{(q_B^{(i)} - \mathbb{E}(q_B^{(i)}))^2}{(\mathbb{E}(q_B^{(i)}))^3}\right) \\ &= \frac{1}{\mathbb{E}(q_B^{(i)})} + \frac{\text{Var}(q_B^{(i)})}{(\mathbb{E}(q_B^{(i)}))^3}. \end{aligned}$$

From [46] (eqs.(5) and (6) in [46]), we easily have $\mathbb{E}(q_B^{(i)}) = e^{-\frac{n^{(i)}}{M}}$ and $\text{Var}(q_B^{(i)}) = \frac{1}{M} e^{-\frac{n^{(i)}}{M}} (1 - (1 + \frac{n^{(i)}}{M}) e^{-\frac{n^{(i)}}{M}})$. Then, we obtain $\mathbb{E}(q_B^{(i)}) \approx e^{\frac{n^{(i)}}{M}} (1 + \frac{1}{M} (e^{\frac{n^{(i)}}{M}} - \frac{n^{(i)}}{M} - 1))$. $\square$

B. FreeRS: Parameter-Free Register Sharing

Data Structure. The pseudo-code for FreeRS is shown as Algorithm 2. FreeRS consists of M registers $R[1], \dots, R[M]$, which are initialized to zero. In addition, FreeRS also uses a hash function $h^*(e)$ to randomly map each user-item pair $e = (s, d)$ to an integer in $\{1, \dots, M\}$ and another function $\rho^*(e)$ that maps e to a random integer in $\{1, 2, \dots\}$ according to a *Geometric*(1/2) distribution. Note that $h^*(e)$ and $\rho^*(e)$ differ from hash functions $h(s)$ and $\rho(s)$ used by vHLL, which map user u to a random integer in $\{1, \dots, M\}$ and $\{1, 2, \dots\}$, respectively.

Update Procedure. When user-item pair $e = (s, d)$ arrives at time t , FreeRS first computes two random variables $h^*(e)$ and $\rho^*(e)$, and then updates $R^{(t)}[h^*(e)]$ as

$$R^{(t)}[h^*(e)] \leftarrow \max\{R^{(t-1)}[h^*(e)], \rho^*(e)\}.$$

Let $q_R^{(t)}$ denote the probability of e changing the value of a register among $R[1], \dots, R[M]$ at time t . Formally, $q_R^{(t)}$ is defined as

$$\begin{aligned} q_R^{(t)} &= \sum_{j=1}^M P(h^*(e) = j \wedge R^{(t)}[j] > R^{(t-1)}[j]) \\ &= \frac{\sum_{j=1}^M 2^{-R^{(t-1)}[j]}}{M}. \end{aligned}$$

Algorithm 2: The pseudo-code for FreeRS.

```

 $R[1, \dots, M] \leftarrow [0, \dots, 0]; q_R \leftarrow 1;$ 
 $\hat{n}_s^{(\text{FreeRS})} \leftarrow 0, s \in S;$ 
foreach  $e = (s, d) \in \Gamma$  do
  if  $\rho^*(e) > R[h^*(e)]$  then
     $q_R \leftarrow q_R + \frac{2^{-\rho^*(e)} - 2^{-R[h^*(e)]}}{M};$ 
     $R[h^*(e)] \leftarrow \rho^*(e);$ 
     $\hat{n}_s^{(\text{FreeRS})} \leftarrow \hat{n}_s^{(\text{FreeRS})} + \frac{1}{q_R};$ 
  end
end

```

Let $\hat{n}_s^{(t, \text{FreeRS})}$ denote the cardinality estimate of user s at time t . When user-item pair $e = (s, d)$ arrives at time t , we update the cardinality estimate of user s as

$$\hat{n}_s^{(t, \text{FreeRS})} \leftarrow \hat{n}_s^{(t-1, \text{FreeRS})} + \frac{\mathbf{1}(R^{(t)}[h^*(e)] \neq R^{(t-1)}[h^*(e)])}{q_R^{(t)}}.$$

For any other user $s' \in S \setminus \{s\}$, we keep its cardinality estimate unchanged, i.e., $\hat{n}_{s'}^{(t, \text{FreeRS})} \leftarrow \hat{n}_{s'}^{(t-1, \text{FreeRS})}$. Similar to $q_B^{(t)}$, we compute $q_R^{(t)}$ incrementally. In detail, we initialize $q_R^{(1)} = 1$ and incrementally compute $q_R^{(t+1)}$ as

$$q_R^{(t+1)} \leftarrow q_R^{(t)} + \frac{2^{-\rho^*(e)} - 2^{-R[h^*(e)]}}{m} \mathbf{1}(R^{(t)}[h^*(e)] \neq R^{(t-1)}[h^*(e)]).$$

Hence, the time complexity of FreeRS for processing each user-item pair is also $O(1)$.

Error Analysis. We derive the error of $\hat{n}_s^{(t, \text{FreeRS})}$ as follows:

Theorem 2: The expectation and variance of $\hat{n}_s^{(t, \text{FreeRS})}$ are

$$\mathbb{E}(\hat{n}_s^{(t, \text{FreeRS})}) = n_s^{(t)},$$

$$\text{Var}(\hat{n}_s^{(t, \text{FreeRS})}) = \sum_{i \in T_s^{(t)}} \mathbb{E}\left(\frac{1}{q_R^{(i)}}\right) n_s^{(t)} \leq n_s^{(t)} \left(\mathbb{E}\left(\frac{1}{q_R^{(t)}}\right) - 1 \right),$$

where

$$\mathbb{E}\left(\frac{1}{q_R^{(i)}}\right) = \sum_{k_1, \dots, k_M \geq 0} \frac{\sum_{n_1 + \dots + n_M = n^{(i)}} \binom{n^{(i)}}{n_1, \dots, n_M} \prod_{j=1}^M \gamma_{n_j, k_j}}{M^{n^{(i)} - 1} \sum_{j=1}^M 2^{-k_j}},$$

with

$$\gamma_{n_j, k_j} = \begin{cases} (1 - 2^{-k_j})^{n_j} - (1 - 2^{-k_j+1})^{n_j}, & n_j > 0, k_j > 0 \\ 0, & n_j > 0, k_j = 0 \\ 1, & n_j = 0, k_j = 0. \end{cases}$$

$\mathbb{E}(\frac{1}{q_R^{(i)}})$ is approximately $\frac{1.386n^{(i)}}{M}$ when $n^{(i)} > 2.5M$.

Proof. Let ξ_e denote an indicator equal to 1 when processing a user-item pair e incurs a value change of any $R[1], \dots, R[M]$, and 0 otherwise. We have

$$\hat{n}_s^{(t, \text{FreeRS})} = \sum_{i \in T_s^{(t)}} \frac{\xi_{e^{(i)}}}{q_R^{(i)}}.$$

For each $\xi_{e(i)}$, we have

$$\mathbb{E}(\xi_{e(i)} | q_R^{(i)}) = q_R^{(i)}, \quad 1 \leq i \leq t,$$

$$\text{Var}(\xi_{e(i)} | q_R^{(i)}) = \mathbb{E}(\xi_{e(i)}^2 | q_R^{(i)}) - (\mathbb{E}(\xi_{e(i)} | q_R^{(i)}))^2 = q_R^{(i)} - (q_R^{(i)})^2.$$

Similar to FreeBS, define

$$Q_R^{(t)} = \{q_R^{(1)}, \dots, q_R^{(t)}\}.$$

Then, we have

$$\mathbb{E}(\hat{n}_s^{(t, \text{FreeRS})} | Q_R^{(t)}) = \mathbb{E}\left(\sum_{i \in T_s^{(t)}} \frac{\xi_{e(i)}}{q_R^{(i)}} \mid Q_R^{(t)}\right) = n_s^{(t)}.$$

Therefore, we obtain

$$\mathbb{E}(\hat{n}_s^{(t, \text{FreeRS})}) = \mathbb{E}(\mathbb{E}(\hat{n}_s^{(t, \text{FreeRS})} | Q_R^{(t)})) = \mathbb{E}(n_s^{(t)}) = n_s^{(t)}.$$

Given $Q_R^{(t)}$, all $\xi_{e(i)}$ ($i \in T_s^{(t)}$) are independent of each other. Similar to FreeBS, the variance of $\hat{n}_s^{(t, \text{FreeRS})}$ given $Q_R^{(t)}$ is

$$\text{Var}(\hat{n}_s^{(t, \text{FreeRS})} | Q_R^{(t)}) = \sum_{i \in T_s^{(t)}} \frac{1}{q_R^{(i)}} - n_s^{(t)}.$$

It is easily shown that $\text{Var}(\mathbb{E}(\hat{n}_s^{(t, \text{FreeRS})} | Q_R^{(t)})) = 0$. Using equation $\text{Var}(X) = \text{Var}(\mathbb{E}(X|Y)) + \mathbb{E}(\text{Var}(X|Y))$, we have

$$\begin{aligned} \text{Var}(\hat{n}_s^{(t, \text{FreeRS})}) &= \mathbb{E}(\text{Var}(\hat{n}_s^{(t, \text{FreeRS})} | Q_R^{(t)})) \\ &= \mathbb{E}\left(\sum_{i \in T_s^{(t)}} \frac{1}{q_R^{(i)}} - n_s^{(t)}\right) = \sum_{i \in T_s^{(t)}} \mathbb{E}\left(\frac{1}{q_R^{(i)}}\right) - n_s^{(t)}. \end{aligned}$$

In what follows we derive the formula for $\mathbb{E}(\frac{1}{q_R^{(i)}})$. Using $h^*(\cdot)$, FreeRS randomly splits stream Γ into M sub-streams Γ_j , $1 \leq j \leq M$. Each $R[j]$ tracks the maximum value of function $\rho^*(\cdot)$ for user-item pairs in sub-stream Γ_j . At time i , assume that n_j distinct user-item pairs have occurred in Γ_j . Then, $P(R^{(i)}[j] = k_j | n_j) = \gamma_{n_j, k_j}$. Therefore,

$$\begin{aligned} P(R^{(i)}[1] = k_1, \dots, R^{(i)}[M] = k_M | n^{(i)}) &= \\ \frac{\sum_{n_1 + \dots + n_M = n^{(i)}} \binom{n^{(i)}}{n_1, \dots, n_M} \prod_{j=1}^M \gamma_{n_j, k_j}}{M^{n^{(i)}}}. \end{aligned}$$

An exact expression for $\mathbb{E}(\frac{1}{q_R^{(i)}})$ is easily derived. However, it is too complex to analyze. Hence, we introduce a method to approximate $\mathbb{E}(\frac{1}{q_R^{(i)}})$. From [19], we have $\mathbb{E}(\frac{\alpha_M M}{q_R^{(i)}}) = \alpha_M M \mathbb{E}(\frac{1}{q_R^{(i)}}) \approx n^{(i)}$ for $n^{(i)} > 2.5M$. Therefore, $\mathbb{E}(\frac{1}{q_R^{(i)}}) \approx \frac{n^{(i)}}{\alpha_M M} \approx \frac{1.386n^{(i)}}{M}$. $\square$

C. Discussions

FreeBS vs CSE. FreeBS outperforms CSE in three aspects: (1) FreeBS can estimate cardinalities up to $\sum_{i=1}^M \frac{M}{i} \approx M \ln M$, which is larger than the maximum cardinality $m \ln m$ allowed by CSE; (2) FreeBS exhibits a smaller estimation error than CSE. From [39], we find that the expectation and the variance of estimation $\hat{n}_s^{(t, \text{CSE})}$ given by CSE are

$$\mathbb{E}(\hat{n}_s^{(t, \text{CSE})}) \approx n_s^{(t)} + \frac{1}{2} \left(\mathbb{E}\left(\frac{1}{q^{(t)}}\right) e^{\frac{n_s^{(t)}}{m}} - \frac{n_s^{(t)}}{m} - 1 \right),$$

$$\text{Var}(\hat{n}_s^{(t, \text{CSE})}) \approx m \left(\mathbb{E}\left(\frac{1}{q^{(t)}}\right) e^{\frac{n_s^{(t)}}{m}} - \frac{n_s^{(t)}}{m} - 1 \right),$$

where $q^{(t)} = \frac{U^{(t)}}{M}$ is the fraction of zero bits in the shared bit array at time t . We find that CSE exhibits a large bias when $n_s^{(t)} \gg m$, while FreeBS is unbiased. When m is large (e.g., m approaches to M), FreeBS and CSE perform nearly the same bit setting operations but use their different cardinality estimators. Then, we have $\mathbb{E}(\frac{1}{q^{(t)}}) \approx \mathbb{E}(\frac{1}{q_B^{(t)}})$. From Theorem 2, therefore, we easily have

$$\begin{aligned} \text{Var}(\hat{n}_s^{(t, \text{CSE})}) &\gtrsim m \left(\mathbb{E}\left(\frac{1}{q^{(t)}}\right) (1 + \frac{n_s^{(t)}}{m}) - \frac{n_s^{(t)}}{m} - 1 \right) \\ &> n_s^{(t)} \mathbb{E}\left(\frac{1}{q^{(t)}}\right) - n_s^{(t)} \gtrsim \text{Var}(\hat{n}_s^{(t, \text{FreeBS})}); \end{aligned}$$

3) FreeBS has complexity $O(1)$ to update user cardinality estimates each time it observes a new user-item pair, while CSE has complexity $O(m)$.

FreeRS vs vHLL. FreeRS and vHLL both can estimate cardinalities up to 2^{2^w} , where w is the number of bits in a register. However, FreeRS outperforms vHLL in two aspects: 1) FreeRS has complexity $O(1)$ to update user cardinality estimates each time it observes a new user-item pair, while vHLL has complexity $O(m)$; 2) FreeRS exhibits a smaller estimation error in comparison with vHLL. From Theorem 2, then we have $\text{Var}(\hat{n}_s^{(t, \text{FreeRS})}) \leq n_s^{(t)} (\mathbb{E}(\frac{1}{q_R^{(t)}}) - 1) \approx n_s^{(t)} (\frac{n^{(t)}}{M \alpha_M} - 1) < \frac{1.386n^{(t)} n_s^{(t)}}{M}$, while the variance of vHLL is

$$\begin{aligned} \text{Var}(\hat{n}_s^{(t, \text{vHLL})}) &\gtrsim \left(\frac{M}{M-m}\right)^2 \times \frac{1.04^2}{m} \times 2n^{(t)} n_s^{(t)} \frac{m}{M} (1 - \frac{m}{M}) \\ &= \frac{2.163n^{(t)} n_s^{(t)}}{M-m} > \frac{2.163n^{(t)} n_s^{(t)}}{M}. \end{aligned}$$

FreeBS vs FreeRS. We observe that 1) FreeBS is faster than FreeRS. For each coming user-item pair e , FreeBS only computes $h^*(e)$ to select and set a bit, but FreeRS needs to compute both $h^*(e)$ and $\rho^*(e)$ to select and update a register; 2) Under the same memory usage, we compare the accuracy of FreeBS using M bits and FreeRS using M/w registers, where w is the number of bits in a register. From Theorems 1 and 2, we have

$$\begin{aligned} \text{Var}(\hat{n}_s^{(t, \text{FreeBS})}) &= \sum_{i \in T_s^{(t)}} \mathbb{E}\left(\frac{1}{q_B^{(i)}}\right) - n_s^{(t)}, \\ \text{Var}(\hat{n}_s^{(t, \text{FreeRS})}) &= \sum_{i \in T_s^{(t)}} \mathbb{E}\left(\frac{1}{q_R^{(i)}}\right) - n_s^{(t)}, \end{aligned}$$

where $\mathbb{E}(\frac{1}{q_B^{(i)}}) \approx e^{\frac{n^{(i)}}{M}}$ and $\mathbb{E}(\frac{1}{q_R^{(i)}}) \approx \frac{1.386wn^{(i)}}{M} < e^{\frac{n^{(i)}}{M}}$ when $\frac{n^{(i)}}{M} \geq 0.772w$. Therefore, FreeRS is more accurate than FreeBS for users not appearing among the first $0.772wM$ distinct user-item pairs presented on stream Γ . Flajolet et al. [19] observe that HLL exhibits a large error for estimating small cardinalities. To solve this problem, they view a register of HLL as a bit of LPC and estimate the cardinality based on

the fraction of registers that equal 0. When $n^{(i)} \ll M/w$, we easily find that $q_R^{(i)}$ is approximately computed as the fraction of registers that equal 0 at time i , and we have $\mathbb{E}(\frac{1}{q_R^{(i)}}) < \mathbb{E}(\frac{1}{q_R^{(j)}})$ because the number of bits in FreeBS is w times larger than the number of registers in FreeRS under the same memory usage. It indicates that FreeBS is more accurate than FreeRS for users whose user-item pairs appear early in stream Γ .

V. EVALUATION

A. Datasets

In our experiments, we used a variety of publicly available real-world datasets to evaluate the performance of our methods in comparison with state-of-the-art methods, which are summarized in Table I. Dataset sanjose (resp. chicago) consists of one-hour passive traffic traces collected from the equinix-sanjose (resp. equinix-chicago) data collection monitors during March 20, 2014. Twitter, Flickr, Orkut and LiveJournal are all graph-datasets where each edge represents social relationship between any two users and may occur more than once in these datasets. Figure 2 shows the CCDFs of user cardinalities for all datasets used in our experiments.

Table I
SUMMARY OF DATASETS USED IN OUR EXPERIMENTS.

dataset	#users	max-cardinality	total cardinality
sanjose [1]	8,387,347	313,772	23,073,907
chicago [1]	1,966,677	106,026	9,910,287
Twitter [31]	40,103,281	2,997,496	1,468,365,182
Flickr [35]	1,441,431	26,185	22,613,980
Orkut [35]	2,997,376	31,949	223,534,301
LiveJournal [35]	4,590,650	9,186	76,937,805

B. Baselines

FreeBS and CSE [50] are bit-sharing algorithms, while FreeRS and vHLL [47] are register-sharing algorithms, where each register consists of 5 bits, i.e., $w = 5$. To compare our methods with these two methods under the same memory size, we let FreeBS and CSE have M bits, whereas FreeRS and vHLL have $M/5$ 5-bit registers. Moreover, both CSE and vHLL use virtual sketches to record the cardinality of each user, and we set m (i.e., the number of bits/registers in the virtual sketch) to be the same for both methods. LPC [46] and HyperLogLog++ [23] (short for HLL++) build a sketch for each user $s \in S$ to record items that s connects to. Specially, HLL++ is an optimized method of HyperLogLog, which uses 6 bits for each register, and implements bias correction and sparse representation strategies to improve cardinality estimation performance. In our experiments, under the same memory size M , we let LPC have $\frac{M}{|S|}$ bits and HLL++ have $\frac{M}{6|S|}$ 6-bit registers for each user respectively.

In this paper, we aim to compute the cardinalities of all users at each time $t = 1, 2, \dots$. At each time t , enumerating each occurred user and computing its cardinality requires time $O(|S^{(t)}|m)$ for methods CSE, vHLL, LPC, and HLL++, where $S^{(t)}$ is the set of occurred users before and include time

Figure 2. CCDFs of user cardinalities.

t . It is computationally intensive and thus is prohibitive for estimating all users' cardinalities over time t . To solve this problem, we allocate each occurred user s a counter $\hat{n}_s$ to keep tracking of u 's cardinality for CSE, vHLL, LPC, and HLL++. For each edge $e^{(t)} = (s^{(t)}, d^{(t)})$ arriving at time t , we only estimate the cardinality of user $s^{(t)}$ for CSE, vHLL, LPC, and HLL++, and then update its counter $\hat{n}_{s^{(t)}}$ and keep the counters of the other occurred users unchanged, which reduces time complexity from $O(|S^{(t)}|m)$ to $O(m)$. To track all users' cardinalities over time, therefore, all methods FreeRS, FreeBS, CSE, vHLL, LPC, and HLL++ require a counter for each user, and we do not consider this memory usage in our comparisons.

C. Metrics

In our experiments, we use a fine-grained metric *relative standard error* (RSE) to evaluate the performance of estimating the cardinality for any user with a particular cardinality n at time t . Smaller RSE indicates better performance of cardinality estimation. Formally, we define

$$\text{RSE}^{(t)}(n) = \frac{1}{n} \sqrt{\frac{\sum_{s \in S} (\hat{n}_s^{(t)} - n)^2 \mathbf{1}(n_s^{(t)} = n)}{\sum_{s \in S} \mathbf{1}(n_s^{(t)} = n)}}.$$

D. Runtime

For CSE, vHLL, LPC and HLL++, its runtime is mainly determined by the number of bits/registers in each user's (virtual) sketch, because estimating a user's cardinality needs to enumerate m bits/registers in the user's sketch, which requires large time complexity $O(m)$. On the contrary, our methods FreeBS and FreeRS have low time complexity $O(1)$ to process each coming element and update the counter for its user. In our experiments, we vary the number of bits/registers in the sketch m to compare the runtime of all six methods. In this case, we record the runtime required for processing each element and updating the cardinality of the user, and then average the update time for all users in the datasets. The experimental results are shown in Figure 3. As m increases, the runtime of all four methods CSE, vHLL, LPC and HLL++ increases. Our methods FreeBS and FreeRS are substantially faster than other

Figure 3. Runtime of our methods FreeBS and FreeRS in comparison with CSE, vHLL, LPC and HLL++ for different m (i.e., the number of bits/registers in the (virtual) sketch) of a user.

four methods for different m . Also we notice that CSE is faster than vHLL and FreeBS is faster than FreeRS, this is because register sharing based methods perform more operations than bitmap sharing methods for processing each element.

E. Accuracy

Whenever an element $e = (s, d)$ arrives, we implemented all six methods to estimate the cardinality of the user s . Figure 4 shows the experimental results when all elements in the dataset Orkut arrive. In our experiments, we fixed the memory size $M = 5 \times 10^8$ bits, and the number of bits/register in the virtual sketch for CSE/vHLL was set to $m = 1,024$. Under the fixed memory size, each user in the dataset Orkut used 167 bits for LPC and 28 6-bit registers for HLL++ respectively. Points close to the diagonal line $\hat{n}_s^{(t)} = n_s^{(t)}$ indicate better cardinality estimation accuracy. From Figure 4, we observe that our methods FreeBS and FreeRS are more accurate than other four methods for different actual cardinalities. CSE and LPC have limited estimation ranges and LPC exhibits extremely bad estimation performance especially for larger cardinalities, therefore we omit the experimental results of LPC in the later experiments. Furthermore, we used the metric RSE to compare the estimation errors of our methods with those of CSE, vHLL and HLL++ at a fine-grained level. Figure 5 shows the RSEs of all methods for all datasets in Table I. We can see that our methods FreeBS and FreeRS are up to 10,000 times more accurate than CSE, vHLL, and HLL++. Bit sharing method FreeBS (resp. CSE) is more accurate than register sharing method FreeRS (resp. vHLL) for characterizing users with small cardinalities. For users with large cardinalities, register sharing method FreeRS (resp. vHLL) outperforms bit sharing method FreeBS (resp. CSE). Specially, the RSE of CSE first decreases and then increases as the actual cardinality increases. This is because CSE has a small estimation range, i.e., $m \ln m$, which is consistent with the original paper [50]. Meanwhile, HLL++ is more accurate than CSE and vHLL for small cardinalities due to its bias correction strategy. Because each user uses fewer registers for HLL++ than vHLL, HLL++ performs larger estimation errors than vHLL for large cardinalities.

Figure 4. (Orkut) Estimated cardinalities vs actual cardinalities for FreeBS, FreeRS, CSE, vHLL, LPC, and HLL++ under the same memory size $M = 5 \times 10^8$ (bits), where the number of bits/registers in the virtual sketch for CSE/vHLL $m = 1,024$.

F. Case Study: Detecting Super Spreaders over Time

We implemented methods FreeBS, FreeRS, CSE, vHLL, and HLL++ to detect super spreaders over time, where a super spreader refers to a user connecting to at least $\Delta n^{(t)}$ items at time t , where $n^{(t)}$ is the sum of all user cardinalities at time t and $0 < \Delta < 1$ is a relative threshold. In our experiments, we set memory size $M = 5 \times 10^8$ bits and the number of bits/registers in the virtual sketch $m = 1,024$. We used two metrics *false negative ratio* (FNR) and *false positive ratio* (FPR) to evaluate performance, where FNR is the ratio of the number of super spreaders not detected to the number of super spreaders, and FPR is the ratio of the number of users that are wrongly detected as super spreaders to the number of all users. Figure 6 shows experimental results of all five methods for detecting super spreaders over time in dataset sanjose. We observe that both FreeBS and FreeRS are more accurate than other three methods for detecting super spreaders over time. For example, FNR and FPR for FreeBS and FreeRS are about 4 to 20 times smaller than other three methods. Table II shows the results for all datasets when all elements arrive. We notice

Figure 5. (All datasets) Cardinality estimation accuracy, where memory size $M = 5 \times 10^8$ (bits), and $m = 1,024$ for CSE/vHLL.

Figure 6. (sanjose) Accuracy of detecting super spreaders over time t , where $\Delta = 5 \times 10^{-5}$, $M = 5 \times 10^8$ (bits), and $m = 1,024$ for CSE/vHLL.

Table II

(ALL DATASETS) PERFORMANCE OF DETECTING SUPER SPREADERS WITH $\Delta = 5 \times 10^{-5}$, $M = 5 \times 10^8$, $m = 1,024$. CSE REPORTED AN EMPTY SET OF USERS FOR TWITTER AND ORKUT DUE TO THE LIMITED ESTIMATION RANGE, THEREFORE WE REPORT THEIR RESULTS ON TWITTER AND ORKUT AS “N/A”.

dataset	FNR					FPR				
	FreeBS	FreeRS	CSE	vHLL	HLL++	FreeBS	FreeRS	CSE	vHLL	HLL++
sanjose	2.54e-3	2.27e-3	1.02e-2	1.12e-2	2.88e-2	1.61e-7	1.76e-7	1.22e-6	1.08e-6	3.38e-6
chicago	2.61e-3	2.55e-3	8.00e-3	8.23e-3	9.06e-3	2.32e-6	2.54e-6	8.46e-6	8.47e-6	1.08e-5
Twitter	2.45e-2	2.27e-2	N/A	6.38e-2	1.04e-1	2.99e-7	2.75e-7	N/A	6.84e-7	1.25e-6
Flickr	5.13e-3	5.55e-3	1.09e-2	1.13e-2	1.38e-2	1.18e-5	1.04e-5	3.57e-5	3.87e-5	4.13e-5
Orkut	1.47e-2	1.53e-2	N/A	7.86e-2	9.71e-2	3.22e-7	3.34e-7	N/A	1.97e-6	2.52e-6
LiveJournal	4.37e-3	4.62e-3	1.04e-2	9.92e-3	1.36e-2	4.33e-7	4.15e-7	1.22e-6	1.13e-6	1.78e-6

that our methods FreeBS and FreeRS outperform other three methods on all datasets.

VI. RELATED WORK

Sketch methods use small amount of memory to quickly build compressive summaries of large data streams, and have been successfully used for applications such as heavy hitter detection [15], [16], [11], [12], [52], [51], heavy change detection [27], [9], [38], super host detection [54], [4], [50],

[42], [43], [44], cardinality distribution estimation [6], [45], network flow size estimation [13], [24], [49], [28], [37], [29], [30], [33], and network traffic entropy estimation [32], [53]. Next, we discuss existing cardinality estimation methods in detail.

Estimating data streams’ cardinalities. To estimate the cardinality of a large data stream (i.e., the number of distinct elements in the data stream), Whang et al. [46] develop the first sketch method LPC. LPC can only estimate cardinal-

ities less than $m \ln m$, where m is the size of the LPC sketch. Therefore, it needs to set a large m to handle data streams with large cardinalities. [17], [5] combine LPC and different sampling methods to enlarge the estimation range. Flajolet and Martin [20] develop a sketch method FM, which uses a register to estimate the data stream's cardinality and provides a cardinality estimation bounded by 2^w , where w is the number of bits in the register. To further improve the accuracy and decrease the memory usage, sketch methods MinCount [3], LogLog [14], HyperLogLog [19], HyperLogLog++ [23], RoughEstimator [26], and HLL-TailCut+ [48] are developed to use a list of m registers and compress the size of each register from 32 bits to 5 or 4 bits under the same estimation range of 2^{32} . [22], [34] develop several cardinality estimators based on order statistics of observed samples. Ting [40] introduces the concept of an area cutting process to generally model the above sketch methods and provides a martingale based estimator to further improve the accuracy of these methods. Chen et.al [7] extend HyperLogLog to estimate the cardinality over sliding windows. Besides these sketch methods, sampling methods Wegman's adaptive sampling [18] and distinct sampling [21] are also developed for estimating a large stream's cardinality, while Flajolet et al. [19] reveal that the memory efficiency of these two sampling methods is even worse than the original FM method. Recently, [8], [41] developed methods using the above sketch methods to estimate the cardinalities of set unions and intersections.

Estimating all user cardinalities. Significant attention has been paid to develop sketch methods to estimate the cardinalities of network hosts (or users) over high speed links. To achieve a desired accuracy, the above methods require a large number of bits/registers for each host because host cardinalities are unknown in advance and vary over a large range. This is not memory efficient because most hosts have small cardinalities. To solve this problem, [54], [50], [43], [47] develop different virtual sketch methods to compress the LPC/HLL sketches of all hosts into a large bit/register array shared by all hosts. Zhao et al. [54] propose a virtual sketch method, which consists of a list of LPC sketches (i.e., a two-dimensional bit array). For each host, they randomly select k (k is usually set to 2 or 3) LPC sketches from the list shared by all hosts. Note that each LPC in the list may be used by more than one hosts, which introduces "noisy" bits in a host's LPCs. To remove the error introduced by "noisy" bits, Zhao et al. [54] develop a method to estimate a host's cardinality based on its k LPC sketches. To further reduce the memory usage, [50], [43] generate each host's virtual LPC sketch by selecting m bits from a large one-dimensional bit array at random. All these virtual sketch methods [50], [43] have to set a large value of m (e.g., thousand) to achieve reasonable accuracy for host cardinalities over a large range. It results in large estimation errors, because most hosts have small cardinalities and many bits in their virtual LPC sketches tend to be contaminated by "noisy" bits. To reduce "noisy" bits in virtual LPC sketches, [39] builds a small regular LPC sketch for each host, which also provides information for estimating the host cardinality.

As we mentioned, these LPC based virtual sketch methods have small estimation ranges bounded by $m \ln m$. To enlarge the estimation range, [47] develop a sketch method vHLL, which generates a virtual HLL sketch by randomly selecting m registers from a large register array shared by all hosts. To achieve desired accuracy for host cardinalities over a large range, [47] also needs to set a large value of m . However, it results in a high computational cost and large estimation errors for network hosts with small cardinalities, in which virtual HLL sketches include many "noisy" registers. In addition, the above virtual sketch methods are customized to non-streaming settings, i.e., estimate host cardinalities at the end of an interval, and are computationally expensive to be extended to streaming settings.

VII. CONCLUSIONS AND FUTURE WORK

In this paper, we develop two novel streaming algorithms FreeBS and FreeRS to accurately estimate user cardinalities over time. Compared to existing bit/register sharing methods using (i.e., selecting) only m bits/registers for each user, FreeBS/FreeRS enables that the number of bits/registers used by a user dynamically increases as its cardinality increases over time and each user can use all shared bits/registers. It is therefore capable to estimate user cardinalities over a large range. For example, existing bit sharing methods can be only used to estimate user cardinalities over the range $[0, m \ln m]$. Our method FreeBS enlarges the range to $[0, M \ln M]$, where M is the total number of bits/registers used by all users. In addition, our algorithms FreeBS and FreeRS exploit dynamic properties of shared bits/registers to significantly improve estimation accuracy. They are simple yet effective, and sharply reduce time complexity of computing all user cardinalities to $O(1)$ each time they observe a new user-item pair. We conduct experiments on real-world datasets, and experimental results demonstrate that our methods FreeBS and FreeRS significantly outperform state-of-the-art methods in terms of accuracy and computational time. In future, we plan to extend our methods to applications such as SDN routers to monitor anomalies.

ACKNOWLEDGMENT

The research presented in this paper is supported in part by National Key R&D Program of China (2018YFC0830500), National Natural Science Foundation of China (U1301254, 61603290, 61602371), the Ministry of Education&China Mobile Research Fund (MCM20160311), the Natural Science Foundation of Jiangsu Province (SBK2014021758), 111 International Collaboration Program of China, the Prospective Joint Research of Industry-Academia-Research Joint Innovation Funding of Jiangsu Province (BY2014074), Shenzhen Basic Research Grant (JCYJ20160229195940462, JCYJ20170816100819428), China Postdoctoral Science Foundation (2015M582663), Natural Science Basic Research Plan in Shaanxi Province of China (2016JQ6034).

REFERENCES

- [1] Internet traces 2014, center for applied internet data analysis (caida). http://www.caida.org/data/passive/passive_2014_dataset.xml.

- [2] M. Abramowitz and I. A. Stegun. *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables*. Dover, 1964.
- [3] Z. Bar-Yossef, T. S. Jayram, R. Kumar, D. Sivakumar, and L. Trevisan. Counting distinct elements in a data stream. In *RANDOM*, 2002.
- [4] J. Cao, Y. Jin, A. Chen, T. Bu, and Z. Zhang. Identifying high cardinality Internet hosts. In *Proceedings of IEEE INFOCOM*, pages 810–818, April 2009.
- [5] A. Chen, J. Cao, L. Shepp, and T. Nguyen. Distinct counting with a self-learning bitmap. *Journal of the American Statistical Association*, 106(495):879–890, 2011.
- [6] A. Chen, L. Li, and J. Cao. Tracking cardinality distributions in network traffic. In *Proceedings of IEEE INFOCOM*, pages 819–827, April 2009.
- [7] W. Chen, Y. Liu, and Y. Guan. Cardinality change-based early detection of large-scale cyber-attacks. In *Proceedings of IEEE INFOCOM*, pages 1788–1796. IEEE, 2013.
- [8] R. Cohen, L. Katzir, and A. Yehezkel. A minimal variance estimator for the cardinality of big data set intersection. In *SIGKDD*, 2017.
- [9] G. Cormode and S. Muthukrishnan. What's hot and what's not: Tracking most frequent items dynamically. In *Proceedings of ACM PODC 2003*, pages 296–306, July 2003.
- [10] G. Cormode and S. Muthukrishnan. An improved data stream summary: the count-min sketch and its applications. *Journal of Algorithms*, 55(1):58–75, 2005.
- [11] G. Cormode, S. Muthukrishnan, and D. Srivastava. Finding hierarchical heavy hitters in data streams. In *Proceedings of VLDB*, pages 464–475, September 2003.
- [12] G. Cormode, S. Muthukrishnan, and D. Srivastava. Diamond in the rough: Finding hierarchical heavy hitters in multi-dimensional data. In *Proceedings of ACM SIGMOD*, pages 155–166, June 2004.
- [13] N. Duffield, C. Lund, and M. Thorup. Estimating flow distributions from sampled flow statistics. In *SIGCOMM*, 2003.
- [14] M. Durand and P. Flajolet. *Loglog Counting of Large Cardinalities*, pages 605–617. Springer Berlin Heidelberg, 2003.
- [15] C. Estan and G. Varghese. New directions in traffic measurement and accounting. In *Proceedings of ACM SIGCOMM*, pages 323–336, October 2002.
- [16] C. Estan and G. Varghese. Automatically inferring patterns of resource consumption in network traffic. In *Proceedings of ACM SIGCOMM*, pages 137–148, August 2003.
- [17] C. Estan, G. Varghese, and M. Fisk. Bitmap algorithms for counting active flows on high speed links. In *Proceedings of ACM SIGCOMM IMC*, pages 182–209, October 2003.
- [18] P. Flajolet. On adaptive sampling. *Computing*, 43(4):391–400, Feb. 1990.
- [19] P. Flajolet, E. Fusy, O. Gandouet, and F. Meunier. Hyperloglog: The analysis of a near-optimal cardinality estimation algorithm. In *AOFA*, 2007.
- [20] P. Flajolet and G. N. Martin. Probabilistic counting algorithms for data base applications. *Journal of Computer and System Sciences*, 31(2):182–209, October 1985.
- [21] P. B. Gibbons. Distinct sampling for highly-accurate answers to distinct values queries and event reports. In *PVLDB*, pages 541–550, San Francisco, CA, USA, 2001.
- [22] F. Giroire. Order statistics and estimating cardinalities of massive data sets. *Discrete Applied Mathematics*, 157(2):406 – 427, 2009.
- [23] S. Heule, M. Nunkesser, and D. P. Woodruff. Hyperloglog in practice: Algorithmic engineering of a state of the art cardinality estimation algorithm. In *EDBT/ICDT*, pages 683–692, 2013.
- [24] N. Hohn and D. Veitch. Inverting sampled traffic. In *Proceedings of ACM SIGCOMM IMC*, pages 222–233, October 2003.
- [25] D. Horvitz and D. Thompson. A generalization of sampling without replacement from a finite universe. *Journal of the American Statistical Association*, 47(260):663–685, 1952.
- [26] D. M. Kane, J. Nelson, and D. P. Woodruff. An optimal algorithm for the distinct elements problem. In *PODS*, pages 41–52, 2010.
- [27] B. Krishnamurthy, S. Sen, Y. Zhang, and Y. Chen. Sketch-based change detection: Methods, evaluation, and applications. In *Proceedings of ACM SIGCOMM IMC*, pages 234–247, October 2003.
- [28] A. Kumar, M. Sung, J. Xu, and J. Wang. Data streaming algorithms for efficient and accurate estimation of flow size distribution. In *Proceedings of ACM SIGMETRICS*, pages 177–188, June 2004.
- [29] A. Kumar, M. Sung, J. Xu, and E. W. Zegura. A data streaming algorithm for estimating subpopulation flow size distribution. In *Proceedings of ACM SIGMETRICS*, pages 61–72, June 2005.
- [30] A. Kumar and J. Xu. Sketch guided sampling—using onling estimates of flow size for adaptive data collection. In *Proceedings of IEEE INFOCOM*, pages 1–11, April 2006.
- [31] H. Kwak, C. Lee, H. Park, and S. Moon. What is twitter, a social network or a news media? In *WWW*, 2010.
- [32] A. Lall, V. Sekar, M. Ogihara, J. Xu, and H. Zhang. Data streaming algorithms for estimating entropy of network traffic. In *Proceedings of ACM SIGMETRICS*, pages 145–156, June 2006.
- [33] P. Lieven and B. Scheuermann. High-speed per-flow traffic measurement with probabilistic multiplicity counting. In *INFOCOM*, pages 1253–1261, 2010.
- [34] J. Lumbroso. An optimal cardinality estimation algorithm based on order statistics and its full analysis. In *AofA*, pages 489–504, 2010.
- [35] A. Mislove, M. Marcon, K. P. Gummadi, P. Druschel, and B. Bhattacharjee. Measurement and analysis of online social networks. In *IMC*, 2007.
- [36] G. Nychis, V. Sekar, D. G. Andersen, H. Kim, and H. Zhang. An empirical evaluation of entropy-based traffic anomaly detection. In *Proceedings of ACM SIGCOMM IMC*, pages 151–156, October 2008.
- [37] B. Ribeiro, T. Ye, and D. Towsley. A resource-minimalist flow size histogram estimator. In *Proceedings of ACM SIGCOMM IMC*, pages 285–290, October 2008.
- [38] R. Schwellen, Z. Li, Y. Chen, Y. Gao, A. Gupta, E. Parsons, Y. Zhang, P. Dinda, M. yang Kao, and G. Memik. Reversible sketches: Enabling monitoring and analysis over high-speed data streams. *IEEE/ACM Transactions on Networking*, 15(5):1059–1072, October 2007.
- [39] J. Tao, P. Wang, X. Guan, and W. Hu. Al-bitmap: Monitoring network traffic activity graphs on high speed links. *Inf. Sci.*, 408:162–175, 2017.
- [40] D. Ting. Streamed approximate counting of distinct elements: Beating optimal batch methods. In *SIGKDD*, pages 442–451, 2014.
- [41] D. Ting. Towards optimal cardinality estimation of unions and intersections with sketches. In *SIGKDD*, 2016.
- [42] P. Wang, X. Guan, W. Gong, and D. Towsley. A new virtual indexing method for measuring host connection degrees. In *Proceedings of IEEE INFOCOM Mini-Conference*, pages 156–160, April 2011.
- [43] P. Wang, X. Guan, T. Qin, and Q. Huang. A data streaming method for monitoring host connection degrees of high-speed links. *IEEE Transactions on Information Forensics and Security*, 6(3-2):1086–1098, 2011.
- [44] P. Wang, X. Guan, D. Towsley, and J. Tao. Virtual indexing based methods for estimating node connection degrees. *Computer Networks*, 56(12):2773–2787, 2012.
- [45] P. Wang, X. Guan, J. Zhao, J. Tao, and T. Qin. A new sketch method for measuring host connection degree distribution. *TIFS*, 9(6):948–960, 2014.
- [46] K. Whang, B. T. Vander-zanden, and H. M. Taylor. A linear-time probabilistic counting algorithm for database applications. *IEEE Transaction of Database Systems*, 15(2):208–229, June 1990.
- [47] Q. Xiao, S. Chen, M. Chen, and Y. Ling. Hyper-compact virtual estimators for big network data based on register sharing. In *SIGMETRICS*, pages 417–428, 2015.
- [48] Q. Xiao, Y. Zhou, and S. Chen. Better with fewer bits: Improving the performance of cardinality estimation of large data streams. In *INFOCOM*, pages 1–9, 2017.
- [49] L. Yang and G. Michailidis. Sample based estimation of network traffic flow characteristics. In *Proceedings of IEEE INFOCOM*, pages 1775–1783, May 2007.
- [50] M. Yoon, T. Li, S. Chen, and J. K. Peir. Fit a spread estimator in small memory. In *Proceedings of IEEE INFOCOM*, pages 504–512, April 2009.
- [51] M. Yu, L. Jose, and R. Miao. Software defined traffic measurement with opensketch. In *Proceedings of USENIX NSDI*, pages 29–42, Berkeley, CA, USA, 2013.
- [52] Y. Zhang, S. Singh, S. Sen, N. G. Duffield, and C. Lund. Online identification of hierarchical heavy hitters: Algorithms, evaluation, and application. In *Proceedings of ACM SIGCOMM IMC*, pages 101–114, October 2004.
- [53] H. Zhao, A. Lall, O. Spatscheck, J. Wang, and J. Xu. A data streaming algorithm for estimating entropies of OD flows. In *Proceedings of ACM SIGCOMM IMC*, pages 279–290, October 2007.
- [54] Q. Zhao, A. Kumar, and J. Xu. Joint data streaming and sampling techniques for detection of super sources and destinations. In *Proceedings of ACM SIGCOMM IMC*, pages 77–90, October 2005.