Weighted Sampling Without Replacement from Data Streams

Vladimir Braverman^{*} Rafail Ostrovsky^{**} Gregory Vorsanger^{***}

Abstract. Weighted sampling without replacement has proved to be a very important tool in designing new algorithms. Efraimidis and Spirakis (IPL 2006) presented an algorithm for weighted sampling without replacement from data streams. Their algorithm works under the assumption of precise computations over the interval $[0, 1]$. Cohen and Kaplan (VLDB 2008) used similar methods for their bottom-k sketches.

Efraimidis and Spirakis ask as an open question whether using finite precision arithmetic impacts the accuracy of their algorithm. In this paper we show a method to avoid this problem by providing a precise reduction from k-sampling without replacement to k-sampling with replacement. We call the resulting method Cascade Sampling.

1 Introduction

Random sampling is a fundamental tool that has many applications in computer science (see e.g., Motwani and Raghavan [12], Knuth [9], Tille [15], and Olken [13]). Random sampling methods are widely used in data stream processing because of their simplicity and efficiency [14, 8, 7, 6, 10, 11]. In a stream, the size of the domain and the probability of sampling an element both change constantly; this makes the process of sampling non-trivial. We distinguish between sampling *with replacement*, where all samples are independent (and thus can be repeated), and sampling *without replacement*, where repetitions are prohibited.

In particular, weighted sampling without replacement has proven to be a very important tool. In weighted sampling, each element is given a weight, where the probability of an element being selected is based on its weight. In their work Efraimidis and Spirakis [5] presented an algorithm for weighted sampling without replacement. Cohen and Kaplan [3] use similar methods for their bottom-k sketches. While their preliminary implementation yielded promising results,

^{*} Johns Hopkins University, Department of Computer Science. Email: vova@cs.jhu.edu. This material is based upon work supported in part by the National Science Foundation under Grant No. 1447639, the Google Faculty Award and DARPA grant N660001-1-2-4014. Its contents are solely the responsibility of the authors and do not represent the official view of DARPA or the Department of Defense.

^{**} University of California Los Angeles, Department of Computer Science and Department of Mathematics, Email: rafail@cs.ucla.edu.

^{***} Johns Hopkins University, Department of Computer Science. gregvorsanger@jhu.edu. This material is based upon work supported in part by Raytheon BBN Technologies.

Efraimidis and Spirakis [5] state, as the main open problem of the paper, “*However, the question if, and to what extent, the finite precision arithmetic affects the algorithms remains an open problem.*”

In this paper we continue this work and provide a new algorithm to avoid the issue of relying on finite precision arithmetic. With this result we show that precision loss is not required in order to sample without replacement. We accomplish this by providing a precise reduction from k -sampling without replacement to k -sampling with replacement, using a special case of k -sampling with replacement, unit sampling (where $k=1$). Additionally, we believe that in the future our method of expressing different random samples via reduction will provide a tool that allows further translation of other sampling methods into a more effective form for streams.

1.1 Related Work

Due to its fundamental nature, the problem of random sampling has received considerable attention in the last few decades.

In 2005, Vitter [16] presented uniform sampling using a reservoir (with and without replacement) over streams. Further, the question of reductions between sampling methods has been addressed before. For instance, Chaudhuri, Motwani and Narasayya [2] briefly discuss reductions for various sampling methods. Cohen and Kaplan [3] use a “mimicking process” in their papers, which is essentially a reduction from sampling without replacement to sampling with replacement.

Chaudhuri, Motwani and Narasayya [2] use the well-known method of “over-sampling”, i.e. we sample the set independently until k distinct elements are obtained. Clearly, this schema does not introduce any precision loss, since unit sampling is used as a black-box.

Unfortunately, the amount of resources required to determine this information is a function of the weight distribution for the data set, and thus can be arbitrarily large.

In particular, consider the case when there is an element with weight that is overwhelmingly larger than the rest of the population. In this case, the number of repetitions found while sampling with replacement is significantly larger than k .

Probably the first effective non-streaming solution for the weighted sampling without replacement problem was the algorithm of Wong and Easton [17]. It is used by many other algorithms (see Olken [13] for the discussion). For data streams, Efraimidis and Spirakis [5] proposed an algorithm that is based on the “exponent method”. The algorithm requires precise computations of random keys $r^{1/w(p)}$, where $r \sim U[0,1]$. The sample generated is composed of the k elements with maximal keys. Cohen and Kaplan [3] used similar methods as a building block for their bottom- k sketches. The bottom- k sketch is an effective construction that has been extensively used for various applications including approximations of aggregative queries over data streams. As Cohen and Kaplan [3] show, these methods are very effective in practical applications and are superior to the sketches that are based on sampling with replacement.

While effective in practice, the algorithms of Efraimidis and Spirakis and Cohen and Kaplan introduce a loss of accuracy, since their techniques require additional floating point arithmetic operations.

1.2 Results

In this paper we show that the tradeoff between precision and performance is not a necessary property of sampling without replacement from data streams and construct a precise streaming reduction from k -sampling without replacement to k -sampling with replacement. This result provides a practical improvement to the algorithms of Efraimidis and Spirakis in cases where high accuracy is required.

Our method yields a surprisingly simple algorithm, given the importance of sampling without replacement and the existence of many previous methods. We call this algorithm Cascade Sampling. In particular, when used with the algorithm from [2] Cascade Sampling requires $O(k)$ memory, constant time per element and the same precision as in [2].

1.3 Intuition

Let A be *any* algorithm that maintains a unit weighted sample from stream D . Similarly to the over-sampling method, we maintain instances of A . Namely, we maintain k instances $A_1, \dots, A_k$. However, we introduce the idea of *stream modification*. That is, instead of applying A independently and symmetrically on D , we apply A_i on the modified stream D_i that does not contain samples of A_j for $j < i$. In particular, A_i may process its input elements in an order *different from the order of their arrival in D* . This simple but novel idea is sufficient to solve the problem. In particular, we can claim that the input of A_i is a random set that precisely matches the definition of weighted sampling without replacement. Since we use A as a black box with only a constant number of auxiliary variables, specifically pointers, the resulting schema is a precise reduction.

2 Definitions

An important building block of our algorithm is the concept of a unit sample, that is, the ability to sample a single element from a set.

Definition 1. Let S be a finite set of elements and let w be a non negative function $w : S \rightarrow R$. A random element X_S with values from S is a **unit weighted random sample** if, for any $a \in S$, $P(X_S = a) = \frac{w(a)}{w(S)}$. Here $w(S) = \sum_{a \in S} w(a)$.

For an algorithm instantiating weighted unit sampling we provide Black-Box WR2 from [2]. Black-Box WR2 is a unit sample when $r = 1$.

Algorithm 1 Black-Box WR2: Algorithm for Weighted Unit Sampling

1. $W \leftarrow 0$.
 2. Initialize reservoir with length $r = 1$, λ_0 .
 3. For each tuple t in stream:
 - (a) Get next tuple t with weight $w(t)$
 - (b) $W \leftarrow W + w(t)$
 - (c) Set $\lambda_0 = t$ with prob. $\frac{w(t)}{W}$
 4. Return λ_0
-

Definition 2. A **data stream** is an ordered, set of elements, $p_1, p_2, \dots, p_n$, that can be observed only once. An algorithm A is a **streaming sampling algorithm** if A outputs a sample using a single pass over the data set.

Definition 3. A set $X = \{X_1, \dots, X_k\}$ is called a **k-sample with replacement** from S if $X_1, \dots, X_k$ are independent random unit samples from S .

Another fundamental sampling method is *weighted sampling without replacement*.

Definition 4. Let S be a finite set such that $|S| \geq k$. An ordered set $X = \{X_1, \dots, X_k\}$ is called a **k-sample without replacement** from S , $|S| \geq k$ if X_1 is a weighted unit sample from S and for any $j > 1$, X_j is a weighted unit sample from $S \setminus \{X_1, \dots, X_{j-1}\}$.

Definition 5. We say that there exists an a reduction from a k -sampling to a unit sampling if for any unit sampling algorithm Λ there exists a k -sampling algorithm $\Upsilon = \Upsilon(\Lambda)$ that uses Λ as a black-box. We say that the reduction is precise if for any Λ that requires memory m and time t :

1. $\Upsilon(\Lambda)$ requires $O(km)$ memory and $O(kt)$ time.
2. $\Upsilon(\Lambda)$ only uses comparisons (in addition to using Λ as a black box).

In other words, $\Upsilon(\Lambda)$ does not introduce any precision loss.

There exists a (trivial) precise reduction from weighted sampling with replacement to unit sampling. In this paper we give the first precise streaming reduction for weighted sampling without replacement to unit sampling.

3 Cascade Sampling

Let S be a finite set such that $|S| \geq k$ and let $a \notin S$. Denote $T = S \cup \{a\}$, and let $w : T \mapsto R^+$ be a function. Let $\{X_1, \dots, X_k\}$ be a k -sample without replacement from S with respect to w . Define an ordered sequence $\{Y_1, \dots, Y_k\}^1$ as follows:

$$Y_1 = \begin{cases} a, & \text{w.p. } \frac{w(a)}{w(T)}; \\ X_1, & \text{otherwise.} \end{cases} \quad (1)$$

¹ Here the additional randomness is independent.

For $i \geq 1$ define²:

$$L_i = \{X_1, \dots, X_i, a\} \setminus \{Y_1, \dots, Y_i\}. \quad (2)$$

We will show that $|L_i| = 1$; assuming that, let Z_i be the single element from L_i , i.e., $L_i = \{Z_i\}$. Put $U_i = T \setminus \{Y_1, \dots, Y_i\}$. Define

$$Y_{i+1} = \begin{cases} Z_i, & \text{w.p. } \frac{w(Z_i)}{w(U_i)}; \\ X_{i+1}, & \text{otherwise.} \end{cases} \quad (3)$$

Lemma 1. *For all $i = 1, \dots, k$ the ordered set $\{Y_1, \dots, Y_i\}$ is an i -sample without replacement from T with respect to w .*

Proof. We prove the lemma by induction on i . For $i = 1$ the statement follows from direct computation and definitions. Assuming that the lemma is correct for i we need to prove that

$$Y_{i+1} \in T \setminus \{Y_1, \dots, Y_i\}, \quad (4)$$

and for any $b \in U_i$:

$$P(Y_{i+1} = b) = \frac{w(b)}{w(U_i)}. \quad (5)$$

To show (4) observe that $\{Y_1, \dots, Y_i\} \subseteq \{X_1, \dots, X_i, a\}$ and $Y_{i+1} \in \{X_{i+1}, Z_i\}$. By definition $X_{i+1} \notin \{X_1, \dots, X_i, a\}$ and $Z_i \notin \{Y_1, \dots, Y_i\}$.

To show (5) fix $\{X_1, \dots, X_i\}$ and $\{Y_1, \dots, Y_i\}$; it follows that Z_i is fixed as well. Denote $V_i = U_i \setminus \{Z_i\}$ and $H_i = S \setminus \{X_1, \dots, X_i\}$; it follows that $H_i = V_i$. For any fixed $b \in V_i$ we have

$$P(Y_{i+1} = b) = P(X_{i+1} = b) \frac{w(V_i)}{w(U_i)} = \frac{w(b)}{w(H_i)} \frac{w(V_i)}{w(U_i)} = \frac{w(b)}{w(U_i)}.$$

The case $b = Z_{i-1}$ is similar.

4 Precise Reduction and Resulting Algorithm

Let A be an algorithm that maintains a unit weighted sample from D . The algorithm from [2] is an example of A but our reduction works with *any* algorithm for unit weighted sampling. We construct an algorithm $\mathcal{Y} = \mathcal{Y}(A)$ such that $\mathcal{Y}$ maintains a k -sample without replacement. Specifically, we maintain k instances of A : $A_1, \dots, A_k$ such that the input of A_i is a random substream of D that is selected in a special way. We denote the input stream for A_i as D_i . Let X_i be the sample produced by A_i . The critical observation is that our algorithm maintains the following invariant: at any moment $D_i = D \setminus \{X_1, \dots, X_{i-1}\}$. Thus, by definition, the weighted sample from D_i is the i -th weighted sample from D when the samples are without replacement.

² Here $\setminus$ denotes the set difference, i.e. $A \setminus B = \{x : x \in A, x \notin B\}$.

Theorem 1. *Algorithm $\Upsilon = \Upsilon(\Lambda)$ maintains a weighted k -sample without replacement from D . If Λ requires space $O(g)$ and time per element $O(h)$, then Υ requires $O(kg)$ space and $O(kh)$ time respectfully. Thus, there exists a precise reduction from k -sampling without replacement to a unit sampling.*

Proof. Follows from the description of the algorithm (See Algorithm 2) and Lemma 1.

Algorithm 2 Cascade Sampling

Input: Data Stream $D = \{p_1, \dots, p_n\}$,

Λ is an algorithm that maintains a unit weighted sample from D ,

$\Lambda_1, \dots, \Lambda_k$ are independent instances of Λ

Output: Weighted k -Sample Without Replacement $\{Y_1, \dots, Y_k\}$

1. For $j = 1, 2, \dots, n$
 - (a) $new = p_j$
 - (b) For $i = 1, \dots, \min\{j, k\}$
 - i. If $(i < j)$ then set $previous = Y_i$ (where Y_i the current output of Λ_i).
 - ii. Feed Λ_i with new
 - iii. If Y_i changes its value to new , then set $new = previous$.
 2. Output $\{Y_1, \dots, Y_k\}$
-

Algorithm 2 provides a solution to the weighted k -Sampling without replacement problem. To better demonstrate the algorithm, we show an example of updating a single unit sample inside of loop (b) in Figure 1. In this example, unit sample λ_1 has currently sampled element a and unit sample λ_2 has currently sampled element b , where a and b are elements that appeared previously in the stream.

4.1 Discussion

There are several directions in which our algorithm can be improved. In particular, run time dependent on the number of samples is one issue for practical datasets with large k . We believe this can be improved by combining several sampling steps into a single step which will be useful for the cases when the element will not be sampled into any of the substreams. This will often be the cases with elements with small weights. Specifically, we ask if it is possible to reduce the total running time from $O(nk)$ to $O(n \log k)$.

Another interesting direction is applying this algorithm to weighted random sampling with a bounded number of replacements as shown in [4]. Finally, this method may also be interesting when applied to the Sliding Window Model [1] and Streams with Deletions [6].

We thank our anonymous reviewers for their helpful suggestions, particularly for suggesting interesting open problems for discussion.

Fig. 1. Updating a Unit Sample

References

1. Vladimir Braverman and Rafail Ostrovsky. Smooth histograms for sliding windows. In *Proceedings of the 48th Annual IEEE Symposium on Foundations of Computer Science*, FOCS '07, pages 283–293, Washington, DC, USA, 2007. IEEE Computer Society.
2. Surajit Chaudhuri, Rajeev Motwani, and Vivek Narasayya. On random sampling over joins. In *Proceedings of the 1999 ACM SIGMOD international conference on Management of data*, SIGMOD '99, pages 263–274, New York, NY, USA, 1999. ACM.
3. Edith Cohen and Haim Kaplan. Tighter estimation using bottom k sketches. *Proc. VLDB Endow.*, 1(1):213–224, August 2008.
4. Pavlos S Efrimidis. Weighted random sampling over data streams. *arXiv preprint arXiv:1012.0256*, 2010.
5. Pavlos S. Efrimidis and Paul G. Spirakis. Weighted random sampling with a reservoir. *Inf. Process. Lett.*, 97(5):181–185, March 2006.
6. Gereon Frahling, Piotr Indyk, and Christian Sohler. Sampling in dynamic data streams and applications. In *Proceedings of the twenty-first annual symposium on Computational geometry*, SCG '05, pages 142–149, New York, NY, USA, 2005. ACM.
7. Phillip B. Gibbons and Yossi Matias. New sampling-based summary statistics for improving approximate query answers. In *Proceedings of the 1998 ACM SIGMOD international conference on Management of data*, SIGMOD '98, pages 331–342, New York, NY, USA, 1998. ACM.

8. Theodore Johnson, S. Muthukrishnan, and Irina Rozenbaum. Sampling algorithms in a stream operator. In *Proceedings of the 2005 ACM SIGMOD international conference on Management of data*, SIGMOD '05, pages 1–12, New York, NY, USA, 2005. ACM.
9. Donald E. Knuth. *The art of computer programming, volume 1 (3rd ed.): fundamental algorithms*. Addison Wesley Longman Publishing Co., Inc., Redwood City, CA, USA, 1997.
10. M. Kolonko and D. Wäsch. Sequential reservoir sampling with a nonuniform distribution. *ACM Trans. Math. Softw.*, 32(2):257–273, June 2006.
11. Kim-Hung Li. Reservoir-sampling algorithms of time complexity $o(n(1 + \log(n/n)))$. *ACM Trans. Math. Softw.*, 20(4):481–493, December 1994.
12. R. Motwani and P. Raghavan. *Randomized algorithms*. Cambridge University Press, New York, NY, 1995.
13. Frank Olken and Doron Rotem. Random sampling from databases: a survey. *Statistics and Computing*, 5(1):25–42, 1995.
14. Mario Szegedy. The DLT priority sampling is essentially optimal. In *Proceedings of the thirty-eighth annual ACM symposium on Theory of computing*, STOC '06, pages 150–158, New York, NY, USA, 2006. ACM.
15. Y. Tille. *Sampling Algorithms*. Springer, Verlag, 2006.
16. J. S. Vitter. *Random sampling with a reservoir*. ACM Transactions on Mathematical Software, v.11 n.1, pp.37–57, 1985.
17. C. K. Wong and M. C. Easton. An efficient method for weighted sampling without replacement. *SIAM J. Comput.*, 9(1):111–113, 1980.