

Optimal Sampling from Sliding Windows

Vladimir Braverman*
University of California Los Angeles
vova@cs.ucla.edu

Rafail Ostrovsky†
University of California Los Angeles
rafail@cs.ucla.edu

Carlo Zaniolo
University of California Los Angeles
zaniolo@cs.ucla.edu

ABSTRACT

A *sliding windows* model is an important case of the streaming model, where only the most “recent” elements remain active and the rest are discarded in a stream. The sliding windows model is important for many applications (see, e.g., Babcock, Babu, Datar, Motwani and Widom (PODS 02); and Datar, Gionis, Indyk and Motwani (SODA 02)). There are two equally important types of the sliding windows model – windows with fixed size, (e.g., where items arrive one at a time, and only the most recent n items remain active for some fixed parameter n), and bursty windows (e.g., where many items can arrive in “bursts” at a single step and where only items from the last t steps remain active, again for some fixed parameter t).

Random sampling is a fundamental tool for data streams, as numerous algorithms operate on the sampled data instead of on the entire stream. Effective sampling from sliding windows is a nontrivial problem, as elements eventually expire. In fact, the deletions are *implicit*; i.e., it is not possible to identify deleted elements without storing the entire window. The implicit nature of deletions on sliding windows does not allow the existing methods (even those that support explicit deletions, e.g., Cormode, Muthukrishnan and Rozenbaum (VLDB 05); Frahling, Indyk and Sohler (SOCG 05)) to be directly “translated” to the sliding windows model. One trivial approach to overcoming the problem of implicit deletions is that of over-sampling. When k samples are required, the over-sampling method maintains $k' > k$ samples in the hope that at least k samples are not expired. The obvious disadvantages of this method are twofold: (a) It introduces additional costs and thus decreases the performance; and (b) The memory bounds are not deterministic, which is typical for

*Supported in part by NSF grant 0830803

†Department of Computer Science and Department of Mathematics, UCLA, Los Angeles, CA 90095, USA. Supported in part by IBM Faculty Award, Xerox Innovation Group Award, NSF grants 0430254, 0716835, 0716389, 0830803 and U.C. MICRO grant.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

ACM PODS '09 Providence, RI, USA

Copyright 200X ACM X-XXXXX-XX-X/XX/XX ...\$5.00.

streaming algorithms (where even small probability events may eventually happen for a stream that is big enough).

Babcock, Datar and Motwani (SODA 02), were the first to stress the importance of improvements to over-sampling. They formally introduced the problem of sampling from sliding windows and improved the over-sampling method for *sampling with replacement*. Their elegant solutions for sampling with replacement are optimal *in expectation*, and thus resolve disadvantage (a) mentioned above. Unfortunately, the randomized bounds do not resolve disadvantage (b) above. Interestingly, all algorithms that employ the ideas of Babcock, Datar and Motwani have the same central problem of having to deal with randomized complexity (see e.g., Datar and Muthukrishnan (ESA 02); Chakrabarti, Cormode and McGregor (SODA 07)). Further, the proposed solutions of Babcock, Datar and Motwani for *sampling without replacement* are based on the criticized over-sampling method and thus do not solve problem (a). Therefore, the question of whether we can solve sampling on sliding windows optimally (i.e., resolving both disadvantages) is implicit in the paper of Babcock, Datar and Motwani and has remained open for all variants of the problem.

In this paper we answer these questions affirmatively and provide optimal sampling schemas for all variants of the problem, i.e., sampling with or without replacement from fixed or bursty windows. Specifically, for fixed-size windows, we provide optimal solutions that require $O(k)$ memory; for bursty windows, we show algorithms that require $O(k \log n)$, which is optimal since it matches the lower bound by Gemulla and Lehner (SIGMOD 08). In contrast to the work of Babcock, Datar and Motwani, our solutions have deterministic bounds. Thus, we prove a perhaps somewhat surprising fact: the memory complexity of the sampling-based algorithm for all variants of the sliding windows model is comparable with that of streaming models (i.e., without the sliding windows). This is the first result of this type, since all previous “translations” of sampling-based algorithms to sliding windows incur randomized memory guarantees only.

Categories and Subject Descriptors. F.2 [ANALYSIS OF ALGORITHMS AND PROBLEM COMPLEXITY.]: Miscellaneous.

Keywords. Data Streams, Sliding Windows, Random Sampling.

1. Introduction

Random sampling and sliding windows are two fundamental concepts for data streams. Sampling is a very natural way to summarize data properties with sublinear space; indeed, it is a key component of many streaming algorithms and techniques. Just to mention a few, the relevant papers include Aggarwal [2]; Alon, Duffield, Lund and Thorup [3]; Alon, Matias and Szegedy [4]; Babcock, Babu, Datar, Motwani and Widom [8]; Babcock, Datar and Motwani [10]; Bar-Yossef [13]; Bar-Yossef, Kumar and Sivakumar [17]; Burriol, Frahling, Leonardi, Marchetti-Spaccamela and Sohler [20]; Chakrabarti, Cormode and McGregor [21]; Chaudhuri and Mishra [26]; Chaudhuri, Motwani and Narasayya [27]; Cohen [29]; Cohen and Kaplan [30]; Cormode, Muthukrishnan and Rozenbaum [32]; Dasgupta, Drineas, Harb, Kumar and Mahoney [35]; Datar and Muthukrishnan [37]; Duffield, Lund and Thorup [38]; Frahling, Indyk and Sohler [43]; Gandhi, Suri and Welzl [46]; Gemulla [47]; Gemulla and Lehner [48]; Gibbons and Matias [49]; Guha, Meyerson, Mishra, Motwani and O’Callaghan [54]; Haas [55]; Kolonko and Wäsch [58]; Li [62]; Palmer and Faloutsos [67]; Szegedy [70]; and Vitter [72]; These papers illustrate the vitality of effective sampling methods for data streams. Among other methods, *uniform random sampling* is the most general and well-understood. Most applications maintain multiple samples using two popular methods: namely, *sampling with replacement* and *sampling without replacement*. The former method assumes independence among samples; the latter forbids repetitions. While sampling without replacement preserves more information, sampling with replacement is sometimes preferred due to its simplicity; thus both schemas are important for applications.

The concept of sliding windows expresses the importance of recent data for applications. In this model, analysis is restricted to the most recent portion of the stream; the outdated elements must not be considered. The significance of sliding windows has been emphasized from the very beginning of data stream research. We cite, e.g., the influential paper of Babcock, Babu, Datar, Motwani and Widom [8]:

Imposing sliding windows on data streams is a natural method for approximation that has several attractive properties. It is well-defined and easily understood: the semantics of the approximation are clear, so that users of the system can be confident that they understand what is given up in producing the approximate answer. It is deterministic, so there is no danger that unfortunate random choices will produce a bad approximation. Most importantly, it emphasizes recent data, which in the majority of real-world applications is more important and relevant than old data: if one is trying in real-time to make sense of network traffic patterns, or phone call or transaction records, or scientific sensor data, then in general insights based on the recent past will be more informative and useful than insights based on stale data. In fact, for many such applications,

sliding windows can be thought of not as an approximation technique reluctantly imposed due to the infeasibility of computing over all historical data, but rather as part of the desired query semantics explicitly expressed as part of the users query.

The importance of the sliding windows model is well illustrated by the considerable amount of relevant papers in both theory and database communities. A small sample subset of relevant papers includes the work of Arasu, Babcock, Babu, Cieslewicz, Datar, Ito, Motwani, Srivastava and Widom [5]; Arasu and Manku [6]; Ayad and Naughton [7]; Babcock, Babu, Datar, Motwani and Thomas [9]; Babcock, Babu, Datar, Motwani and Widom [8]; Babcock, Datar and Motwani [10, 11]; Babcock, Datar, Motwani and O’Callaghan [12]; Das, Gehrke and Riedewald [34]; Datar, Gionis, Indyk and Motwani [36]; Datar and Motwani, Chapter 8, [1]; Datar and Muthukrishnan [37]; Feigenbaum, Kannan and Zhang [39]; Gibbons and Tirthapura [50]; Golab, DeHaan, Demaine, Lopez-Ortiz and Munro [51]; Golab and Özsu [52]; Lee and Ting [61]; Li, Maier, Tufte, Papadimos and Tucker [64]; and Tatbul and Zdonik [71].

Two types of sliding windows are widely recognized. *Fixed-size* (or *sequence-based*) windows define a fixed amount of the most recent elements to be active. For instance an application may restrict an analysis to the last trillion of the elements. Fixed-size windows are important for applications where the arrival rate of the data is fixed (but still extremely fast), such as sensors or stock market measurements. In *bursty* (or *timestamp-based*) windows the validity of an element is defined by an additional parameter such as a timestamp. For instance, an application may restrict an analysis to elements that arrived within the last hour. Bursty windows are important for applications with asynchronous data arrivals, such as networking or database applications.

The importance of both concepts raises two natural questions of *optimal sampling from sliding windows*:

Question 1.1. *Is it possible to maintain a uniform random sampling from sliding windows using provably optimal memory bounds?*

Question 1.2. *Is sampling from sliding windows algorithmically harder than sampling from the entire stream?*

In this paper, perhaps somewhat surprisingly, we definitively answer both questions. Informally, what we show is that it is possible to “translate” (with optimal deterministic memory bounds for all sampling-based algorithms) sampling with and without replacement on the entire stream to sampling with or without replacement in all variants of the sliding windows model. We state precise results in Theorems 2.1, 2.2, 3.6, 4.4 and 5.1.

1.1 Discussion and Related Work

In spite of their apparent simplicity, both Questions 1.1 and 1.2 have no trivial solution. Indeed, the sliding windows

model implies eventual deletions of samples; thus, none of the well-known methods for insertion-only streams (such as the reservoir method [72]) are applicable. Moreover, the deletions are *implicit*, i.e., they are not triggered by an explicit user’s request. Thus, the algorithms for streams with explicit deletions (such as [32, 43]) do not help. Let us illustrate the inherent difficulty of sampling from sliding windows with the following example. Consider the problem of maintaining a single sample from a window of the last 2^{50} elements. Assume that the 100-th element is picked as an initial sample. Eventually, the $(2^{50} + 100)$ -th element arrives, in which case the sample is outdated. But at this time, the data has already been passed and cannot be sampled.

Babcock, Datar and Motwani [10] were the first to address the problem of sampling from sliding windows. They criticized “naive” solutions and stressed the importance of further improvements; we repeat their important arguments here. Periodic sampling is a folklore method for fixed-size windows. This method replaces an expired sample with the newly arrived element. While this method gives worst-case optimal memory bounds, the predictability of samples is not acceptable for many applications (see also [48, 55]). Another obvious approach is over-sampling using the Bernoulli method. If k samples are required, then each element is picked w.p. $O(\frac{k \log n}{n})$, independently, in the hope that at least k samples are not expired. This method has two drawbacks: first, the expected memory is $O(k \log n)$, where n is a window size; and second, with small but positive probability, the sample size will be smaller than k .

The key idea of the elegant algorithms in [10] is a “successors list”; in fact, this idea has been used in almost all subsequent papers. The successors list method suggests backing up a sample with a list of active successors. When a sample expires, the next successor in the list becomes a sample; thus a sample is available at any moment. Based on this idea, Babcock, Datar and Motwani built solutions for sampling with replacement. For sequence-based windows of size n , their *chain sampling* algorithm picks a successor among n future elements, and stores it as it arrives. They show that such schema has an expected memory bound of $O(k)$ and with high probability will not exceed $O(k \log n)$. For timestamp-based windows, their *priority sampling* method associates a priority with every new element. A priority is a random number from $(0, 1)$; a sample is an element p with highest priority; and a sample’s successor is an element with the highest priority among all elements that arrived after p . Priority sampling requires $O(k \log n)$ memory words in expectation and with high probability. There is a lower bound of $\Omega(k \log n)$ for timestamp-based windows that was shown in [48]. Thus, the methods of Babcock, Datar and Motwani are *optimal in expectation*. However, the inherent problem of the replacement method is that the size of the successors list is itself a random variable; thus this method cannot provide *worst-case optimal* bounds. Moreover, Babcock, Datar and Motwani suggested over-sampling as a solution for sampling without replacement; thus the problem of further improvements is implicitly present in their paper. In his excellent survey, Haas [55] gave a further detailed discussion of their solutions.

Zhang, Li, Yu, Wang and Jiang [73] provide an adaptation of reservoir sampling to sliding windows. However, their approach implies storing the window in memory; thus it is applicable only for small windows. In an important work, Gemulla and Lehner [48] addressed the question of sampling without replacement for timestamp-based windows. They suggest a natural extension of priority sampling by maintaining a list of elements with k -highest priorities. This gives an expected $O(k \log n)$ solution. However, their memory bounds are still randomized. Gemulla [47] and Gemulla and Lehner [48] recently addressed the problem of random sampling from timestamp-based windows with a bounded memory. This setting is different from the original problem of Babcock, Datar and Motwani [10]. Namely, it introduces additional uncertainty in the following sense: there is no guarantee that a sample is available at any moment. They provide a “local” lower bound on the success probability that depends on the window’s data. However, there is no “global” lower bound; as Gemulla [47] states in his thesis:

we cannot guarantee a global lower bound other than 0 that holds at any arbitrary time without a-priori knowledge of the data stream.

Thus, the problem of optimal sampling for the entire lifespan of the stream from sliding windows remained an open problem for all versions until today. We stress that, while this problem is important in its own right, it also has further implication for many other problems. Indeed, uniform random sampling is a key tool for many streaming problems (see, e.g., [66]). “Translations” to sliding windows using previous methods introduce randomized complexity instead of deterministic memory bounds (see, e.g., [21]).

1.2 Our Contribution

In this paper we answer affirmatively to Questions 1.1 and 1.2 for all variants of the problem, i.e., for sampling with and without replacement from fixed-size or bursty windows. Our solutions have provable optimal memory guarantees and are stated precisely in Theorems 2.1, 2.2, 3.6, 4.4 and 5.1. In particular, we give $O(k)$ bounds for fixed-size windows (for sampling with or without replacement) and $O(k \log n)$ bounds for bursty windows (for sampling with or without replacement). This is a strict improvement over previous methods that closes the gap between randomized and deterministic complexity, an important fact in its own right.

Furthermore, we prove a perhaps somewhat surprising result: that the memory complexity of the sampling-based algorithm for all variants of the sliding windows model is comparable with that of streaming models (i.e., without the sliding windows). This is the first result of this type, since all previous “translations” of sampling-based algorithms to sliding windows incur randomized memory guarantees only.

Finally, this paper introduces surprisingly simple (yet novel) techniques that are different from all previous approaches. In particular, we reduce the problem of sampling without replacement to the problem of sampling with replacement for

all variants of the sliding windows model. This may be of independent interest, since the former method is a more general case than the latter; thus our paper also proves equivalence for sliding windows, as discussed in the next section.

1.3 High-Level Ideas of Our Approach

We start by describing four key ideas: *equivalent-width partitions*, *covering decomposition*, *generating implicit events*, and *black-box reduction*. Also, we outline our approach by giving high-level descriptions of the most important steps.

Equivalent-Width Partitions. Our methods for the sequence-based windows are based on a surprisingly simple (yet novel) idea of *equivalent-width partitions*: consider sets A, B, C such that $C \subseteq B \subseteq A \cup C$ and $A \cap C = \emptyset$ and $|B| = |A|$. Our goal is to obtain a sample from B , given samples from A and C . We use the following rule: if a sample from A belongs to B , then we assign it to be a sample from B ; otherwise, we assign a sample from C to be a sample from B . The direct computations show the correctness of this schema, i.e., the result is always a uniform sample from B .

As a next step, observe that the above idea can be applied to the sliding windows model. We partition (logically) the entire stream into disjoint intervals (we call them *buckets*) of size n , where n is the size of the window. For each bucket we maintain a random sample using any one-pass algorithm (e.g., the reservoir sampling method). If the window coincides with the most recent bucket, then our task is easy; we assign this bucket’s sample to be the output. Otherwise, the window intersects the two most recent buckets. It must be the case that the most recent bucket is “partial”; i.e., not all elements have arrived yet. But this case matches precisely our key idea: the most recent bucket corresponds to C , our window corresponds to B and the second-most recent bucket corresponds to A . We thus can apply the above rule and obtain our sample using only samples from the two buckets. Since we need only these samples, the optimality of our schema is straightforward.

The above idea can be generalized to the sampling without replacement. Indeed, we show that, given k -samples without replacement from A and C , we take the portion of A ’s sample that belongs to B and complete it with the random portion from C ’s sample. We show that the result is a k -sample without replacement from B . As before, we apply this idea to sliding windows; the detailed proofs can be found in the main body of this paper.

Covering Decomposition and Generating Implicit Events.

For timestamp-based windows, the size of the window is unknown; moreover, it was shown (see, e.g., [36]) that the size of the window cannot be computed precisely with sub-linear memory. This negative result is a key difference between timestamp-based windows and all other models, such as insertion-only streams and streams with explicit deletions (the turnstile model). In fact, this negative result is one of

the main reasons for the randomized bounds in previous solutions. Indeed, it is not clear at all how to obtain uniformity if even the size of the sampled domain is unknown.

Our key observation is that it is possible to sample from a window without even a knowledge of its size. As before, consider disjoint sets A, B, C such that $C \subseteq B \subseteq A \cup C$ and $A \cap C = \emptyset$. In the current scenario we do not assume that $|A| = |B|$ and still obtain samples from B . We show that if $|A| \leq |B|$, and it is possible to generate random events w.p. $\frac{|A|}{|B|}$, then it is possible to “combine” the samples from A and C into a sample from B . The new rule is a generalization of our above ideas. We assign the sample from A to be a sample from B if the A ’s sample belongs to B (for technical reasons, we decrease the probability of this event by $\frac{|A|}{|B|}$ multiplicative factor). Otherwise, we assign the sample from C to be the sample from B . We show that this rule gives a uniform sample from B .

To apply this idea to sliding windows, we need to overcome two problems. First, we must be able to maintain such an A and C (as before we associate B with our window). This task is nontrivial, since the size of the window is unknown. Our second key idea is a novel *covering decomposition* structure. Using this structure, we are able to maintain such an A and C at any moment. In fact, we obtain an important property that $|A| \leq |C|$. This structure can be seen as a modification and generalization of the smooth histogram [19] method.

Second, we need the ability to generate events w.p. $\frac{|A|}{|B|}$ which is still an unknown probability since $|B|$ is the size of our window. Our third key idea is a novel technique that we call *generating implicit events*. At the heart of our technique lies the idea of gradually decreasing the probabilities, starting from 1, until we achieve the desired probability of $|A|/|B|$. In particular, we show that it is possible to generate a non-uniform distribution over the elements of A , where the probability of picking an element is a function of the element’s timestamp (or index). The function is constructed in such a way that the probability of picking an element among the last i elements of A is equal to $\frac{i}{|C|+i}$. That is, the probability of picking an expired element is $\frac{|C|}{|B|}$. Since $|A| \leq |C|$ and since we know the values of $|A|$ and $|C|$, it is possible to generate events w.p. $\frac{|A|}{|B|}$. The details can be found in the main body of this paper.

Black-Box Reduction. Finally, we show that a k -sample without replacement may be generated from k independent samples, R_0, \dots, R_{k-1} . We apply our fourth key idea, a *black-box reduction* from sampling without replacement to sampling with replacement. The novelty of our approach is based on sampling from different domains; in fact, R_i samples all but i last active elements. Such samples can be generated if, in addition, we store the last k elements.

Independence of Disjoint Windows. Our algorithms generate independent samples for non-overlapping windows. The independency follows from the nice property of the reservoir

algorithm (that we use to generate samples in the buckets). Let R_1 be a sample generated for the bucket B , upon arrival of i elements of B . Let R_2 be a fraction of the final sample (i.e., the sample when the last element of B arrives) that belongs to the last $|B| - i$ elements. The reservoir algorithm implies that R_1 and R_2 are independent. Since the rest of the buckets contain independent samples as well, we conclude that our algorithms are independent for non-overlapping windows.

1.4 Roadmap and Notations

We use the following notations throughout our paper. We denote by D a stream and by $p_i, i \geq 0$ its i -th element. For $0 \leq x < y$ we define $[x, y] = \{i, x \leq i \leq y\}$. Finally, bucket $B(x, y)$ is the set of all stream elements between p_x and p_{y-1} : $B(x, y) = \{p_i, i \in [x, y - 1]\}$.

Our bounds are expressed in memory words; that is we assume that a single memory word is sufficient to store a stream element or its index or a timestamp.

Section 2 presents sampling for sequence-based windows, with and without replacement. Sections 3 and 4 are devoted to sampling for timestamp-based windows, with and without replacement. Section 5 outlines possible applications for our approach. Due to the lack of space, some proofs are omitted from the main body of the paper, but they can all be found in the appendix.

2. Equivalent-Width Partitions and Sampling for Sequence-Based Windows

2.1 Sampling With Replacement

Let n be the predefined size of a window. We say that a bucket is *active* if all its elements have arrived and at least one element is non-expired. We say that a bucket is *partial* if not all of its elements have arrived. We show below how to create a single random sample. To create a k -random sample, we repeat the procedure k times, independently.

We divide D into buckets $B(in, (i + 1)n), i = 0, 1, \dots$. At any point in time, we have exactly one active bucket and at most one partial bucket. For every such bucket B , we independently generate a single sample, using the reservoir algorithm [72]. We denote this sample by X_B .

Let B be a partial bucket and $C \subseteq B$ be the set of all arrived elements. The properties of the reservoir algorithm imply that X_B is a random sample of C .

Below, we construct a random sample Z of all non-expired elements. Let U be the active bucket. If there is no partial bucket, then U contains only all non-expired elements. Therefore, $Z = X_U$ is a valid sample. Otherwise, let V be the partial bucket. Let $U_e = \{x : x \in U, x \text{ is expired}\}, U_a = \{x : x \in U, x \text{ is non-expired}\}, V_a = \{x : x \in V, x \text{ arrived}\}$.

Note that $|V_a| = |U_e|$ and let $s = |V_a|$. Also, note that

our window is $U_a \cup V_a$ and X_V is a random sample of V_a . The random sample Z is constructed as follows. If X_U is not expired, we put $Z = X_U$, otherwise $Z = X_V$. To prove the correctness, let p be a non-expired element. If $p \in U_a$, then $P(Z = p) = P(X_U = p) = \frac{1}{n}$. If $p \in V_a$, then

$$P(Z = p) = P(X_U \in U_e, X_V = p) =$$

$$P(X_U \in U_e)P(X_V = p) = \frac{s}{n} \frac{1}{s} = \frac{1}{n}.$$

Therefore, Z is a valid random sample. We need to store only samples of active or partial buckets. Since the number of such buckets is at most two and the reservoir algorithm requires $\Theta(1)$ memory, the total memory of our algorithm for k -sample is $\Theta(k)$. Thus,

Theorem 2.1. *It is possible to maintain k -sampling with replacement for sequence-based windows using $O(k)$ memory words.*

2.2 Sampling Without Replacement

We can generalize the idea above to provide a k -random sample without replacement. In this section k -sample means k -random sampling without replacement.

We use the same buckets $B(in, (i + 1)n), i = 0, 1, \dots$. For every such bucket B , we independently generate a k -sample X_B , using the reservoir algorithm.

Let B be a partial bucket and $C \subseteq B$ be the set of all arrived elements. The properties of the reservoir algorithm imply that either $X_B = C$, if $|C| < k$, or X_B is a k -sample of C . In both cases, we can generate an i -sample of C using X_B only, for any $0 < i \leq \min(k, |C|)$.

Our algorithm is as follows. Let U be the active bucket. If there is no partial bucket, then U contains only all active elements. Therefore, we can put $Z = X_U$. Otherwise, let V be the partial bucket. We define U_e, U_a, V_a, s as before and construct Z as follows. If all elements of X_U are not expired, $Z = X_U$. Otherwise, let i be the number of expired elements, $i = |U_e \cap X_U|$. As we mentioned before, we can generate an i -sample of V_a from X_V , since $i \leq \min(k, s)$. We denote this sample as X_V^i and put

$$Z = (X_U \cap U_a) \cup X_V^i.$$

We will prove now that Z is a valid random sample. Let $Q = \{p_{j_1}, \dots, p_{j_k}\}$ be a fixed set of k non-expired elements such that $j_1 < j_2 < \dots < j_k$. Let $i = |Q \cap V_a|$, so $\{p_{j_1}, \dots, p_{j_{k-i}}\} \subseteq U_a$ and $\{p_{j_{k-i+1}}, \dots, p_{j_k}\} \subseteq V_a$. If $i = 0$, then $Q \subseteq U$ and

$$P(Z = Q) = P(X_U = Q) = \frac{1}{\binom{n}{k}}.$$

Otherwise, by independency of X_U and X_V^i

$$P(Z = Q) =$$

$$P(|X_U \cap U_e| = i, \{p_{j_1}, \dots, p_{j_{k-i}}\} \subseteq X_U,$$

$$X_V^i = \{p_{j_{k-i+1}}, \dots, p_{j_k}\}) =$$

$$P(|X_U \cap U_e| = i, \{p_{j_1}, \dots, p_{j_{k-i}}\} \subseteq X_U)^*$$

$$P(X_V^i = \{p_{j_{k-i+1}}, \dots, p_{j_k}\}) =$$

$$\frac{\binom{s}{i}}{\binom{n}{k}} * \frac{1}{\binom{s}{i}} = \frac{1}{\binom{n}{k}}.$$

Therefore, Z is a valid random sample of non-expired elements. Note that we store only samples of active or partial buckets. Since the number of such buckets is at most two and the reservoir algorithm requires $O(k)$ memory, the total memory of our algorithm is $O(k)$. Thus,

Theorem 2.2. *It is possible to maintain k -sampling without replacement for sequence-based windows using $O(k)$ memory words.*

3. Sampling With Replacement for Timestamp-Based Windows

Let $n = n(t)$ be the number of non-expired elements. For each element p , timestamp $T(p)$ represents the moment of p 's entrance. For a window with (predefined) parameter t_0 , p is active at time t if $t - T(p) < t_0$. We show below how to create a single random sample. To create a k -random sample, we repeat the procedure k times, independently.

3.1 Notations

A bucket structure $BS(x, y)$ is a set

$$\{p_x, x, y, T(x), R_{x,y}, Q_{x,y}, r, q\},$$

where $T(x)$ is a timestamp of p_x , $R_{x,y}$ and $Q_{x,y}$ are independent random samples from $B(x, y)$ and r, q are indexes of the picked (for random samples) elements. We denote by $N(t)$ the size of D at the moment t and by $l(t)$ the index of the earliest active element. Note that $N(t) \leq N(t+1)$, $l(t) \leq l(t+1)$ and $T(p_i) \leq T(p_{i+1})$.

3.2 Covering Decomposition

Let $a \leq b$ be two indexes. A covering decomposition of a bucket $B(a, b)$, $\zeta(a, b)$, is an ordered set of bucket structures with independent samples inductively defined below.

$$\zeta(b, b) := BS(b, b+1),$$

and for $a < b$,

$$\zeta(a, b) := \langle BS(a, c), \zeta(c, b) \rangle,$$

where $c = a + 2^{\lfloor \log(b+1-a) \rfloor - 1}$. Note that

$$|\zeta(a, b)| = O(\log(b-a)),$$

so $\zeta(a, b)$ uses $O(\log(b-a))$ memory.

Given p_{b+1} , we inductively define an operator $Incr(\zeta(a, b))$ as follows.

$$Incr(\zeta(b, b)) := \langle BS(b, b+1), BS(b+1, b+2) \rangle.$$

For $a < b$, we put

$$Incr(\zeta(a, b)) := \langle BS(a, v), Incr(\zeta(v, b)) \rangle,$$

where v is defined below.

If $\lfloor \log(b+2-a) \rfloor = \lfloor \log(b+1-a) \rfloor$, then we put $v = c$, where $BS(a, c)$ is the first bucket structure of $\zeta(a, b)$. Otherwise, we put $v = d$, where $BS(c, d)$ is the second bucket structure of $\zeta(a, b)$. (Note that $\zeta(a, b)$ contains at least two buckets for $a < b$.)

We show how to construct $BS(a, d)$ from $BS(a, c)$ and $BS(c, d)$. We have in this case $\lfloor \log(b+2-a) \rfloor = \lfloor \log(b+1-a) \rfloor + 1$, and therefore $b+1-a = 2^i - 1$ for some $i \geq 2$.

Thus $c - a = 2^{\lfloor \log(2^i - 1) \rfloor - 1} = 2^{i-2}$ and

$$\lfloor \log(b+1-c) \rfloor = \lfloor \log(b+1-a - (c-a)) \rfloor =$$

$$\lfloor \log(2^i - 2^{i-2} - 1) \rfloor = i - 1.$$

Thus $d - c = 2^{\lfloor \log(b+1-c) \rfloor - 1} = 2^{i-2} = c - a$. Now we can create $BS(a, v)$ by unifying $BS(a, c)$ and $BS(c, d)$: $BS(a, v) = \{p_a, d - a, R_{a,d}, Q_{a,d}, r', q'\}$. We put $R_{a,d} = R_{a,c}$ with probability $\frac{1}{2}$ and $R_{a,d} = R_{c,d}$ otherwise. Since $d - c = c - a$, and $R_{c,d}, R_{a,c}$ are distributed uniformly, we conclude that $R_{a,d}$ is distributed uniformly as well. $Q_{a,d}$ is defined similarly and r', q' are indexes of the chosen samples. Finally, the new samples are independent of the rest of ζ 's samples. Note also that $Incr(\zeta(a, b))$ requires $O(\log(b-a))$ operations.

Lemma 3.1. *For any a and b , $Incr(\zeta(a, b)) = \zeta(a, b+1)$.*

Lemma 3.2. *For any t with a positive number of active elements, we are able to maintain one of the following:*

1. $\zeta(l(t), N(t))$,

or

2. $BS(y_t, z_t), \zeta(z_t, N(t))$,

where $y_t < l(t) \leq z_t$, $z_t - y_t \leq N(t) + 1 - z_t$ and all random samples are independent.

3.3 Generating Implicit Events

We use the following notations for this section. Let $B_1 = B(a, b)$ and $B_2 = B(b, N(t) + 1)$ be two buckets such that p_a is expired, p_b is active and $|B_1| \leq |B_2|$. Let BS_1 and BS_2 be corresponding bucket structures, with independent random samples R_1, Q_1 and R_2, Q_2 . We put $\alpha = b - a$ and $\beta = N(t) + 1 - b$. Let γ be the (unknown) number of non-expired elements inside B_1 , so $n = \beta + \gamma$. We stress that α, β are known and γ is unknown.

Lemma 3.3. *It is possible to generate a random sample $Y = Y(Q_1)$ of B_1 , with the following distribution:*

$$P(Y = p_{b-i}) = \frac{\beta}{(\beta+i)(\beta+i-1)}, \quad 0 < i < \alpha,$$

$$P(Y = p_a) = \frac{\beta}{\beta + \alpha - 1}.$$

Y is independent of R_1, R_2, Q_2 and can be generated within constant memory and time, using Q_1 .

Lemma 3.4. *It is possible to generate a zero-one random variable X such that $P(X = 1) = \frac{\alpha}{\beta + \gamma}$. X is independent of R_1, R_2, Q_2 and can be generated using constant time and memory.*

Lemma 3.5. *It is possible to construct a random sample V of all non-expired elements using only the data of BS_1, BS_2 and constant time and memory.*

3.4 Main Results

Theorem 3.6. *We can maintain a random sample over all non-expired elements using $\Theta(\log n)$ memory.*

PROOF. By using Lemma 3.2, we are able to maintain one of two cases. If case 1 occurs, we can combine random variables of all bucket structures with appropriate probabilities and get a random sample of all non-expired elements. If case 2 occurs, we use notations of Section 3.3, interpret the first bucket as B_1 and combine buckets of covering decomposition to generate samples from B_2 . Properties of the second case imply $|B_1| \leq |B_2|$ and therefore, by using Lemma 3.5, we are able to produce a random sample as well. All procedures described in the lemmas require $\Theta(\log n)$ memory. Therefore, the theorem is correct. \square

Lemma 3.7. *The memory usage of maintaining a random sample within a timestamp-based window has a lower bound $\Omega(\log(n))$.*

4. Black-Box Reduction

In this section, we present black-box reduction from sampling without replacement to sampling with replacement. As a result, we obtain an optimal algorithm for sampling without replacement for timestamp-based windows. Informally, the idea is as follows. We maintain k independent random samples R_0, \dots, R_{k-1} of active elements, using the algorithm from Section 3. The difference between these samples and the k -sample with replacement is that R_i samples all active elements except the last i . This can be done using $O(k + k \log n)$ memory. Finally, a k -sample without replacement can be generated using R_0, \dots, R_{k-1} only.

Let us describe the algorithm in detail. First, we construct R_i . To do this, we maintain an auxiliary array with the last i elements. We repeat all procedures in Section 3, but we “delay” the last i elements. An element is added to covering decomposition only when more than i elements arrive after it. We prove the following variant of Lemma 3.2.

Lemma 4.1. *Let $0 < i \leq k$. For any t with more than i active elements, we are able to maintain one of the following:*

1. $\zeta(l(t), N(t) - i),$

or

2. $BS(y_t, z_t), \zeta(z_t, N(t) - i),$

where $y_t < l(t) \leq z_t$ and $z_t - y_t \leq N(t) + 1 - i - z_t$ and all random samples of the bucket structures are independent.

The proof is presented in the appendix. The rest of the procedure remains the same. Note that we can use the same array for every i , and therefore we can construct R_0, \dots, R_{k-1} using $\Theta(k + k \log n)$ memory.

In the reminder of this section, we show how R_0, \dots, R_{k-1} can be used to generate a k -sample without replacement. We denote by R_i^j a i -random sample without replacement from $[1, j]$.

Lemma 4.2. *R_{a+1}^{b+1} can be generated using independent R_a^b, R_1^{b+1} samples only.*

Lemma 4.3. *R_k^n can be generated using only independent samples $R_1^n, R_1^{n-1}, \dots, R_1^{n-k+1}$.*

PROOF. By using Lemma 4.2, we can generate R_2^{n-k+2} using R_1^{n-k+1} and R_1^{n-k+2} . We can repeat this procedure and generate $R_j^{n-k+j}, 2 \leq j \leq k$, using $R_{j-1}^{n-k+j-1}$ (that we already constructed by induction) and R_1^{n-k+j} . For $j = k$ we have R_k^n . \square

By using Lemma 4.3, we can generate a k -sample without replacement using only R_0, \dots, R_{k-1} . Thus, we have proved

Theorem 4.4. *It is possible to maintain k -sampling without replacement for timestamp-based windows using $O(k \log n)$ memory words.*

5. Applications

Consider that algorithm Λ is sampling-based, i.e., it operates on a uniformly chosen subset of D instead of the whole stream. Such an algorithm can be immediately transformed to sliding windows by replacing the underlying sampling method with our algorithms. We obtain the following general result and illustrate it with the examples below.

Theorem 5.1. *For the sampling-based algorithm Λ that solves problem P , there exists an algorithm Λ' that solves P on sliding windows. The memory guarantees are preserved for sequence-based windows and have a multiplicative overhead of $\log n$ for timestamp-based windows.*

Frequency moment is a fundamental problem in data stream processing. Given a stream of elements, such that $p_j \in [m]$, the frequency x_i of each $i \in [m]$ is defined as $|\{j | p_j = i\}|$ and the k -th frequency moment is defined as $F_k = \sum_{i=1}^m x_i^k$. The first algorithm for frequency moments for $k > 2$ was proposed in the seminal paper of Alon, Matias and Szegedy [4]. They present an algorithm that uses $O(m^{1-\frac{1}{k}})$ memory. Numerous improvements to lower and upper bounds have been reported, including the works of Bar-Yossef, Jayram, Kumar and Sivakumar [14], Chakrabarti, Khot and Sun [23], Coppersmith and Kumar [33], and Ganguly[44]. Finally, Indyk and Woodruff [57] and later Bhuvanagiri, Ganguly, Kesh and Saha [18] presented algorithms that use $\tilde{O}(m^{1-\frac{1}{k}})$ memory and are optimal. The algorithm of Alon, Matias

and Szegedy [4] is sampling-based, thus we can adapt it to sliding windows using our methods. The memory usage is not optimal, however this is the first algorithm for frequency moments over sliding windows that works for all k . Recently Braverman and Ostrovsky [19] adapted the algorithm from [18] to sliding windows, producing a memory-optimal algorithm that uses $\tilde{O}(m^{1-\frac{2}{k}})$. However, it involves k^k multiplicative overhead, making it infeasible for large k ; thus these results generally cannot be compared. We have

Corollary 5.2. *For any $k > 2$, there exists an algorithm that maintains an (ϵ, δ) -approximation of the k -th frequency moment over sliding windows using $\tilde{O}(m^{1-\frac{1}{k}})$ bits.*

Recently, numerous graph problems were addressed in the streaming environment. Stream elements represent edges of the graph, given in arbitrary order (we refer readers to [20] for a detailed explanation of the model). One of the fundamental graph problems is estimating a number of small cliques in a graph, in particular the number of triangles. Effective solutions were proposed by Jowhari and Ghodsi [59], Bar-Yosseff, Kumar and Sivakumar [15] and Buriol, Frahling, Leonardi, Marchetti-Spaccamela and Sohler [20]. The last paper presented an (ϵ, δ) -approximation algorithm that uses $O(1 + \frac{\log |E|}{|E|} \frac{1}{\epsilon^2} \frac{|T_1|+2|T_2|+3|T_3|}{|T_3|} \log \frac{2}{\delta})$ memory ([20], Theorem 2) that is the best result so far. Here, $|T_i|$ represents the number of node-triplets having i edges in the induced sub-graph. The algorithm is applied on a random sample collected using the reservoir method. By replacing the reservoir sampling with our algorithms, we obtain the following result.

Corollary 5.3. *There exists an algorithm that maintains an (ϵ, δ) -approximation of the number of triangles over sliding windows. For sequence-based windows it uses $O(1 + \frac{\log |E_W|}{|E_W|} \frac{1}{\epsilon^2} \frac{|T_1|+2|T_2|+3|T_3|}{|T_3|} \log \frac{2}{\delta})$ memory, where E_W is the set of active edges. Timestamp-based windows adds a multiplicative factor of $\log n$.*

Following [20], our method is also applicable for incidence streams, where all edges of the same vertex come together.

The entropy of a stream is defined as $H = -\sum_{i=1}^m \frac{x_i}{N} \log \frac{x_i}{N}$, where x_i is as above. The entropy norm is defined as $F_H = \sum_{i=1}^m x_i \log x_i$. Effective solutions for entropy and entropy norm estimations were recently reported by Guha, McGregor and Venkatasubramanian [53]; Chakrabarti, Do Ba and Muthukrishnan [22]; Harvey, Nelson and Onak [56]; Chakrabarti, Cormode and McGregor [21]; and Zhao, Lall, Ogihara, Spatscheck, Wang and Xu [74].

The paper of Chakrabarti, Cormode and McGregor presents an algorithm that is based on a variation of reservoir sampling. The algorithm maintains entropy using $O(\epsilon^{-2} \log \delta^{-1})$ that is nearly optimal. The authors also considered the sliding window model and used a variant of priority sampling [10] to obtain the approximation. Thus, the worst-case memory guarantees are not preserved for sliding windows. By replacing priority sampling with our methods, we obtain

Corollary 5.4. *There exists an algorithm that maintains an*

(ϵ, δ) -approximation of entropy on sliding windows using $O(\epsilon^{-2} \log \delta^{-1} \log n)$ memory.

Moreover, our methods can be used with the algorithm from [22] to obtain $\tilde{O}(1)$ memory for large values of the entropy norm. This algorithm is based on reservoir sampling and thus can be straightforwardly implemented in sliding windows. As a result, we build the first solutions with provable memory guarantees on sliding windows.

Our algorithms can be naturally extended to some biased functions. Biased sampling [2] is non-uniform, giving larger probabilities for more recent elements. The distribution is defined by a biased function. We can apply our methods to implement step biased functions, maintaining samples over each window with different lengths and combining the samples with corresponding probabilities. Our algorithm can extend the ideas of Feigenbaum, Kannan, Strauss and Viswanathan [42] for testing and spot-checking to sliding windows. Finally, we can apply our tools to the algorithm of Procopiuc and Procopiuc for density estimation [69], since it is based on the reservoir algorithm as well.

6. References

- [1] C. Aggarwal (editor), *Data Streams: Models and Algorithms*, Springer Verlag, 2007.
- [2] C. Aggarwal, "On biased reservoir sampling in the presence of stream evolution", *Proceedings of the 32nd international conference on Very large data bases*, pp. 607–618, 2006.
- [3] N. Alon, N. Duffield, C. Lund, M. Thorup, "Estimating arbitrary subset sums with few probes". *Proceedings of the twenty-fourth ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems*, pp. 317–325, 2005.
- [4] N. Alon, Y. Matias, M. Szegedy, "The space complexity of approximating the frequency moments". *Proceedings of the twenty-eighth annual ACM symposium on Theory of computing*, pp. 20–29, 1996.
- [5] A. Arasu, B. Babcock, S. Babu, J. Cieslewicz, M. Datar, K. Ito, R. Motwani, U. Srivastava, J. Widom, "STREAM: The Stanford Data Stream Management System", Book Chapter, "Data-Stream Management: Processing High-Speed Data Streams", Springer-Verlag, 2005.
- [6] A. Arasu, G. S. Manku, "Approximate counts and quantiles over sliding windows", *Proceedings of the twenty-third ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems*, 2004.
- [7] A. M. Ayad, J. F. Naughton, "Static optimization of conjunctive queries with sliding windows over infinite streams", *Proceedings of the 2004 ACM SIGMOD international conference on Management of data*, 2004.
- [8] B. Babcock, S. Babu, M. Datar, R. Motwani, J. Widom, "Models and issues in data stream systems", *Proceedings of the twenty-first ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems*, 2002.
- [9] B. Babcock, S. Babu, M. Datar, R. Motwani, D. Thomas, "Operator scheduling in data stream systems", *The VLDB Journal The International Journal on Very Large Data Bases*, v.13 n.4, pp.333–353, 2004.

- [10] B. Babcock, M. Datar, R. Motwani, "Sampling from a moving window over streaming data", *Proceedings of the thirteenth annual ACM-SIAM symposium on Discrete algorithms*, pp.633–634, 2002.
- [11] B. Babcock, M. Datar, R. Motwani, "Load Shedding for Aggregation Queries over Data Streams", *Proceedings of the 20th International Conference on Data Engineering*, 2004.
- [12] B. Babcock, M. Datar, R. Motwani, L. O'Callaghan, "Maintaining variance and k-medians over data stream windows", *Proceedings of the twenty-second ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems*, pp.234–243, 2003.
- [13] Z. Bar-Yossef, "Sampling lower bounds via information theory", *STOC*, 2003.
- [14] Z. Bar-Yossef, T. S. Jayram, R. Kumar, D. Sivakumar, "An Information Statistics Approach to Data Stream and Communication Complexity", *Proceedings of the 43rd Symposium on Foundations of Computer Science*, pp. 209–218, 2002.
- [15] Z. Bar-Yossef, R. Kumar, D. Sivakumar, "Reductions in streaming algorithms, with an application to counting triangles in graphs", *Proceedings of the thirteenth annual ACM-SIAM symposium on Discrete algorithms*, pp.623–632, 2002.
- [16] Z. Bar-Yossef, T. S. Jayram, R. Kumar, D. Sivakumar, L. Trevisan, "Counting Distinct Elements in a Data Stream", *Proceedings of the 6th International Workshop on Randomization and Approximation Techniques*, pp.1-10, 2002.
- [17] Z. Bar-Yossef, R. Kumar, D. Sivakumar, "Sampling algorithms: lower bounds and applications", *STOC*, 2001.
- [18] L. Bhuvanagiri, S. Ganguly, D. Kesh, C. Saha, "Simpler algorithm for estimating frequency moments of data streams", *Proceedings of the seventeenth annual ACM-SIAM symposium on Discrete algorithm*, pp.708–713, 2006.
- [19] V. Braverman, R. Ostrovsky, "Smooth histograms on stream windows", *Proceedings of the 48th Symposium on Foundations of Computer Science*, 2007.
- [20] L. S. Buriol, G. Frahling, S. Leonardi, A. Marchetti-Spaccamela, C. Sohler, "Counting triangles in data streams", *Proceedings of the twenty-fifth ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems*, pp.253–262, 2006.
- [21] A. Chakrabarti, G. Cormode, A. McGregor, "A near-optimal algorithm for computing the entropy of a stream". In *Proceedings of ACM-SIAM Symposium on Discrete Algorithms*, 2007.
- [22] A. Chakrabarti, K. Do Ba, S. Muthukrishnan, "Estimating Entropy and Entropy Norm on Data Streams", In *Proceedings of the 23rd International Symposium on Theoretical Aspects of Computer Science*, 2006.
- [23] A. Chakrabarti, S. Khot, X. Sun, "Near-optimal lower bounds on the multi-party communication complexity of set-disjointness", *Proceedings of the 18th Annual IEEE Conference on Computational Complexity*, 2003.
- [24] K. L. Chang, R. Kannan, "The space complexity of pass-efficient algorithms for clustering", in *ACM-SIAM Symposium on Discrete Algorithms*, 2006, pp. 1157-1166.
- [25] M. Charikar, C. Chekuri, T. Feder, R. Motwani, "Incremental clustering and dynamic information retrieval", *SIAM J. Comput.*, 33 (2004), pp. 1417-1440.
- [26] K. Chaudhuri, N. Mishra, "When Random Sampling Preserves Privacy", *CRYPTO*, 2006.
- [27] S. Chaudhuri, R. Motwani, V. Narasayya, "On random sampling over joins", *Proceedings of the 1999 ACM SIGMOD international conference on Management of data*, pp.263-274, 1999.
- [28] Y. Chi, H. Wang, P. S. Yu, R. R. Muntz, "Moment: Maintaining Closed Frequent Itemsets over a Stream Sliding Window", *Fourth IEEE International Conference on Data Mining (ICDM'04)*, pp. 59–66, 2004.
- [29] E. Cohen, "Size-estimation framework with applications to transitive closure and reachability," *Journal of Computer and System Sciences*, v.55 n.3, pp.441–453, 1997.
- [30] E. Cohen, H. Kaplan, "Summarizing data using bottom-k sketches," *Proceedings of the twenty-sixth annual ACM symposium on Principles of distributed computing*, 2007.
- [31] G. Cormode, M. Datar, P. Indyk, S. Muthukrishnan, "Comparing Data Streams Using Hamming Norms (How to Zero In)", *IEEE Transactions on Knowledge and Data Engineering*, v.15 n.3, pp.529–540, 2003.
- [32] G. Cormode, S. Muthukrishnan, I. Rozenbaum, "Summarizing and mining inverse distributions on data streams via dynamic inverse sampling", *Proceedings of the 31st international conference on Very large data bases*, 2005.
- [33] D. Coppersmith, R. Kumar, "An improved data stream algorithm for frequency moments", *Proceedings of the fifteenth annual ACM-SIAM symposium on Discrete algorithms*, pp.151–156, 2004.
- [34] A. Das, J. Gehrke, M. Riedewald, "Semantic Approximation of Data Stream Joins", *IEEE Transactions on Knowledge and Data Engineering*, v.17 n.1, pp.44–59, 2005.
- [35] A. Dasgupta, P. Drineas, B. Harb, R. Kumar, M. W. Mahoney, "Sampling algorithms and coresets for l_p regression", *SODA*, 2008.
- [36] M. Datar, A. Gionis, P. Indyk, R. Motwani, "Maintaining stream statistics over sliding windows: (extended abstract)", *Proceedings of the thirteenth annual ACM-SIAM symposium on Discrete algorithms*, pp.635–644, 2002.
- [37] M. Datar, S. Muthukrishnan, "Estimating Rarity and Similarity over Data Stream Windows", *Proceedings of the 10th Annual European Symposium on Algorithms*, pp.323–334, 2002.
- [38] N. Duffield, C. Lund, M. Thorup, "Flow sampling under hard resource constraints", *ACM SIGMETRICS Performance Evaluation Review*, v.32 n.1, 2004.
- [39] J. Feigenbaum, S. Kannan, and J. Zhang, "Computing diameter in the streaming and sliding-window models", *Algorithmica*, 41:25–41, 2005.
- [40] J. Feigenbaum, S. Kannan, A. McGregor, S. Suri, J. Zhang, "Graph distances in the streaming model: the value of space", *SODA*, 2005.
- [41] J. Feigenbaum, S. Kannan, A. McGregor, S. Suri, J. Zhang, "On graph problems in a semi-streaming model", *Theor. Comput. Sci.*, 2005.
- [42] J. Feigenbaum, S. Kannan, M. Strauss, M. Viswanathan, "Testing and Spot-Checking of Data Streams", *Algorithmica*, 34(1): 67-80, 2002.
- [43] G. Frahling, P. Indyk, C. Sohler, "Sampling in dynamic data streams and applications", *Proceedings of the twenty-first annual symposium on Computational geometry*, 2005.

- [44] S. Ganguly, “Estimating Frequency Moments of Update Streams using Random Linear Combinations”, *Proceedings of the 8th International Workshop on Randomized Algorithms*, pp. 369–380, 2004.
- [45] S. Ganguly, “Counting distinct items over update streams”, *Theoretical Computer Science*, pp.211–222, 2007.
- [46] S. Gandhi, S. Suri, E. Welzl, “Catching elephants with mice: sparse sampling for monitoring sensor networks”, *SenSys*, 2007.
- [47] R. Gemulla, “Sampling Algorithms for Evolving Datasets”, PhD Dissertation.
- [48] R. Gemulla and W. Lehner, “Sampling time-based sliding windows in bounded space”, *In Proc. of the 2008 ACM SIGMOD Intl. Conf. on Management of Data*, pp. 379–392.
- [49] P. B. Gibbons, Y. Matias, “New sampling-based summary statistics for improving approximate query answers”, *Proceedings of the 1998 ACM SIGMOD international conference on Management of data*, pp.331–342, 1998.
- [50] P. B. Gibbons, S. Tirthapura, “Distributed streams algorithms for sliding windows”, *Proceedings of the fourteenth annual ACM symposium on Parallel algorithms and architectures*, pp.10–13, 2002.
- [51] L. Golab, D. DeHaan, E. D. Demaine, A. Lopez-Ortiz, J. I. Munro, “Identifying frequent items in sliding windows over on-line packet streams”, *Proceedings of the 3rd ACM SIGCOMM conference on Internet measurement*, 2003.
- [52] L. Golab, M. T. Özsu, “Processing sliding window multi-joins in continuous queries over data streams”, *Proceedings of the 29th international conference on Very large data bases*, pp.500–511, 2003.
- [53] S. Guha, A. McGregor, S. Venkatasubramanian, “Streaming and sublinear approximation of entropy and information distances”, *Proceedings of the seventeenth annual ACM-SIAM symposium on Discrete algorithm*, pp.733-742, 2006.
- [54] S. Guha, A. Meyerson, N. Mishra, R. Motwani, L. O’Callaghan, “Clustering Data Streams: Theory and Practice”, *IEEE Trans. on Knowledge and Data Engineering*, vol. 15, 2003.
- [55] P. J. Haas, “Data stream sampling: Basic techniques and results”, *In M. Garofalakis, J. Gehrke, and R. Rastogi (Eds.), Data Stream Management: Processing High Speed Data Streams*, Springer.
- [56] N. Harvey, J. Nelson, K. Onak, “Sketching and Streaming Entropy via Approximation Theory”, *The 49th Annual Symposium on Foundations of Computer Science (FOCS 2008)*.
- [57] P. Indyk, D. Woodruff, “Optimal approximations of the frequency moments of data streams”, *Proceedings of the thirty-seventh annual ACM symposium on Theory of computing*, pp.202–208, 2005.
- [58] M. Kolonko, D. Wäsch, “Sequential reservoir sampling with a nonuniform distribution”, v.32, i.2, pp.257–273, 2006.
- [59] H. Jowhari, M. Ghodsi, “New streaming algorithms for counting triangles in graphs”, *Proceedings of the 11th COCOON*, pp. 710–716, 2005.
- [60] L. K. Lee, H. F. Ting, “Frequency counting and aggregation: A simpler and more efficient deterministic scheme for finding frequent items over sliding windows”, *Proceedings of the twenty-fifth ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems (PODS ’06)*, pp. 290–297, 2006.
- [61] L. K. Lee, H. F. Ting, “Maintaining significant stream statistics over sliding windows”, *Proceedings of the seventeenth annual ACM-SIAM symposium on Discrete algorithm*, pp.724–732, 2006.
- [62] K. Li, “Reservoir-sampling algorithms of time complexity $O(n(1 + \log(N/n)))$ ”, *ACM Transactions on Mathematical Software (TOMS)*, v.20 n.4, pp.481–493, Dec. 1994.
- [63] J. Li, D. Maier, K. Tufte, V. Papadimos, P. A. Tucker, “Semantics and Evaluation Techniques for Window Aggregates in Data Streams”, *SIGMOD*, 2005.
- [64] J. Li, D. Maier, K. Tufte, V. Papadimos, P. A. Tucker, “No pane, no gain: efficient evaluation of sliding-window aggregates over data streams”, *ACM SIGMOD Record*, v.34 n.1, 2005.
- [65] G. S. Manku, R. Motwani, “Approximate frequency counts over data streams”, *In Proceedings of the 28th International Conference on Very Large Data Bases*, pp.356–357, 2002.
- [66] S. Muthukrishnan, “Data Streams: Algorithms And Applications” *Foundations and Trends in Theoretical Computer Science*, Volume 1, Issue 2.
- [67] C. R. Palmer, C. Faloutsos, “Density biased sampling: an improved method for data mining and clustering”, *Proceedings of the 2000 ACM SIGMOD international conference on Management of data*, pp.82–92, 2000
- [68] V. Paxson, G. Almes, J. Mahdavi, M. Mathis, “Framework for IP performance metrics”, RFC 2330, 1998.
- [69] C. Procopiuc, O. Procopiuc, “Density Estimation for Spatial Data Streams”, *Proceedings of the 9th International Symposium on Spatial and Temporal Databases*, pp.109–126, 2005.
- [70] M. Szegedy, “The DLT priority sampling is essentially optimal”, *Proceedings of the thirty-eighth annual ACM symposium on Theory of computing*, pp.150–158, 2006.
- [71] N. Tatbul, S. Zdonik, “Window-aware load shedding for aggregation queries over data streams”, *Proceedings of the 32nd international conference on Very large data bases*, 2006.
- [72] J. S. Vitter, “Random sampling with a reservoir”, *ACM Transactions on Mathematical Software (TOMS)*, v.11 n.1, pp.37–57, 1985.
- [73] L. Zhang, Z. Li, M. Yu, Y. Wang, Y. Jiang, “Random sampling algorithms for sliding windows over data streams”, *Proc. of the 11th Joint International Computer Conference*, pp. 572–575, 2005.
- [74] H. Zhao, A. Lall, M. Ogihara, O. Spatscheck, J. Wang, J. Xu, “A data streaming algorithm for estimating entropies of od flows”, *Proceedings of the 7th ACM SIGCOMM conference on Internet measurement*, 2007.

APPENDIX

Lemma 3.1. For any a and b , $\text{Incr}(\zeta(a, b)) = \zeta(a, b + 1)$.

PROOF. We prove the lemma by induction on $b - a$. If $a = b$ then, since $b + 1 = b + 2^{\lfloor \log((b+1)+1-b) \rfloor - 1}$, we have, by definition of $\zeta(b, b + 1)$,

$$\begin{aligned} \zeta(b, b + 1) &= \langle BS(b, b + 1), \zeta(b + 1, b + 1) \rangle = \\ &\langle BS(b, b + 1), BS(b + 1, b + 2) \rangle = \text{Incr}(\zeta(b, b)). \end{aligned}$$

We assume that the lemma is correct for $b - a < h$ and prove it for $b - a = h$. Let $BS(a, v)$ be the first bucket of $Incr(\zeta(a, b))$. Let $BS(a, c)$ be the first bucket of $\zeta(a, b)$. By definition, if $\lfloor \log(b + 2 - a) \rfloor = \lfloor \log(b + 1 - a) \rfloor$ then $v = c$. We have

$$v = c = a + 2^{\lfloor \log(b+1-a) \rfloor - 1} = a + 2^{\lfloor \log(b+2-a) \rfloor - 1}.$$

Otherwise, let $BS(c, d)$ be the second bucket of $\zeta(a, b)$. We have from above $\lfloor \log(b + 2 - a) \rfloor = \lfloor \log(b + 1 - a) \rfloor + 1$, $d - c = c - a$ and $v = d$. Thus

$$\begin{aligned} v = d = 2c - a &= 2 \left(a + 2^{\lfloor \log(b+1-a) \rfloor - 1} \right) - a = \\ &= a + 2^{\lfloor \log(b+1-a) \rfloor} = a + 2^{\lfloor \log(b+2-a) \rfloor - 1}. \end{aligned}$$

In both cases $v = a + 2^{\lfloor \log((b+1)+1-a) \rfloor - 1}$ and, by definition of ζ

$$\zeta(a, b + 1) = \langle BS(a, v), \zeta(v, b + 1) \rangle.$$

By induction, since $b - v < h$, we have $Incr(\zeta(v, b)) = \zeta(v, b + 1)$. Thus

$$\begin{aligned} \zeta(a, b + 1) &= \langle BS(a, v), \zeta(v, b + 1) \rangle = \\ &= \langle BS(a, v), Incr(\zeta(v, b)) \rangle = Incr(\zeta(a, b)). \end{aligned}$$

□

Lemma 3.2. *For any t with a positive number of active elements, we are able to maintain one of the following:*

1. $\zeta(l(t), N(t))$,
or
2. $BS(y_t, z_t), \zeta(z_t, N(t))$,
where $y_t < l(t) \leq z_t$, $z_t - y_t \leq N(t) + 1 - z_t$ and all random samples are independent.

PROOF. We prove the lemma by induction on t . First we assume that $t = 0$. If no element arrives at time 0, the stream is empty and we do nothing. Otherwise, we put $\zeta(0, 0) = BS(0, 1)$, and for any $i, 0 < i \leq N(0)$ we generate $\zeta(0, i)$ by executing $Incr(\zeta(0, i - 1))$. Therefore, at the end of this step, we have $\zeta(0, N(0)) = \zeta(l(0), N(0))$. So, the case 1 is true.

We assume that the lemma is correct for t and prove it for $t + 1$.

1. If for t the window is empty, then the procedure is the same as for the basic case.
2. If for t we maintain case 1, then we have three sub-cases.
 - (a) If $p_{l(t)}$ is not expired at the moment $t + 1$, then $l(t + 1) = l(t)$. Similar to the basic case, we apply the $Incr$ procedure for every new element with index $i, N(t) < i \leq N(t + 1)$. Due to the properties of $Incr$, we have at the end $\zeta(l(t + 1), N(t + 1))$. Therefore case 1 is true for $t + 1$.
 - (b) If $p_{N(t)}$ is expired, then our current bucket structures represent only expired elements. We delete them and apply the procedure for the basic case.

- (c) The last sub-case is the one when $p_{N(t)}$ is not expired and $p_{l(t)}$ is expired. Let $\langle BS_1, \dots, BS_h \rangle$, ($BS_i = BS(v_i, v_{i+1})$) be all buckets of $\zeta(l(t), N(t))$. Since $p_{N(t)}$ is not expired, there exists exactly one bucket structure, BS_i , such that p_{v_i} is expired and $p_{v_{i+1}}$ is not expired. We can find it by checking all the bucket structures, since we store timestamps for p_{v_i} s. We put

$$y_{t+1} = v_i, z_{t+1} = v_{i+1}.$$

We have by definition

$$\begin{aligned} \zeta(z_{t+1}, N(t)) &= \zeta(v_{i+1}, N(t)) = \\ &= \langle BS_{i+1}, \dots, BS_k \rangle. \end{aligned}$$

Applying the $Incr$ procedure to all new elements, we construct $\zeta(z_{t+1}, N(t + 1))$. Finally, we have:

$$z_{t+1} - y_{t+1} = v_{i+1} - v_i = 2^{\lfloor \log(N(t)+1-v_i) \rfloor - 1} \leq$$

$$\frac{1}{2}(N(t) + 1 - v_i) = \frac{1}{2}(N(t) + 1 - y_{t+1}).$$

Therefore $z_{t+1} - y_{t+1} \leq N(t) + 1 - z_{t+1} \leq N(t + 1) + 1 - z_{t+1}$. Thus, case 2 is true for $t + 1$. We discard all non-used bucket structures BS_1, \dots, BS_{i-1} .

3. Otherwise, for t we maintain case 2. Similarly, we have three sub-cases.

- (a) If p_{z_t} is not expired at the moment $t + 1$, we put $y_{t+1} = y_t, z_{t+1} = z_t$. We have

$$z_{t+1} - y_{t+1} = z_t - y_t \leq$$

$$N(t) + 1 - z_t \leq N(t + 1) + 1 - z_{t+1}.$$

Again, we add the new elements using the $Incr$ procedure and we construct $\zeta(z_{t+1}, N(t + 1))$. Therefore case 2 is true for $t + 1$.

- (b) If $p_{N(t)}$ is expired, we apply exactly the same procedure as for 2b.
- (c) If p_{z_t} is expired and $p_{N(t)}$ is not expired, we apply exactly the same procedure as for 2c.

Therefore, the lemma is correct. □

Lemma 3.3. *It is possible to generate a random sample $Y = Y(Q_1)$ of B_1 , with the following distribution:*

$$P(Y = p_{b-i}) = \frac{\beta}{(\beta + i)(\beta + i - 1)}, \quad 0 < i < \alpha,$$

$$P(Y = p_a) = \frac{\beta}{\beta + \alpha - 1}.$$

Y is independent of R_1, R_2, Q_2 and can be generated within constant memory and time, using Q_1 .

PROOF. Let $\{H_j\}_{j=1}^{\alpha-1}$ be a set of zero-one independent random variables such that

$$P(H_j = 1) = \frac{\alpha\beta}{(\beta + j)(\beta + j - 1)}.$$

Let $D = B_1 \times \{0, 1\}^{\alpha-1}$ and Z be the random vector with values from D , $Z = \langle Q_1, H_1, \dots, H_{\alpha-1} \rangle$. Let $\{A_i\}_{i=1}^{\alpha}$ be a set of subsets of D :

$$A_i = \{ \langle q_{b-i}, a_1, \dots, a_{i-1}, 1, a_{i+1}, \dots, a_{\alpha-1} \rangle \mid a_j \in \{0, 1\}, j \neq i \}.$$

Finally we define Y as follows

$$Y = \begin{cases} q_{b-i}, & \text{if } Z \in A_i, \quad 1 \leq i < \alpha, \\ q_a, & \text{otherwise.} \end{cases}$$

Since Q_1 is independent of R_1, R_2, Q_2 , Y is independent of them as well. We have

$$\begin{aligned} P(Y = p_{b-i}) &= P(Z \in A_i) = \\ P(Q_1 = q_{b-i}, H_i = 1, H_j \in \{0, 1\} \text{ for } j \neq i) &= \\ P(Q_1 = q_{b-i})P(H_i = 1)P(H_j \in \{0, 1\} \text{ for } j \neq i) &= \\ P(Q_1 = q_{b-i})P(H_i = 1) &= \\ \frac{1}{\alpha} \frac{\alpha\beta}{(\beta+i)(\beta+i-1)} &= \frac{\beta}{(\beta+i)(\beta+i-1)}. \end{aligned}$$

Also,

$$\begin{aligned} P(Y = p_a) &= 1 - \sum_{i=1}^{\alpha-1} P(Y = p_{b-i}) = \\ 1 - \sum_{i=1}^{\alpha-1} \frac{\beta}{(\beta+i)(\beta+i-1)} &= \\ 1 - \beta \sum_{i=1}^{\alpha-1} \left(\frac{1}{\beta+i-1} - \frac{1}{\beta+i} \right) &= \\ 1 - \beta \left(\frac{1}{\beta} - \frac{1}{\beta+\alpha-1} \right) &= \frac{\beta}{\beta+\alpha-1}. \end{aligned}$$

By definition of A_i , the value of Y is uniquely defined by Q_1 and exactly one H . Therefore, the generation of the whole vector Z is not necessary. Instead, we can calculate Y by the following simple procedure. Once we know the index of Q_1 's value, we generate the corresponding H_i and calculate the value of Y . We can omit the generation of other H s, and therefore we need constant time and memory. \square

Lemma 3.4. *It is possible to generate a zero-one random variable X such that $P(X = 1) = \frac{\alpha}{\beta+\gamma}$. X is independent of R_1, R_2, Q_2 and can be generated using constant time and memory.*

PROOF. Since γ is unknown, it cannot be generated by flipping a coin; a slightly more complicated procedure is required.

Let $Y(Q_1)$ be the random variable from Lemma 3.3. We have

$$P(Y \text{ is not expired}) = \sum_{i=1}^{\gamma} P(Y = q_{b-i}) =$$

$$\sum_{i=1}^{\gamma} \frac{\beta}{(\beta+i)(\beta+i-1)} =$$

$$\beta \sum_{i=1}^{\gamma} \left(\frac{1}{\beta+i-1} - \frac{1}{\beta+i} \right) =$$

$$\beta \left(\frac{1}{\beta} - \frac{1}{\beta+\gamma} \right) = \frac{\gamma}{\beta+\gamma}.$$

Therefore $P(Y \text{ is expired}) = \frac{\beta}{\beta+\gamma}$.

Let S be a zero-one variable, independent of R_1, R_2, Q_2, Y such that

$$P(S = 1) = \frac{\alpha}{\beta}.$$

We put

$$X = \begin{cases} 1, & \text{if } Y \text{ is expired AND } S = 1, \\ 0, & \text{otherwise.} \end{cases}$$

We have

$$P(X = 1) = P(Y \text{ is expired}, S = 1) =$$

$$P(Y \text{ is expired})P(S = 1) = \frac{\beta}{\beta+\gamma} \frac{\alpha}{\beta} = \frac{\alpha}{\beta+\gamma}.$$

Since Y and S are independent of R_1, R_2, Q_2 , X is independent of them as well. Since we can determine if Y is expired within constant time, we need a constant amount of time and memory. \square

Lemma 3.5. *It is possible to construct a random sample V of all non-expired elements using only the data of BS_1, BS_2 and constant time and memory.*

PROOF. Our goal is to generate a random variable V that chooses a non-expired element w.p. $\frac{1}{\beta+\gamma}$. Let X be the random variable generated in Lemma 3.4. We define V as follows.

$$V = \begin{cases} R_1, & R_1 \text{ is not expired AND } X = 1, \\ R_2, & \text{otherwise.} \end{cases}$$

Let p be a non-expired element. If $p \in B_1$, then since X is independent of R_1 , we have

$$P(V = p) = P(R_1 = p, X = 1) =$$

$$P(R_1 = p)P(X = 1) = \frac{1}{\alpha} \frac{\alpha}{\beta+\gamma} = \frac{1}{\beta+\gamma} = \frac{1}{n}.$$

If $p \in B_2$, then

$$P(V = p) =$$

$$(1 - P(R_1 \text{ is not expired})P(X = 1))P(R_2 = p) =$$

$$\left(1 - \frac{\gamma}{\alpha} \frac{\alpha}{\beta+\gamma} \right) \frac{1}{\beta} = \frac{1}{\beta+\gamma} = \frac{1}{n}.$$

\square

Lemma 3.7. *The memory usage of maintaining a random sample within the time-based window has a lower bound $\Omega(\log(n))$.*

PROOF. Let D be a stream with the following property. For timestamp i , $0 \leq i \leq 2t_0$, we have 2^{2t_0-i} elements and for $i > 2t_0$, we have exactly one element per timestamp.

For timestamp $0 \leq i \leq t_0$, the probability of choosing p with $T(p) = i$ at the moment $t_0 + i - 1$ is

$$\frac{2^{2t_0-i}}{\sum_{j=i}^{i+t_0-1} 2^{2t_0-j}} = \frac{2^{2t_0-i}}{2^{t_0-i+1} \sum_{j=0}^{t_0-1} 2^{t_0-j-1}} = \frac{2^{t_0-1}}{\sum_{j=0}^{t_0-1} 2^j} = \frac{2^{t_0-1}}{2^{t_0} - 1} > \frac{1}{2}.$$

Therefore, the expected number of distinct timestamps that will be picked between moments $t_0 - 1$ and $2t_0 - 1$ is at least $\sum_{i=t_0-1}^{2t_0-1} \frac{1}{2} = \frac{t_0+1}{2}$. So, with a positive probability we need to keep in memory at least $\frac{t_0}{2}$ distinct elements at the moment t_0 . The number of active elements n at this moment is at least 2^{t_0} . Therefore the memory usage at this moment is $\Omega(\log n)$, with positive probability. We can conclude that $\log(n)$ is a lower bound for memory usage. \square

Lemma 4.1. *Let $0 < i \leq k$. For any t with more than i active elements, we are able to maintain one of the following:*

1. $\zeta(l(t), N(t) - i)$,
or
2. $BS(y_t, z_t), \zeta(z_t, N(t) - i)$,
where $y_t < l(t) \leq z_t$ and $z_t - y_t \leq N(t) + 1 - i - z_t$ and all random samples of the bucket structures are independent.

PROOF. The proof is the same as in Lemma 3.2, except for cases 1, 2b, 3b. For these cases, when the current window is empty, we keep it empty unless more than i elements are active. We can do this using our auxiliary array. Also, when new elements arrive, some of them may already be expired (if we kept them in the array). We therefore cannot apply the *Incr* procedure for any “new” element. Instead, we should first skip all expired elements and then apply *Incr*. The rest of the proof remains the same. \square

Lemma 4.2. R_{a+1}^{b+1} can be generated using independent R_a^b, R_1^{b+1} samples only.

PROOF. The algorithm is as follows.

$$R_{a+1}^{b+1} = \begin{cases} R_a^b \cup \{b+1\}, & \text{if } R_1^{b+1} \in R_a^b, \\ R_a^b \cup R_1^{b+1}, & \text{otherwise.} \end{cases}$$

Let $X = \{x_1, \dots, x_{a+1}\}$ be a set of points from $[1, b+1]$, such that $x_1 < x_2 < \dots < x_a < x_{a+1}$.

If $x_{a+1} < b+1$, then we have

$$P(R_{a+1}^{b+1} = X) =$$

$$P\left(\bigcup_{j=1}^{a+1} (R_1^{b+1} = x_j \cap R_a^b = X \setminus \{x_j\})\right) =$$

$$\sum_{j=1}^{a+1} P(R_1^{b+1} = x_j) P(R_a^b = X \setminus \{x_j\}) =$$

$$(a+1) \frac{1}{b+1} \frac{1}{\binom{b}{a}} = \frac{1}{\binom{b+1}{a+1}}.$$

Otherwise,

$$P(R_{a+1}^{b+1} = X) =$$

$$P(R_a^b = X \setminus \{b+1\}, R_1^{b+1} \in X) =$$

$$\frac{1}{\binom{b}{a}} \frac{a+1}{b+1} = \frac{1}{\binom{b+1}{a+1}}.$$

\square