

# KLL<sup>±</sup>: Approximate Quantile Sketches over Dynamic Datasets

Fuheng Zhao    Sujaya Maiyya    Ryan Wiener    Divyakant Agrawal    Amr El Abbadi  
Department of Computer Science, University of California Santa Barbara  
{fuheng\_zhao, sujaya\_maiyya, ryanlwiener, agrawal, amr}@cs.ucsb.edu

## ABSTRACT

Recently the long standing problem of optimal construction of quantile sketches was resolved by **K**arnin, **L**ang, and **L**iberty using the KLL sketch (FOCS 2016). The algorithm for KLL is restricted to online insert operations and no delete operations. For many real-world applications, it is necessary to support delete operations. When the data set is updated dynamically, i.e., when data elements are inserted and deleted, the quantile sketch should reflect the changes. In this paper, we propose  $KLL^\pm$  – the first quantile approximation algorithm to operate in the *bounded deletion* model to account for both inserts and deletes in a given data stream.  $KLL^\pm$  extends the functionality of KLL sketches to support arbitrary updates with small space overhead. The space bound for  $KLL^\pm$  is  $O(\alpha^{1.5}/\epsilon \log^2 \log(1/\epsilon\delta))$ , where  $\epsilon$  and  $\delta$  are constants that determine precision and failure probability, and  $\alpha$  bounds the number of deletions with respect to insert operations. The experimental evaluation of  $KLL^\pm$  highlights that with minimal space overhead,  $KLL^\pm$  achieves comparable accuracy in quantile approximation to KLL.

## 1. INTRODUCTION

With the rise of big data in companies such as Google, Amazon, and Facebook, managing hundreds of terabytes to petabytes of data has become a necessity for day to day operations. To make use of this data, it is crucial to develop a deeper understanding of the underlying distributions of the data in real datasets. In particular, techniques such as quantile approximations, a non-parametric representation, are widely used to characterize data distributions. For large amounts of data, one-pass<sup>1</sup> algorithms are desirable, and many well-known data sketches are based on one-pass algorithms. For instance, the HyperLogLog [10] sketch<sup>2</sup> is a one-pass algorithm for the *count-distinct* problem; the Bloom

<sup>1</sup>Where input data is read only once.

<sup>2</sup>The term *sketch* refers to the algorithm and data structures that can extract valuable information through one pass of the data.

This work is licensed under the Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License. To view a copy of this license, visit <http://creativecommons.org/licenses/by-nc-nd/4.0/>. For any use beyond those covered by this license, obtain permission by emailing [info@vldb.org](mailto:info@vldb.org). Copyright is held by the owner/author(s). Publication rights licensed to the VLDB Endowment.

*Proceedings of the VLDB Endowment*, Vol. 12, No. xxx

ISSN 2150-8097.

DOI: <https://doi.org/10.14778/xxxxxx.xxxxxx>

Filter [3] data sketch is a one-pass algorithm for the *set membership* problem. Several one-pass quantile approximation algorithms [20, 11, 8, 19, 26, 27, 1, 12] have been proposed to guarantee high precision while using small memory footprints. Since one-pass algorithms only read each element from the input only once, these one-pass algorithms are not only useful for large databases, but are also naturally applicable in the *streaming* data settings.

Approximate quantile problems are developed for a variety of settings. For example, approximate quantiles are considered in the streaming data settings [19, 20], sliding window settings [2], and distributed settings [26]. Moreover, approximate quantile sketches operate mainly within the *insertion-only* model and the *turnstile* model. The insertion-only model, also referred as the cash register model, consists of a stream of only insert operations; whereas the turnstile model consists of a stream of insert and delete operations such that deletes are performed on previously inserted items [27]. Quantile sketches such as the GK sketch [12] and Q-digest [26] operate in the insertion-only model; and quantile sketches such as Range Subset Sum [11], Dyadic Count-Min [8], and Dyadic Count Sketch [27] operate in the turnstile model. Turnstile model quantile summaries have significant usage in database and networking applications, and are used in most Database Management Systems (DBMSs) which monitor and maintain a variety of order statistics such as quantiles over the contents of database relations [24]; in value range partitioning for parallel database [25]; and in financial services [16]. Sketches proposed for the turnstile model assume a fixed universe in order to tolerate an arbitrary number of deletions and thus incur higher space and update complexity compared to data-driven sketch. The additional challenge of tolerating an arbitrary amount of deletion for the turnstile model is arguably an infeasible task for data-driven sketches which stores a subset of items chosen from the input. Several researchers [11, 27] have noted that if one first inserts  $n$  elements and then deletes all but one element, the data-driven sketch has no information about which element will survive since the data-driven sketch only keeps a subset of items. Wang et al. [27] conjecture the impossibility of any data-driven sketch to support the turnstile model where all elements can be deleted. Therefore, we focus on developing a data-driven sketch in the context of the *bounded deletion* model.

Recently, Jayaram et al. [18] observed that in practice many turnstile models only incur a fraction of deletions and thus proposed an intermediate model, called the *bounded deletions* model, which naturally lies in between the insertion-

only stream model and the turnstile model. Jayaram et al. [18] use parameter  $\alpha$  to denote different degrees of deletions such that at most  $1-1/\alpha$  of prior insertions are deleted; when  $\alpha = 1$  the input dataset becomes the insertion-only model. The authors [18] show significant space improvements for many fundamental streaming problems within the bounded deletion model over the turnstile model, such as heavy hitters and inner product estimation, and identify the bounded deletion model as being particularly useful for applications such as computing differences between network traffic patterns, identifying new trending searches,  $L_0$  estimation of moving sensors (monitoring wildlife or water flow pattern), and completing synchronization in database analytical context.

Furthermore the bounded deletion model is important for maintaining approximate quantiles in real-world applications such as summarizing product sales in electronic commerce platforms and ranking in standardized testings. Many companies need accurate quantile information to identify current status and forecast demands of product life cycles [14]. After customers purchase some products, a certain percentage of the customers may return the product and submit refund requests, hence the quantile summaries should reflect these changes. However, it is highly unlikely that the majority of these customers will return their purchases and in most cases a bounded deletion model can be assumed. In the context of standardized testing such as SAT, GRE, and LSAT, quantiles are seen as a more descriptive measurement compared to the test scores, since the difficulty of the exam varies. In these standardized testing settings, students may request to regrade their exam only once to verify any machine errors in answer sheet scanning or human errors in grading the essay. The updated grades may change the underlying score distribution. Thus, a quantile summary with  $\alpha = 2$ , is sufficient to reflect all changes even if all students require a regrade<sup>3</sup>. In Section 4.3, we also show how to leverage  $KLL^\pm$  with  $\alpha = 2$  to maintain fixed-size sliding window quantile approximation over datasets.

In this paper, we extend the original optimal quantile sketch  $KLL$  [19], which operates in the insertion-only model, and present  $KLL^\pm$  quantile sketch that supports bounded deletions in which the total number of deleted items are less than a certain threshold of the total inserted items, and the deleted items are previously inserted into the data set. The  $KLL$  algorithm is online updatable for insert operations, but not for delete operations. Without delete operations, a  $KLL$  sketch cannot be maintained for a dynamic dataset that is constantly updated as an expensive scan operation will be needed to recompute the  $KLL$ -sketch, when deletions are applied. Our  $KLL^\pm$  expands the functionality of the original  $KLL$  sketch by supporting bounded deletions; and if the administrator of the big data, knows a priori that deletions are not insensitively large compared to the insertions, such as some corrections or adjustments on previously inserted items, then  $KLL^\pm$  can approximate quantiles with small space and high accuracy. In summary, the main contributions in this paper are: (i) Presenting the  $KLL^\pm$  sketch, the first data-driven quantile sketch that operates in the bounded deletion model; (ii) Providing thorough mathematical proofs and analysis to guarantee the correctness of

$KLL^\pm$  sketches; (iii) Applying the  $KLL^\pm$  sketches to maintain approximate quantile estimations in the fixed-size sliding window over datasets; (iv) Evaluating  $KLL^\pm$  and comparing it with  $KLL$ , two parallel  $KLL$  strategy, and DCS sketch [27] through various experiments; and (v) Incorporating several optimization strategies [17] proposed for the original  $KLL$  sketch [19], such as implementing the algorithm with fixed memory size.

This paper is structured as follows: Section 2 discusses background of quantile sketches and gives an overview of previous algorithms. Section 3 introduces the  $KLL^\pm$  quantile sketch for datasets with an incoming input stream in the bounded deletion model, and along with the proposed algorithm for compaction referred to as Conditional Compaction algorithm. Section 4 analyzes the space complexity of our sketch, and demonstrates a potential use case of  $KLL^\pm$  to maintain fixed-size sliding window quantile approximation over datasets. Section 5 presents the results of an evaluation conducted using synthetic and real world datasets and compares  $KLL^\pm$  with the state of the art  $KLL$  sketch (insertion-only) and DCS sketch (turnstile) [27]. Finally, Section 6 summarizes our contributions and concludes this work.

## 2. BACKGROUND

Given a multiset of  $n$  elements  $S = \{s_1, \dots, s_n\}$ , the rank of an element  $s_i \in S$  is the number of elements in  $S$  that are less than or equal to  $s_i$ , and  $R(s_i)$  is the function specific to a set  $S$  that takes as input an element  $s_i$  and returns its rank. The quantile of an element  $s_i \in S$  is defined as  $R(s_i)/n$ . Equivalently, the  $\phi$ -quantile of set  $S$ ,  $Q(\phi)$  returns the element  $s_i$  such that  $Q(R(s_i)/n) = s_i$ . Typically, quantile is represented as a fraction  $\phi \in (0, 1]$ . The relationship between the rank and quantile of an element can be represented as:  $R(s_i) = \phi \cdot n$ , where  $\phi$  is a non-zero fraction less than or equal to one and  $n$  is the number of elements in the dataset. Equivalently,  $\phi$  elements are less than or equal to  $s_i$  and  $(1 - \phi) \times n$  elements are greater than  $s_i$ . The most familiar quantile value is  $\phi = 0.5$  which is referred to as the *median* of the dataset.

### 2.1 Deterministic and Randomized algorithms

Computing the true median value, especially for large data sets, is memory intensive. Munro and Paterson [23] formally proved that to find the true median of a set of size  $n$  with  $p$  sequential passes of the input requires at least  $\Omega(n^{1/p})$  memory. Thus, to determine the true median using a one-pass algorithm requires memory linear to the size of the set. With limited memory and for large data sets, calculating the exact quantiles is infeasible. An alternate and more practical approach to the problem is to approximate the quantiles, represented as  $\epsilon$  approximation  $\phi$  quantile, where  $\epsilon$  is the precision value.

The *deterministic*  $\epsilon$  approximation  $\phi$  quantile algorithms take as input a quantile query  $\phi$  and a precision value  $\epsilon$  and output an answer  $x$  such that  $x$ 's quantile must be in the range  $[\phi - \epsilon, \phi + \epsilon]$ . An alternative approach proposes a family of *randomized* algorithms where the output answer  $x$  is said to be within the  $[\phi - \epsilon, \phi + \epsilon]$  range with a high probability [17, 19, 20]. These algorithms provide guarantees by bounding the failure probability to at most  $\delta$  such that the user has  $1 - \delta$  confidence that the sketch's output is  $\epsilon$  approximation.

<sup>3</sup><https://collegereadiness.collegeboard.org/sat/scores/verifying-scores>

Sketch	Space	Update Time	Randomization	Model	Framework
GK Sketch [12]	$O(\frac{1}{\epsilon} \log(\epsilon n))$	$O(\log \frac{1}{\epsilon} + \log \log(\epsilon n))$	Deterministic	Insertion-Only	Data-Driven
Q-digest [26]	$O(\frac{1}{\epsilon} \log U)$	$O(\log \frac{1}{\epsilon} + \log \log U)$	Deterministic	Insertion-Only	Universe-Driven
MRL99 [21]	$O(\frac{1}{\epsilon} \log^2(\frac{1}{\epsilon}))$	$O(\log(\frac{1}{\epsilon}))$	Randomized	Insertion-Only	Data-Driven
Mergeable KLL [19]	$O(\frac{1}{\epsilon} \log^2 \log(\frac{1}{\epsilon}))$	$O(\frac{1}{\epsilon})$ $O(\log(\frac{1}{\epsilon}))$ [17]	Randomized	Insertion-Only	Data-Driven
RSS [11]	$O(\frac{1}{\epsilon} \log^2 U \log \frac{\log U}{\epsilon})$	$O(\frac{1}{\epsilon} \log^2 U \log \frac{\log U}{\epsilon})$	Randomized	Turnstile	Universe-Driven
DCM [8]	$O(\frac{1}{\epsilon} \log^2 U \log \frac{\log U}{\epsilon})$	$O(\log U \log \frac{\log U}{\epsilon})$	Randomized	Turnstile	Universe-Driven
DCS [27]	$O(\frac{1}{\epsilon} \log^{1.5} U \log^{1.5} \frac{\log U}{\epsilon})$	$O(\log U \log(\frac{\log U}{\epsilon}))$	Randomized	Turnstile	Universe-Driven
Mergeable KLL $^\pm$	$O(\frac{\alpha^{1.5}}{\epsilon} \log^2 \log \frac{1}{\epsilon})$ s.t. $D \leq (1 - 1/\alpha)I$	$O(\frac{\alpha^{1.5}}{\epsilon})$ $O(\log(\frac{\alpha^{1.5}}{\epsilon}))$ [17]	Randomized	Bounded Deletions	Data-Driven

Table 1: Comparison between different quantile sketches

## 2.2 Quantile Sketch: Insertion-only Model

Greenwald and Khanna [12] developed an algorithm, referred to as the GK sketch that tracks a sorted subset of elements in the input data stream such that these elements provide lower and upper bounds for each quantile individually instead of maintaining a single bound over all quantiles. The GK sketch is deterministic and requires  $O(1/\epsilon \log(\epsilon N))$  space in the worst case. It is conjectured [1], however, that the GK sketch is not fully *mergeable* – the property of merging two sketches on two different datasets is to create one combined or merged sketch that can then be used for quantile computation over the union of the two datasets. Mergeability allows users to compute sketches over multiple data partitions independently and to combine them in parallel to compute the summaries. The merged sketch should be as accurate as a single sketch over the entire data set. This is a crucial requirement in distributed settings, where data is often stored distributedly across different machines. Creating sketches independently and then merging them avoids the communication costs and large latency of transferring large amounts of data to a central repository. Furthermore, this also makes the quantile computation highly scalable in the context of very large datasets.

Karnin et al. [19] extended the GK sketch and presented the KLL sketch, an asymptotically optimal but non-mergeable sketch with  $O(1/\epsilon \log \log(1/\delta\epsilon))$  space. Karnin et al. also presented a mergeable KLL sketch with  $O(1/\epsilon \log^2 \log(1/\delta\epsilon))$  space bounds. The core building block for the KLL quantile sketch is called the compactor, first introduced in [20]. KLL can be seen as an array of  $H$  compactors as depicted in Figure 1. One of the main contributions by Karnin et al. [19] is to obtain the optimal sketch size by having different capacity compactors at different heights and exponentially decreasing the capacity of compactors at lower heights.

When a *compactor* $[h]$ 's size reaches its capacity, i.e.,  $k_h = c_h$ , the compactor performs a compaction process, in which it pushes  $k_h/2$  elements from *compactor* $[h]$  to *compactor* $[h+1]$ . Hence, these  $k_h/2$  elements update their weight to  $2 \cdot w_h$ , i.e.,  $2^h$ . The compaction process introduces error since each compaction pushes only half of the elements to the next level. For example, consider a compaction of two elements A, B with weight  $w_h = 1$  in which  $\text{rank}(A)=1$  and  $\text{rank}(B)=2$ . A compaction pushes either A or B, and the compacted element's weight becomes  $w_{h+1} = 2$ . If element A is chosen then the sketch loses the information about element B and believes element A appeared twice, hence  $\text{rank}(A)=2$  and  $\text{rank}(B)=2$  in which case we introduced  $+1$




Figure 1: An illustration of a set of compactors. As the height decreases from  $H$  to 1, the capacity of the compactor decreases from  $k$  to  $c^{H-1}k$ .  $k$  is a constant given by the user, and  $c$  is a constant between 0.5 and 1.

( $+w_H$ ) error for rank A. If element B is chosen then the sketch loses the information about element A and believes element B appeared twice, hence  $\text{rank}(A)=0$  and  $\text{rank}(B)=2$  in which case we introduced  $-1$  ( $-w_H$ ) error for rank A. In both cases, we introduce no error for rank B. Therefore the rank estimation before and after a compaction process differs by at most  $w_h$  as shown in Figure 2. Agarwal et al. [1] suggested that by removing the odd or even indexed elements with equal probability, the expected error becomes zero. This ensures that the total error is bounded by  $\delta$  for  $O(1/\epsilon)$  quantiles.

In summary, Karnin et al. [19] (i) exponentially decrease the compactor capacity; (ii) replace compactors of capacity 2 with a sampler, which randomly selects one element from  $2^{w_H - \log_{1/c} k}$  elements; (iii) keep the size of the top  $O(\log \log 1/\delta)$  compactors fixed (similar to the MRL sketch [20]); (iv) replace the top  $O(\log \log 1/\delta)$  compactors with a GK sketch [12]. Recently, Ivkin et al. [17] extended the theoretical development of Karnin et al. [19] to achieve practical improvements for implementing KLL sketches. In summary the contributions are: (i) Implementing the algorithm with memory limit parameter, denoted by  $k$ , as opposed to  $\delta$  (failure probability) and  $\epsilon$  (precision) parameters; (ii) extending the functionality to handle weighted data streams; and (iii) various optimization strategies to reduce memory footprint and update time. These optimization strategies are orthogonal to our work, and in Appedix F we show how to integrate our work with the Sweep Compactor [17] to

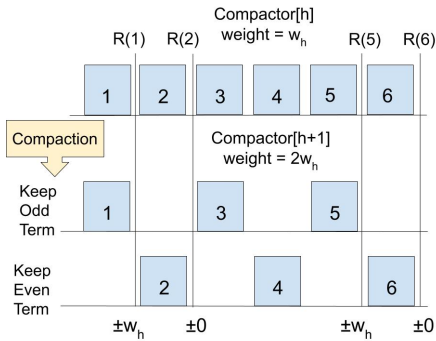


Figure 2:  $Compactor[h]$  reaches its capacity of 6 and performs a compaction. If the odd indexed terms are chosen then the compaction contributes  $+w_h$  error to  $R(1)$ ,  $R(3)$ , and  $R(5)$ . If the even indexed terms are chosen then the compaction contributes  $-w_h$  error to  $R(1)$ ,  $R(3)$ , and  $R(5)$ . In both cases the compaction contributes no error to  $R(2)$ ,  $R(4)$ , and  $R(6)$ .

improve the worst case update time.

### 2.3 Quantile Sketch: Turnstile Model

In the turnstile model, the input data stream consists of both insertions and deletions and a deletion cannot delete an element that does not exist. This is also known as the strict definition for the turnstile model. Gilbert et al. [11] were the first to propose quantile sketches for the turnstile model and introduced quantile sketches to support both insertions and deletions in database management systems (DBMSs). The authors [11] presented the novel Random Subset Sums (RSS) sketch, which breaks down the universe  $U$  into dyadic intervals and maintains the frequency estimations of elements for each layer, using total space of  $O(\frac{1}{\epsilon^2} \log^2 U \log \frac{\log U}{\epsilon})$ , with update time  $O(\frac{1}{\epsilon^2} \log^2 U \log(\frac{\log U}{\epsilon}))$ . The dyadic structure decomposes the universe into  $\log U$  layers such that in each  $i$  layer the universe is partitioned into  $U/2^i$  intervals of size  $2^i$ . The top most layer represents the interval:  $[1, U]$ , the second top most layer represents two intervals:  $[1, U/2]$ ,  $[U/2+1, U]$ , the third top most layer represents four intervals:  $[1, U/4]$ ,  $[U/4+1, U/2]$ ,  $[U/2+1, 3U/4]$ ,  $[3U/4+1, U]$  and so on until the bottom layer representing all the elements in  $U$ . Thus, to find the rank estimate of an element  $x \in U$ , we can decompose the interval  $[1, x]$  into  $\log x$  number of disjoint dyadic intervals, and query each of these intervals to the frequency estimation sketch of their corresponding layers and then sum all the estimations to obtain the rank estimation.

Later Cormode et al. [8] proposed the DCM (Dyadic Count Min) sketch with overall space  $O(\frac{1}{\epsilon} \log^2 U \log(\frac{\log U}{\epsilon}))$  and update time  $O(\log U \log(\frac{\log U}{\epsilon}))$ . DCM sketch replaced the frequency estimation sketch in each layer with a Count Min Sketch, in which Count Min sketch uses small space to output the frequency estimation of each element with an additive factor of  $\epsilon$  with high probability. Recently, Wang et al. [27] proposed DCS (Dyadic Count Sketch) sketch which replaced the Count Min sketch with Count sketch [5]. Count sketch is similar to Count Min sketch as it has small space overhead to output frequency estimation for each element with an additive factor of  $\epsilon$  in high probability. While Count Min sketch's estimation is biased toward overestima-

tion, Count sketch gives an unbiased frequency estimation by reporting the median. Wang et al. [27] point out that the property of unbiased frequency estimation from the Count sketch is appealing to the quantile problem. Estimations, by sketches for each dyadic layers, may give positive errors or negative errors, and these errors may cancel out each other. Thus, DCS sketches further improve the space bound to  $O(\frac{1}{\epsilon} \log^{1.5} U \log^{1.5}(\frac{\log U}{\epsilon}))$  with the same update time complexity as DCM sketches.

These three data sketches (RSS, DCM, and DCS) guarantee that deleting a previously inserted element has no impact on the space or accuracy of the sketch. In fact, Gilbert et al. [11] compared the RSS sketch with 2-parallel GK sketches (one for insertions and one for deletions), and found that when deletions are relatively small, the simple 2-parallel GK sketches exceed the accuracy of RSS by two orders of magnitude. When the number of deletions are significantly large, RSS is more accurate than the 2-parallel GK sketches.

During the past three decades of research, various quantile summaries have been developed through a variety of models and frameworks. In Table 1, we compare the difference and similarity among several different sketches. Quantile sketches algorithms can be categorized using data-driven or universe-driven framework [7]. Algorithms in the data-driven framework keep a subset of items that appeared in the stream and maintain their statistics. On the other hand, algorithms in the universe-driven framework maintain attributes over the universe and have space and update time bounds that depend on the size of the universe  $|U|$ . For instance, Shrivastava et al. [26] proposed a novel quantile sketch for the insertion-only model and universe-driven framework with  $O(1/\epsilon \log(|U|))$  space. Our algorithm uses the the bounded deletion model and data-driven framework, more precisely the comparison framework<sup>4</sup>. The benefit for sketches in the data-driven framework is that these sketches make no assumptions on the universe size, hence can handle dataset with attributes involving variable-length strings while their space and update time are independent from the universe size; Recent surveys by Wang et al. [27], Greenwald et al. [13], Chen et al. [6], and Cormode et al. [9] provide more comprehensive background on quantile sketches.

## 3. $KLL^\pm$ QUANTILE SKETCH

The challenges for supporting arbitrary number of delete operations in the turnstile model motivated researchers [11, 8, 27] to explore universe-driven algorithms which necessitate the sketch size and update time to be dependent on the size of the universe, which can be quite large. Our main goal is to extend data-driven algorithms to the bounded deletion model where the input consists of  $I$  insertions and at most  $(1 - 1/\alpha) \times I$  deletions. We propose the  $KLL^\pm$  Quantile Sketch, a generalization of the KLL sketch to maintain quantile information in the bounded deletion model using as small of a memory footprint as possible.

### 3.1 Basic Structure

Similar to MRL [20] and KLL [19],  $KLL^\pm$  can be seen as an array of  $H$  compactors where  $H$  denote the total height. Each compactor is identified by its height, denoted as  $compactor[h]$ . The topmost compactor and the bottom compactor have height  $H$  and 1, respectively. Each  $compactor[h]$

<sup>4</sup>Only comparisons are applied on elements.

has a limited capacity  $c_h = c^{H-h}k$  where  $k$  is the capacity of *compactor*[ $H$ ]. Each *compactor*[ $h$ ] contains  $k_h$  elements such that  $k_h < c_h$  and each element in *compactor*[ $h$ ] has weight  $w_h = 2^{h-1}$ . Once a compactor is full,  $k_h = c_h$ , the compactor will go through a compaction to free spaces for new elements.

### 3.2 Differentiating Inserts & Deletes

In the insertion-only model, the incoming data stream consists of data that is being inserted. To extend this model to contain both insert and delete operations, we first differentiate between *insert* and *delete* operations by using one additional bit (representing the sign) with the data value. We assume that deletions are negative numbers with negative weights and insertions are positive numbers with positive weights. Let function *sign*(*element*) return 1 for inserted element and return -1 for deleted element. The weight of an element is defined based on the sign function and the height of the compactor. For any element in *compactor*[ $h$ ]:  $weight(element) = sign(element) * 2^{h-1}$ . Hence, the deleted items are recorded with negative weights and inserted items with positive weights.

### 3.3 Conditional Compaction

Once a compactor is full, i.e.,  $k_h = c_h$ , the compactor needs to be compacted to free space for new incoming elements. The presence of inserted and deleted elements requires a new compaction algorithm. This section presents a novel compaction algorithm, Conditional Compaction, described in Algorithm 1. In the conditional compaction algorithm we assume that the capacity of the input compactor is even. If the capacity of the compactor is odd then the first or the last element is randomly kept to ensure that an even sized number of elements are compacted.

Before compacting each *compactor*[ $h$ ], the elements inside the *compactor*[ $h$ ] are sorted (the sorting order does not matter), line 1 of Algorithm 1, After sorting, Conditional Compaction Algorithm follows the following compaction process:

**Discard**( $-e_i, e_i$ ): For every element  $-e_i$ , if there is a matching  $e_i$ , then *Discard*( $-e_i, e_i$ ). This is because an insert  $e_i$  and a delete  $e_i$  operation cancel each other and any such matching pairs of inserts and deletes are discarded in the compaction process. Since the compactor is sorted<sup>5</sup>, a classical two-pointer technique<sup>6</sup>, can be used to find all matching pairs of  $\langle delete(e_i), insert(e_i) \rangle$ . If such pairs exist then *Discard*( $-e_i, e_i$ ) results in the removal of matched pairs which in turn creates free space in the compactor and hence the algorithm skips the rest of the compaction process. A crucial property of compaction is to push elements to the compactor at the next level. Since in this case a compactor will not push any elements to the compactor at the next level, this is not a full compaction.

**Push**( $e_i, e_{i+1}$ ): If no  $\langle delete(e_i), insert(e_i) \rangle$  pairs are found and discarded, then the algorithm performs the push operation. For every pair of adjacent elements,  $e_i$  and  $e_{i+1}$  with the same sign (i.e., they are both inserts or deletes) at level

<sup>5</sup>If compactor is not sorted, we need to keep a map while scanning the compactor to find all matching pairs.

<sup>6</sup>Put one pointer at beginning and one at the end, increments or decrements the pointers to find all matching pairs in a single scan.

$h$ , the algorithm decides on a *random offset* (line 6) whether the first or the second element of the pair is pushed to the next level,  $C[h+1]$ .

**Keep**( $e_i, e_{i+1}$ ): If the pair of adjacent elements,  $e_i$  and  $e_{i+1}$ , have different signs (i.e.,  $e_i$  is a delete operation and  $e_{i+1}$  is an insert operation), then both elements are retained in the compactor at level  $h$ . Note that there will be at most one such mismatched pair with different signs.

The rationale for keeping adjacent elements with different signs is based on the following observation. Assume there is a  $\langle delete(B), insert(A) \rangle$  pair in *compactor*[ $h$ ]. This information implies that the input contains  $w_h$  insertions and at least one of these insertions is element A, and also  $w_h$  deletions and at least one of these deletions is element B. If we randomly push either insert A or delete B into *compactor*[ $h+1$ ] which contains elements with weight of  $2w_h$ , then the sketch not only lost the information about  $w_h$  deletions or  $w_h$  insertions respectively, but also introduce false information on the total number of insertions and deletions.

---

#### Algorithm 1: Conditional Compaction in KLL<sup>±</sup>

---

```

1  sort( $C[h]$ );
2   $Discarded = discard((-x, x) \text{ pairs } \in C[h])$ ;
3  if  $Discarded$  then
4  | // No compaction needed
5  | return;
6   $offset = random(0, 1)$ ;
7  for  $i = 0; i < C[h].length; i = i + 2$  do
8  | if  $sign(C[h][i]) == sign(C[h][i + 1])$  then
9  | | push( $C[h][i + offset]$ ) to  $C[h+1]$ ;
10 | else
11 | | keep( $C[h][i], C[h][i + 1]$ ) in  $TEMP$ ;
12 end
13  $C[h].clear()$ ;
14  $C[h] \leftarrow TEMP$ ;
15 return;

```

---

Two main distinctions that arise in the compaction process of KLL<sup>±</sup> are: (i) discarding matching pairs to free space before the push operations, and (ii) changing the minimum capacity of a compactor to at least three (from two in KLL). The first modification reduces the number of push operations by removing matched pairs which cancel each other. Later, we prove that discard operations do not introduce error and push operations can introduce at most  $w_h$  error.

The second modification changes the minimal capacity of a capacitor to three. This modification is necessary because if the minimum capacity of a compactor is two and these elements neither have the same operation nor are a matched pair, then the Conditional Compaction would retain both elements in the compactor; and hence the compactor remains full even after the compaction process. Increasing the minimum capacity of a compactor to three guarantees a reduction in size after a compaction, when the compactor becomes full. If the compactor contains matched pairs, then the discarded operation can be applied to free space. If there is no matched pair, then a full compactor with capacity three can have four possible cases in the compaction process, as shown in Table 2. Table 2 also describes the compaction strategy for each of these cases, assuming elements are in the sorted

Table 2: Compaction for a full compactor of size 3 consisting of sorted elements  $e_1$ ,  $e_2$ , and  $e_3$  and no matched pair.

3 Inserts +++	50%: push( $e_1$ or $e_2$ ) and keep( $e_3$ ) 50%: keep( $e_1$ ) and push( $e_2$ or $e_3$ )
2 Inserts and 1 Delete -++	keep( $e_1$ ) and push( $e_2$ or $e_3$ )
1 Insert and 2 Deletes --+	push( $e_1$ or $e_2$ ) and keep( $e_3$ )
3 Deletes ---	50%: push( $e_1$ or $e_2$ ) and keep( $e_3$ ) 50%: keep( $e_1$ ) and push( $e_2$ or $e_3$ )

order and no matching pairs. When the compactor is full, one element out of two elements with the same operations is randomly chosen and pushed to the next level, and the element not been pushed is removed.

### 3.4 Estimating Ranks and Quantiles

The Output Operation, described in Algorithm 2, summarizes the current snapshot of the sketch. It uses all the compactors inside the sketch to obtain the quantile information. First, the total weight of the sketch is calculated by summing the weights of all elements, where  $h$  is the height of the compactor and  $k_h$  is the number of elements in compactor[ $h$ ] (line 9):  $TotalWeight = \sum_{h=1}^H \sum_{i=1}^{k_h} weight(e_{h,i})$

Second, a map *ItemWeightSortedMap* of  $\langle item, weights \rangle$  pairs is created to contain the final aggregate weights of each element in the sketch. For each element, the aggregated weights is calculated by incrementing or decrementing the corresponding weights, based on the height of the compactor in which the element was encountered, and whether the corresponding operation was an insert or a delete.

Third, since *ItemWeightSortedMap* is sorted in ascending order by item value, map *Result* calculates the estimated rank of each item by summing through all weights of elements whose values are less than or equal to the item. To transform the rank of an item into the quantile of an item, the rank information is divided by the total weight. The output operation returns a map that contains the quantile information for each item in the sketch. To estimate the rank of element,  $x$ , that is not in the output, we can use the rank of the largest element that is less than  $x$  as the estimation. For items with estimate rank less than zero or larger than the total weight, we treat their rank as zero or as the total weight since all deleted items are previously inserted and hence all ranks should be within  $[0, TotalWeight]$ .

### 3.5 An Illustrative Example

Assume a  $KLL^\pm$  sketch with  $k = 6$ ,  $c = 2/3$  in which the topmost compactor has capacity 6, i.e.,  $k$ , and the second topmost compactor has capacity 4, i.e.,  $ck$ . Items 1 through 8 are inserted and then items 7, 3, 2, and 1 are deleted, in that order, as illustrated in Figure 3. After the deletions, the multiset is left with  $\{4,5,6,8\}$  and their respective ranks are  $R(4)=1$ ,  $R(5)=2$ ,  $R(6)=3$ ,  $R(8)=4$ . In the example shown in Figure 3, the sketch needs to be compacted three times. Assume the first compaction has offset of 1; the second compaction has offset of 0; the third compaction again has offset of 1. The first conditional compaction occurs after element 6 is inserted, and pushes the 2, 4, and 6 from compactor[1] to compactor[2]. Compactor[2] maintains full capacity of 6, while the capacity of compactor[1] is now reduced to 4.

### Algorithm 2: Output Operation in $KLL^\pm$

```

1 Result = map();
2 // A map sorted in ascending Item Order
3 ItemWeightSortedMap =
  OrderedMap< item, weights >;
4 TotalWeight = 0;
5 for all compactor[h] in Sketch do
6   for item in compactor[h] do
7     weight = sign(item) 2h-1;
8     ItemWeightSortedMap[(abs(item))] += weight;
9     TotalWeight += weight;
10  end
11 end
12 PrevW = 0;
13 //traverse ItemWeightSortedMap in ascending order
  for < item, weights > in ItemWeightSortedMap do
14   Result[item] = (weights+PrevW)/TotalWeight;
15   PrevW += weights;
16 end
17 //result contains the quantile info for each item.
18 return Result;
```

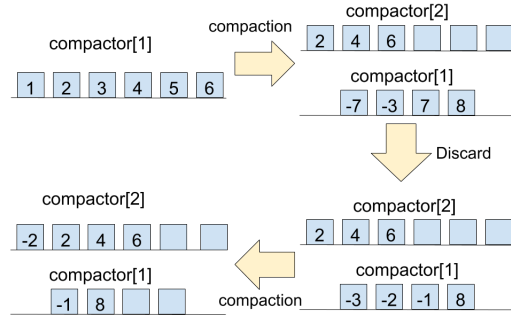


Figure 3: Items from 1 to 8 are inserted and then items 7,3,2,1 are deleted. The resulting dataset is  $\{4,5,6,8\}$ ;  $KLL^\pm$  estimate  $\mathcal{R}(4) = 1$ ,  $\mathcal{R}(5) = 1$ ,  $\mathcal{R}(6) = 3$ ,  $\mathcal{R}(8) = 4$ .

The second conditional compaction occurs after elements 7, 8, -7 and -3 are added. This compaction finds a matching pair of -7 and 7 and hence discards the pair, leaving compactor[1] with -3 and 8. The third conditional compaction occurs after elements -2 and -1 are added. This compaction pushes -2 to compactor[2], and -1 and 8 are both kept in compactor[1] since they have different signs. Based on the current sketch, the original data set are now represented as  $\{-1, 4, 4, 6, 6, 8\}$ . Finally, to estimate ranks,  $KLL^\pm$  computes the weight of all  $(abs(element), weight(element))$  pairs. Then sort based on the element's value:  $(1, -1)$ ,  $(2, 0)$ ,  $(4, 2)$ ,  $(6, 2)$ ,  $(8, 1)$ , and compute their corresponding ranks:  $\mathcal{R}(1) = 0$ ,  $\mathcal{R}(2) = 0$ ,  $\mathcal{R}(4) = 1$ ,  $\mathcal{R}(6) = 3$ , and  $\mathcal{R}(8) = 4$ . To report  $\mathcal{R}(5)$ , we use the rank of largest element that is less than 5, i.e.  $\mathcal{R}(5) = \mathcal{R}(4) = 1$ , differs from true  $R(5)$  by 1. As a result, the output of the sketch in term of rank is:  $\mathcal{R}(4) = 1$ ,  $\mathcal{R}(5) = 1$ ,  $\mathcal{R}(6) = 3$ ,  $\mathcal{R}(8) = 4$ .

### 3.6 Error in Compaction

This section discusses the error introduced by Conditional



Compaction. Essentially, the algorithm introduces errors when pushing one out of two elements and does not introduce any error when discarding matched elements or when retaining a pair of non-matching elements. Consider the  $\text{Push}(e_1, e_2)$  operation in  $\text{compactor}[h]$  with weight  $w_h$ : at random, either  $e_1$  or  $e_2$  is pushed into  $\text{compactor}[h+1]$  with weight  $w_{h+1}$ . Pushing elements to higher weight results in  $+w_h$ ,  $0$ , or  $-w_h$  error. By removing the odd or even items with equal probability, the expected error becomes zero [1].

LEMMA 1. *Hoeffding's Inequality. Let  $X_1, \dots, X_n$  be independent random variables such that  $-w_i \leq X_i \leq w_i$  and the expected value  $E(X_i) = 0$  for  $i = 1, 2, \dots, n$ . Then for any  $t > 0$  we have:*

$$\Pr\left[\left|\sum_{i=1}^n X_i\right| > t\right] \leq 2\exp\left(-\frac{t^2}{2\sum_{i=1}^n w_i^2}\right)$$

In Condition Compaction, Algorithm 1 discards matched pairs of inserted and deleted elements. We need to establish that this does not impact the overall error and make sure the expected error in one compaction is still zero in order to apply Hoeffding's inequality [15] to bound the total errors. Proposition 1 is proved by induction in Appendix A.

PROPOSITION 1. *Discarding matched inserted and deleted elements within a compactor at level  $h$  during a compaction cycle does not introduce any error to the sketch's total error.*

The Conditional Compaction algorithm has three main components that can potentially introduce error: (i)  $\text{Push}(e_i$  or  $e_{i+1})$ , (ii)  $\text{Keep}(e_i$  and  $e_{i+1})$ , and (iii)  $\text{Discard}(-e_i$  and  $e_i)$ . Pushing one of two elements from  $\text{compactor}[h]$  introduces  $+w_h$ ,  $0$ ,  $-w_h$  error; keeping elements in their own compactor introduces no error; discarding the matched pairs of inserted and deleted elements introduces no error (as proved in Proposition 1). Hence, we can assert that the expected error is still 0 when pushing even or odd index terms with equal probability, and thus we can apply Hoeffding's inequality to bound the probability of total error exceeding  $\epsilon \cdot (I - D)$ , where  $\epsilon$  is the desired precision and  $I - D$  is the size of the dataset after the deletions, as shown next.

To show that the  $\text{KLL}^\pm$  sketch ensures that the probability of the total error exceeding  $\epsilon \cdot (I - D)$  is less than a small constant probability,  $\delta$ . Let random variable  $X_{i,h}$  denote the error introduced by the  $i^{\text{th}}$  compaction in  $\text{compactor}[h]$  and let  $m_h$  be the number of compactions that occurred at  $\text{compactor}[h]$ . Note  $X_{i,h}$  can be 1, 0, or -1. Therefore, the total error is computed as:

$$\text{Err} = \sum_{h=1}^H \sum_{i=1}^{m_h} w_h X_{i,h} \quad (1)$$

Since we know that  $E(X_i, h) = 0$  by keeping even and odd terms with equal probability, Hoeffding's inequality can be applied:

$$\Pr(|\text{Err}| > \epsilon(I - D)) \leq 2\exp\left(-\frac{\epsilon^2(I - D)^2}{2\sum_{h=1}^H \sum_{i=1}^{m_h} w_h^2}\right) \leq \delta$$

Where  $\delta$  is a small constant probability.  $\delta$  denotes the maximum failure probability; hence, when  $\delta$  is small, the algorithm has high confidence for success.

### 3.7 Space Bound

In this section, we analyze the space bound and approximation guarantees for  $\text{KLL}^\pm$  with an array of  $H$  compactors. The minimum capacity of compactors is three to handle both insert and delete operations. To guarantee any compactor will have capacity greater than or equal to three, a compactor at height  $h$  has capacity of  $\max(\lfloor c^{H-h}k \rfloor, 3)$ . Hence, as the total height  $H$  increases, there will be a stack of bottom compactors with capacity three. In particular, for compactors with small  $h$  such that  $\lfloor c^{H-h}k \rfloor \leq 3$ , the compactors will have capacity of three.

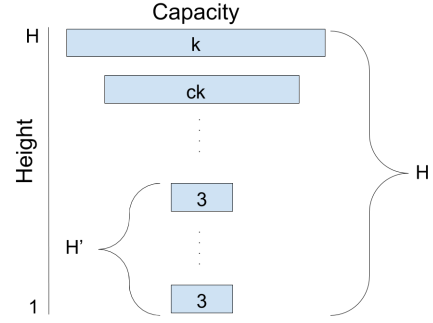


Figure 4:  $H'$  denotes the height of the bottom compactors with size 3, and  $H$  is the height at the highest level. Numbers associated with each compactor is its capacity.

Assume that the  $\text{KLL}^\pm$  sketch has  $k \geq 4$  and  $c \in (0.5, 1)$  with  $c_h = \max(\lfloor c^{H-h}k \rfloor, 3)$ , and the incoming data have  $I$  insertions and  $D$  deletions such that  $D \leq (1 - 1/\alpha)I$ . Let  $r = \frac{D}{I}$  be the rate of deletions such that  $r \leq (1 - 1/\alpha)$  ensuring that the incoming updates have bounded deletions compared to insertions.  $H$  is the height of the compactor at the highest level;  $H'$  is the height of the bottom compactors with capacity three (see Figure 4) such that  $H' \leq H - 1$ .

Because each compactor will undergo compaction when it becomes full, every compactor of capacity three contains at most two items, and hence the total weight of these bottom compactors is:

$$W_{\text{bottom}} \leq \sum_{h=1}^{H'} 2w_h = \sum_{h=1}^{H'} 2^h \leq 2^{H'+1}$$

The total weights of the top compactors is:

$$W_{\text{top}} \leq \sum_{h=H'+1}^H (k_h - 1)w_h \leq (k - 1)2^H$$

Combining the total weights of the bottom  $H'$  and top  $H - H'$  compactors results in an upper bound on the number of items in the dataset:  $I - D \leq 2^{H'+1} + (k - 1)2^H \leq k2^H$ . We can then bound the total number of items:

$$n = I + D = (1 + r)I \leq k2^H \frac{1+r}{1-r} \quad (2)$$

Since the topmost  $\text{compactor}[H]$  exists, this implies that  $\text{compactor}[H-1]$  has been compacted. Thus,  $n \geq c_{H-1}w_{H-1} \geq ck2^{H-2} = k2^H(c/4)$ . By moving  $H$  to the left hand side and  $n$  to the right hand side, we get:

$$H \leq \log_2(4n/c) = \log_2(n/c) + 2$$

The *keep* operation and *discard* operation introduce no error as shown in Proposition 1, and only the push operation will introduce at most  $\pm w_h$  error. Also, because the conditional compaction Algorithm 1 will perform no compaction if there are any discard operations, the push operations are performed on compactor[h] that contains  $k_h = c_h$  elements. Then let  $m_h$  denote the number of compactations occurred for compactor at height  $h$ , and the total error introduced in the sketch is upper bounded by the number of compactations:  $Err = \sum_{h=1}^H \sum_{i=1}^{m_h} w_h X_{i,h}$  where  $X_{i,h}$  is a random variable denote the error introduced by the  $i^{th}$  compaction in which  $E(X_{i,h}) = 0$  and  $|X_{i,h}| \leq 1$ , shown in Equation 1. We can bound the number of compactations,  $m_h$ , at height  $h$ , where each compaction is performed on  $k_h = c_h$  elements with weight  $w_h$  as  $m_h \leq \frac{n}{c_h w_h} \leq \frac{2n}{k^{2H}} (2/c)^{H-h}$  (derivation in Appendix B) by substituting in Equation 2,  $m_h \leq 2 \frac{1+r}{1-r} (2/c)^{H-h}$ . Note in this upper bound, we did not consider discard operations; discard operations only reduce the number of compactations. We now apply Hoeffding's inequality to bound the probability for compactors to introduce more than  $\epsilon \cdot (I - D)$  error:

$$P(|Err| > \epsilon(I - D)) \leq 2 \exp\left(-\frac{\epsilon^2(I - D)^2}{2 \sum_{h=1}^H \sum_{i=1}^{m_h} w_h^2}\right) \leq \delta$$

The denominator can be expanded to  $\sum_{h=1}^H \sum_{i=1}^{m_h} w_h^2 = \sum_{h=1}^H m_h w_h^2 \leq \frac{1+r}{1-r} \frac{16n^2}{c(2c-1)k^2}$ , where  $r = \frac{D}{I} \leq (1 - 1/\alpha)$  (Full derivation included in Appendix C). Let  $C = \frac{c(2c-1)}{32}$ , we get the form:

$$P(|Err| > \epsilon \left(\frac{1-r}{1+r} n\right)) \leq 2 \exp\left(-\left(\frac{1-r}{1+r}\right)^3 C \epsilon^2 k^2\right) \quad (3)$$

Since  $0 \leq r \leq (1 - \frac{1}{\alpha})$ , we can derive  $1 + r \leq 2 - \frac{1}{\alpha}$ ,  $1 - r \geq \frac{1}{\alpha}$ , and hence  $\frac{1-r}{1+r} \geq \frac{1}{2\alpha-1}$ . Setting  $k = (2\alpha - 1)^{1.5} / (\epsilon \sqrt{C}) \sqrt{\ln(2/\delta)}$  suffices to bound the failure probability by  $\delta$ . Therefore, the total space used is:

$$\sum_{h=1}^H c_h \leq 3H + k \sum_{h=1}^H c^{H-h} = O(k + \log(n/k)) \quad (4)$$

**THEOREM 1.** *In the bounded deletion model  $D \leq (1 - 1/\alpha)I$ , there exists a data sketch that computes an  $\epsilon$  approximation for the rank of a single item with probability  $1 - \delta$  whose space complexity is  $O(\frac{\alpha^{1.5}}{\epsilon} \sqrt{\log(1/\delta)} + \log(\frac{\epsilon n}{\alpha^{1.5}}))$ .*

In order to have  $\epsilon$  approximations for all items, the failure probability need to be decreased into  $\epsilon\delta$ , bounding the failure probability for approximating a set of  $O(1/\epsilon)$  items. Thus, to solve all quantile approximation, the space complexity becomes:

$$O\left(\frac{\alpha^{1.5}}{\epsilon} \sqrt{\log(1/\epsilon\delta)} + \log\left(\frac{\epsilon n}{\alpha^{1.5}}\right)\right)$$

Note that the space complexity includes  $\log(\epsilon n)$  term and this is due to the stack of compactors with size 3. We can further reduce the space complexity by replacing the bottom compactors with samplers as discussed in the next section.

## 4. SAMPLER AND MERGEABILITY

In this section, we describe how to incorporate the sampling approaches from the original KLL [19] to further reduce the space bound while still maintaining full mergeability. We also discuss the merge operation between two KLL $^\pm$

sketches to construct a combined sketch over the union of the underlying datasets, and, in the end, propose a new algorithm that leverages KLL $^\pm$  sketches to maintain quantile approximations in fixed-size sliding window.

### 4.1 Sampler-based Bottom Compactors

From the Conditional Compaction Algorithm, when a bottom compactor is full, an inserted item or a deleted item is randomly chosen out of two insertions or two deletions. This is equivalent to replacing the bottom  $H'$  compactors with two *samplers* (one for insertions and one for deletions) to simulate the bottom compactors and consume  $O(1)$  space. Karnin et al. [19] introduced the mergeable sampler and prove the correctness of the sampling schema. The sampler has an associated height  $h$ , and  $h$  increases over more and more inputs. The sampler outputs an item with weight  $2^h$  as input to the compactor of level  $h + 1$ . Therefore, the sketch only contains compactors with heights greater than  $H'$  (from  $H' + 1$  to  $H$ ).

More precisely [19], the sampler has an associated height  $h$  and stores one item with weight at most  $2^{h-1}$ . An update is processed as follows: Let  $w_{internal}$  be the weight of the internal item stored in the sampler and  $w_{new}$  be the weight of the newly arriving item. This  $w_{new}$  is 1 for items from the input; However, when merging two sketches (as in Section 4.2),  $w_{new}$  are used to represent the item weight from another sampler. If  $w_{internal} + w_{new} < 2^h$ , then the sampler stores the new item with probability  $w_{new}/(w_{internal} + w_{new})$ . If  $w_{internal} + w_{new} = 2^h$ , then the sampler pushes the stored item into *compactor*[ $h + 1$ ] and resets  $w_{internal}$  to zero. If  $w_{internal} + w_{new} > 2^h$ , then the sampler pushes the item that has larger weight with probability  $\max(w_{internal}, w_{new})/2^h$ ; the sampler also updates the stored item to the item that has smaller weight and resets  $w_{internal}$  to  $\min(w_{internal}, w_{new})$ . The last case is necessary to support merge operations. Hence for a sampler at height  $h$ , this approach takes an input of  $W$  items where  $2^{h-1} < W \leq 2^h$ . With probability  $W/2^h$ , it pushes one of the observed items chosen at random and otherwise pushes nothing. Lemma 2 in Karnin et al. [19] establishes that the sampler's push operation introduces an unbiased error of  $2^h Y_{h,i}$  where  $Y_{h,i}$  is a random variable with  $E(Y_{h,i}) = 0$  and  $|Y_{h,i}| \leq 1$ .

In KLL $^\pm$ , the bottom compactors of capacity 3 are replaced by two samplers, one for insertions and one for deletions. The insert and delete samplers share the same height thus ensuring that when an item is output to the compactor, the item will have the same weight independent of whether it is an insertion or a deletion. Since there are total  $I$  insertions and  $D$  deletions, as the shared samplers' height increase the insert sampler at height  $h$  at most pushes  $\frac{I}{2^h}$  elements of weight  $2^h$  to the first compactor and similarly the delete sampler of height  $h$  at most pushes  $\frac{D}{2^h}$  elements of weight  $2^h$  to the first compactor. Hence there are at most  $\frac{n}{2^h}$  elements pushed for sampler height  $h$ . The total errors introduced are the sum of all the errors produced by the push operations from both of these two samplers. Given the probability of the samplers total error exceeding  $\epsilon(I - D)$  should be less than or equal to  $\delta$ , setting  $k = (2\alpha - 1)^{1.5} / \epsilon \sqrt{\log(1/\epsilon\delta)}$  is sufficient, as shown in Appendix D.

**THEOREM 2.** *In the bounded deletion model, there exists a data sketch that computes an  $\epsilon$  approximation for the rank*



of a single item with probability  $1 - \delta$  whose space complexity is  $O(\alpha^{1.5}/\epsilon\sqrt{\log(1/\epsilon\delta)})$ .

The total sketch size becomes  $\sum_{h=H'}^H c_h = \sum_{h=H'}^H c^{H-h} k \leq \frac{k}{1-c} = O(k)$ ,  $c \in (0.5, 1)$ , and  $O(k) = O(\alpha^{1.5}/\epsilon\sqrt{\log(1/\delta)})$ . Moreover, fixing the capacity of the top  $O(\log\log(1/\delta))$  compactors to  $k$  as in KLL[19] can be applied to reduce the  $\text{KLL}^\pm$  sketch size to  $O(\frac{\alpha^{1.5}}{\epsilon}\log^2\log\frac{1}{\epsilon})$  and still remain mergeable. In  $\text{KLL}^\pm$  with input size of  $I + D$ , the probability of total error in rank estimation exceeding  $\epsilon(I - D)$  should be less than or equal to  $\delta$ . The proof of this space bounds is shown in Appendix E. Lastly, Karnin et al. [19] suggest to replace the top compactors by the GK sketch to obtain an asymptotically optimal space bound while sacrificing mergeability. However, since the GK sketch does not support deletes, this optimization is not applicable in  $\text{KLL}^\pm$ .

## 4.2 Mergeability

$\text{KLL}^\pm$  sketches with the same deletion upper bound  $\alpha$  are fully mergeable, and the merge operation is described in Algorithm 3. The mergeability of  $\text{KLL}^\pm$  sketches is desirable for distributed settings. The samplers in Section 4.1 are designed to support merge operations by supporting weighted updates. To merge two sketches  $\langle S_A, S_B \rangle$ , first combine the samplers from both sketches. Assume  $S_A$ 's samplers have height  $h_A$  and  $S_B$ 's samplers have height  $h_B$  such that  $h_A \geq h_B$ .  $S_B$  feeds its stored insert and delete items from its insert and delete samplers, respectively, with their internal weights into  $S_A$ 's samplers, and then all compactors with height less than or equal to  $h_A$  in  $S_B$  feed their items into the appropriate sampler in  $S_A$  with proper weights, calculated using their height ( $w_h = 2^{h-1}$ ). Compactors with height larger than  $h_A$  are appended to the same height compactors in  $S_A$ . Lastly, each compactor that contains elements more than its capacity is compacted and the new maximal  $H$  is computed after all compactations are completed i.e., using the combined length of these two sketches.

---

### Algorithm 3: Merge Two $\text{KLL}^\pm$ Sketches

---

```

1 if  $S_A.\text{SamplerHeight} < S_B.\text{SamplerHeight}$  then
2   | swap( $S_A, S_B$ )
3  $S_{A,\text{InsertSampler}}.\text{update}(S_{B,\text{InsertSampler}})$ ;
4  $S_{A,\text{DeleteSampler}}.\text{update}(S_{B,\text{DeleteSampler}})$ ;
5 for all compactor[ $h$ ] in  $S_B$  do
6   | if  $h \leq S_A.\text{SamplerHeight}$  then
7     | for item in compactor[ $h$ ] do
8       | if insert(item) then
9         |   |  $S_{A,\text{InsertSampler}}.\text{update}(\text{item}, 2^{h-1})$ ;
10        |   | if delete(item) then
11          |   |  $S_{A,\text{DeleteSampler}}.\text{update}(\text{item}, 2^{h-1})$ ;
12        |   | end
13      | else
14        |  $S_A.\text{compactor}[h].\text{concatenate}(\text{compactor}[h])$ ;
15 end
16  $S_A.\text{compaction}()$ ;
17 return  $S_A$ ;

```

---

## 4.3 Fixed-Size Sliding Window

We consider a potential use-case for  $\text{KLL}^\pm$  sketch to maintain  $\epsilon$ -approximate quantiles in a fixed-size sliding window




Figure 5: Grey blocks are expired sketches; yellow block is the current *active* sketch; blue block is the current *under-construction* sketch; green block is the current *backup* sketch.

over append only datasets, e.g., transaction logs. In the fixed-size sliding window, the window boundary synchronously shifts over the data, and the sketch reports the quantile information for the most recent  $w$  items.

In the fixed-size sliding window, the input items are insertions, and the expired elements (elements except the recent  $w$  items) need to be deleted. Since the  $\text{KLL}^\pm$  sketch only tolerates bounded deletions, it cannot directly support such a sliding window setting as the ratio between deletions and insertions will eventually approach to one. However, we observe that by partitioning the input data into *three* overlapping blocks, we can ensure the deletions are bounded for each block. The three blocks or data partitions are termed as *active*, *under-construction*, and *backup* sketches as shown in Figure 5. Each block can be represented by one  $\text{KLL}^\pm$  sketch to support bounded deletions.

Initially all sketches contain no elements. When the active sketch has inserted  $\frac{w}{2}$  elements, new incoming elements are inserted into both the active sketch and backup sketch. Once the active sketch contains  $w$  elements and the backup sketch contains  $\frac{w}{2}$  elements, incoming elements are inserted into the backup and under-construction sketches and the outdated (oldest) elements are deleted from the active sketch. When the backup sketch has  $w$  elements, the active sketch, which has deleted  $\frac{w}{2}$  elements, becomes expired. The backup sketch now becomes the new active sketch and the under-construction sketch becomes the new backup sketch. Since deleting the expired sketch (previously active sketch) saves space, the freed space is used to construct the new under-construction sketch.

At any point in time, the quantile information is reported of the most recent  $w$  elements by merging the active and under-construction sketches, i.e., the yellow and blue blocks shown in Figure 5. Also note that, since deletions only arrive when the sketch has already inserted  $w$  items, all sketches will have a deletions to insertions ratio of  $r = \frac{d}{w}$ , and because a sketch expires after at most  $\frac{w}{2}$  deletions, i.e.,  $d \leq \frac{w}{2}$ ,  $r \leq \frac{0.5w}{w} \leq \frac{1}{2}$  and hence setting  $\alpha = 2$  is sufficient to ensure that  $r \leq \frac{1}{2}$ . In fact, the  $\alpha$  can also be changed to other values; however, the number of backup sketches will also change accordingly. Thus, we have demonstrated a randomized algorithm to maintain quantile summary in fixed-size sliding window with three  $O(\frac{\alpha^{1.5}}{\epsilon}\log^2\log\frac{1}{\epsilon})$  sketches and  $\alpha = 2$ .

## 5. EXPERIMENTS

This section experimentally evaluates  $KLL^\pm$  for datasets that experience an input data stream that consists of both insertions and bounded deletions of elements.  $KLL^\pm$  is the first quantile sketch algorithm to operate in the bounded deletion model and the experiments aim to identify the overhead incurred in accounting for bounded deletes compared to other sketches:

- **KLL:** Since KLL is insertion-only, the input stream only inserts those elements that are left after all the deletions.
- **Two-Parallel KLL:** Two independent KLL sketches: one for insertions and one for deletions; then aggregate their estimations to approximate quantiles.
- **DCS<sup>7</sup>:** A universe-driven sketch that assumes a bounded universe to tolerate an arbitrary number of deletions.

### 5.1 Experimental Setup

We set  $c = 2/3$  for all the experiments, and implemented  $KLL^\pm$  by enhancing the Python 3.7.6 code-base of the original KLL sketch algorithm presented in [17, 19]. The changes incorporated in our implementation are: (i) the minimum capacity of a compactor is increased from 2 to 3; (ii) bottom compactors of capacity 3 are replaced by 2 samplers, one for insertions and one for deletions; and (iii) the compaction algorithm is modified to implement our new Conditional Compaction algorithm.

The metric for accuracy measurement is the Kolmogorov-Smirnov divergence [4] – the *maximum* deviation among all quantile queries – a measurement widely used to perform comparisons between CDFs with different distributions [17]. For each experiment, the maximum errors are averaged over 5 independent runs. These experiments are performed on a machine with 2.9 GHz quad-core Intel Core i7 processor and 16GB of 2133MHz LPDDR3 on-board memory.

### 5.2 Data Sets

The experimental evaluation is conducted using both synthetic data sets and real world data sets consisting of elements that are inserted and deleted. For the synthetic data, we consider three different distributions:

- **Uniform Distribution:** The insertions are randomly generated from a discrete uniform distribution, and the deletions are uniformly chosen from the insertions.
- **Zipf Distribution:** The elements drawn from bounded universe and the frequencies of elements follow the Zipf Law [28], in which the frequency of element with rank  $R$ :  $f(R, s) = \frac{\text{constant}}{R^s}$  where  $s$  indicate skewness. Deletions are uniformly chosen from the insertions.
- **Binomial Distribution:** The elements are generated according to the binomial distribution with parameters  $n$  and  $p$  where  $p$  is the probability of success in  $n$  independent Bernoulli trials.

In addition to the synthetic data sets, we used the following real world Wiki dataset<sup>8</sup>:

<sup>7</sup>See chapter 4 in [9] for implementation details.

<sup>8</sup><https://dumps.wikimedia.org/other/pagecounts-ez/>

- **Page View Statistics for Wikimedia (Wiki) [22]:** This is an extensive data set consisting of page count files from 2007 to 2016, wherein the experiments use the most recent, i.e., the 2016 page count files. The 2016 page count files include 8 months of  $\langle projectname, pagename, \#requests, \#bytes \rangle$  tuples. The data are aggregated by day and within each day, data are sorted on  $projectname$  and  $pagename$ . In the experiments, deleted items are random uniformly chosen from the inserted items and each update is a concatenation of  $projectname$  and  $pagename$  and the comparison model is lexicographic.

We also conducted experiments by exploring two additional properties of the data sets:

- **Sorted Dataset:** Input is sorted in descending order such that insertions arrive before deletions; The deletions are uniformly chosen from the insertions.
- **Shuffled Dataset:** The insertions are randomly shuffled and the deletions are also randomly shuffled and uniformly chosen from the insertions.

### 5.3 Evaluation

All experiments measure the maximum error by Kolmogorov Smirnov divergence. The y-axis depicts the average of maximum quantile error over 5 independent runs: lower y-axis values indicate better accuracy. Most of the following experiments evaluate the error value in approximating quantiles while increasing the sketch size in which the x-axis denotes the sketch size. Except for Section 5.3.6, we assume all insertions arrive before any deletions into the sketch. This input pattern is in fact an adverse pattern as the discard operation will find less matched pairs.

#### 5.3.1 $KLL^\pm$ vs. Two-Parallel KLL vs. DCS

This experiment compares the accuracy among  $KLL^\pm$ , two-parallel KLL method, and DCS under the same memory budget. The experiment measures the quantile estimation when the underlying dataset is entirely changed: We first insert one million elements from the binomial distribution with parameter of  $B(2^{16}, 0.5)$  in which the elements should be densely centered at value  $2^{15}$ . Then another one million elements are inserted from a uniform distribution in Figure 6(a), or from a zipf distribution with skewness factor of 0.5 (moderate skew) in Figure 6(b). Both the uniform distribution and zipf distribution assume a bounded universe where  $|U| = 2^{16}$ , and all inputs are randomly shuffled. Finally, all elements from the binomial distribution are deleted. Since the total insertions are two millions and the total deletions are one million, the delete:insert ratio is 0.5.

Figure 6(a) and (b) show that both data driven sketch approaches,  $KLL^\pm$  and Two Parallel KLL, perform significantly better than the universe driven DCS sketch. Furthermore, although the maximum error of  $KLL^\pm$  and Two Parallel Method decreases as the sketch size increases,  $KLL^\pm$  has less maximum error across all sketch sizes. This finding is expected since  $KLL^\pm$  makes a best effort to apply discard operations thus catching cancellations early on and reducing the number of compactions, whereas in the two-parallel KLL method each sketch has no knowledge about the other and accumulates all the errors. We also observe that DCS performs worse on the zipf distribution compared to a uniform




Figure 6: Comparison of  $KLL^\pm$  with Two Parallel KLL and DCS when the underlying data distribution is entirely changed from Binomial to (a) Uniform and to (b) Zipf (.5). (c)  $KLL^\pm$  accuracy with different delete:insert ratios.

distribution. The skewness in the input data distribution affects the performance of DCS in which DCS’s accuracy decreases when the skewness increases. This observation for universe-driven DCS sketch is in consistent with the theoretical expectations [27]. On the other hand, the performance of both data-driven sketches,  $KLL^\pm$  and two-parallel KLL, is not affected by the skew in the input.

### 5.3.2 Error correlations: deletion ratio & sketch size

We experimentally verified that the delete to insert ratio on the data set affects the accuracy of the sketch. By scaling the sketch size by a factor of  $(2\alpha - 1)^{1.5}$  according to Section 4,  $KLL^\pm$  can keep the errors to the same level compare with errors from KLL with no deletions, as shown in Figure 6 (c). This scaling factor depends on the insert:delete ratio  $r$ . For example, when  $r = 0.5$ , we set  $\alpha = 2$  to increase the  $KLL^\pm$  sketch size. Note  $(2\alpha - 1)^{1.5}$  is actually an upper bound, as in the proof we make no assumptions on the number of discard operations which is affected by the input pattern. Hence, we expect  $KLL^\pm$  with scaled sketch sizes to perform no worse than the original KLL with no deletions.

The input stream contains one million elements in random shuffled order and drawn from a uniform distribution. The deletions are uniformly chosen from the insertions. This experiment shows the interplay of space versus accuracy when the deletes in the input stream increase and verifies the theoretical claim made in Section 4.1 that increasing the sketch size in accordance with the increase in delete:insert ratios keeps the error roughly constant. Figure 6 (c) shows that the higher delete:insert ratio leads to larger errors, while increasing the sketch size with the increase in deletions ensure the accuracy is no worse than the original KLL with no deletions. For instance, when KLL uses  $k = 512$ , the corresponding average maximum error value is 0.0028. For  $r = 0.25$ , if the space used by  $KLL^\pm$  is scaled with  $\alpha = \frac{4}{3}$  then  $k$  increased to 1102 and the corresponding averaged maximum error value is 0.0022. Similarly for  $r = 0.5$ , scaling the space used by  $KLL^\pm$  with  $\alpha = 2$ ,  $k$  increases to 2661 and the corresponding averaged maximum error value is 0.0019. This verifies the theoretical expectation that scaling the sketch size according to the delete:insert ratio keeps the error roughly constant.

### 5.3.3 Different Data Distributions

In this experiment, we further demonstrate the trade-off between space and accuracy on different data distributions.

Figures 7(a) through 7(c) depict the quantile approximation errors for two types of data distributions: for synthetic datasets *uniform* and *zipf* and for real world data *Wikimedia page view statistics*. For each type of data distribution, the experiment also plots the behaviour of  $KLL^\pm$  and KLL when the input stream is sorted vs. shuffled. While the universe-driven DCS’s accuracy is independent from the input pattern and the number of deletions [27], the DCS plots in Figure 7(a) and (b) are used as a comparison reference for  $KLL^\pm$  with large delete:insert ratio, namely 90%. In Figure 7 (a) and (b), note that even when a significant number of items are deleted and the delete:insert ratio becomes 90%,  $KLL^\pm$  still performs well compare to DCS.

The behavior of all sketches is consistent across all data distributions: using more sketch space leads to smaller errors. Furthermore, the experiment shows that  $KLL^\pm$  behaves worse on shuffled streams compared to sorted streams, across all types of data distributions. This finding is consistent with [17] where it is observed that the randomness within the stream affects the accuracy of the KLL sketch.

### 5.3.4 Update Time

In this section, we experimentally compare the update time among KLL,  $KLL^\pm$ , and DCS sketches. The items in the input stream are shuffled uniform distribution of  $|U| = 2^{16}$ . We also include the update time of DCS for  $U = 2^{32}$ . All sketches share the same space budget. In Figure 8(a), the y-axis is the update-time and the x-axis is the stream length, smaller y-value implies faster update time per item. The result is aligned with our expectation, namely that  $KLL^\pm$  incurs slightly more time than KLL using the same memory budget, because  $KLL^\pm$  makes a best effort to apply discard operations before compaction. In DCS, on the other hand, the universe size affects the update time as a larger universe size leads to slower update time per item.

### 5.3.5 Error Sensitivity to stream length

This experiment demonstrates that the quantile approximation error of  $KLL^\pm$  is independent of the input stream length, as shown in Figure 8(b). In this figure, the x-axis denotes the input stream length in which  $N = (I+D)$ , where  $N$  is the total number of operations in the stream, and  $r$  determines the percentage of elements that are deleted. For this experiment, KLL and  $KLL^\pm$  both uses parameter  $k = 1024$ . The insertions are randomly shuffled items from uniform distribution and the deletions are randomly chosen from the




Figure 7: Trade-off between space and accuracy on different data distributions




Figure 8: (a) Update time per item in KLL,  $KLL^\pm$ , and DCS with universe size of  $2^{16}$  and  $2^{32}$ ; (b) Maximum error of shuffled uniform streams with varying stream length; (c) Interleaved deletion pattern.

previously inserted items. This experiment highlights that, for a given delete:insert ratio, the error remains roughly constant with the increase in stream length and hence is independent of the input stream size.

### 5.3.6 Interleaved Deletions

All prior experiments were performed under the assumption that the input consists of all inserted items first, followed by deleted items. In this experiment, we explore the performance of  $KLL^\pm$  with inputs consisting of interleaved inserts and deletes. The input stream is a shuffled uniform distribution. The input is divided into a number of *folds*, such as 1, 10, 100, and 1000 folds, in which 1 fold means the whole input stream consist of a single pair of insertions and deletions (i.e.,  $\langle \text{all inserts, all deletes} \rangle$ ). Similarly, a 100 folds means the stream consists of 100 pairs of  $(\frac{I}{100}, \frac{D}{100})$  sub-streams in which  $I$  is the total number of insertions and  $D$  is the total number of deletions. For each pair of  $\langle \text{insertions, deletions} \rangle$ , the deletions are uniformly chosen from the insertions. More folds imply deleted items are closer to their corresponding inserted items. In Proposition 1, we showed that discard operations introduce no error. When deleted items are closer to the inserted items, we expect more discard operations leading to fewer errors. Figure 8(c) shows that when number of folds are small, the averaged maximum error is higher, and when number of folds increases and hence the number of interleavings is greater,  $KLL^\pm$  improves its performance by applying more discard operations, thus reducing the overall averaged maximum error. On the other hand, the performance of Two Parallel KLL sketches

will not improve in the interleaved deletion pattern, as the insertions and deletions are separately managed. All experiments in Section 5.3.1 have insertions before any deletions. When the insertions and deletions are mixed, the maximum error of the  $KLL^\pm$  decreases which reflects more realistic real world scenarios.

## 6. CONCLUSIONS

Quantile approximations have an important role in both research as well as real world systems. Many algorithms have been proposed to approximate quantiles for the the insertion-only and the turnstile models. In this work, we propose a data-driven algorithm  $KLL^\pm$  to approximate quantiles in the bounded deletions model. To our knowledge, this is the first work to account for bounded deletions for approximating quantiles. The experimental evaluations of  $KLL^\pm$  highlight that the accuracy provided by the quantile approximations of  $KLL^\pm$  is significantly better than the state of the art DCS sketch even when a significant fraction (90%) of elements are deleted. We also demonstrate that the accuracy of  $KLL^\pm$  is not affected by the underlying distribution of the data which is not the case with the universe-driven sketch such as DCS. Furthermore, the experiments highlight that  $KLL^\pm$  has much faster update times compared to DCS. These characteristics of  $KLL^\pm$  makes it a practical choice for real world applications. Finally, we also demonstrated that the deletion property of  $KLL^\pm$  can be leveraged for maintaining quantile estimation in fixed-sized sliding windows over datasets.

## 7. REFERENCES

- [1] P. K. Agarwal, G. Cormode, Z. Huang, J. Phillips, Z. Wei, and K. Yi. Mergeable summaries. In *Proceedings of the 31st ACM SIGMOD-SIGACT-SIGAI symposium on Principles of Database Systems*, pages 23–34, 2012.
- [2] A. Arasu and G. S. Manku. Approximate counts and quantiles over sliding windows. In *Proceedings of the twenty-third ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems*, pages 286–296, 2004.
- [3] B. H. Bloom. Space/time trade-offs in hash coding with allowable errors. *Communications of the ACM*, 13(7):422–426, 1970.
- [4] F. P. Cantelli. Sulla determinazione empirica delle leggi di probabilita. *Giorn. Ist. Ital. Attuari*, 4(421-424), 1933.
- [5] M. Charikar, K. Chen, and M. Farach-Colton. Finding frequent items in data streams. In *International Colloquium on Automata, Languages, and Programming*, pages 693–703. Springer, 2002.
- [6] Z. Chen and A. Zhang. A survey of approximate quantile computation on large-scale data. *IEEE Access*, 8:34585–34597, 2020.
- [7] G. Cormode, T. Johnson, F. Korn, S. Muthukrishnan, O. Spatscheck, and D. Srivastava. Holistic udafs at streaming speeds. In *Proceedings of the 2004 ACM SIGMOD international conference on Management of data*, pages 35–46, 2004.
- [8] 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.
- [9] G. Cormode and K. Yi. *Small Summaries for Big Data*. Cambridge University Press, 2020.
- [10] P. Flajolet, É. Fusy, O. Gandouet, and F. Meunier. Hyperloglog: the analysis of a near-optimal cardinality estimation algorithm. 2007.
- [11] A. C. Gilbert, Y. Kotidis, S. Muthukrishnan, and M. J. Strauss. How to summarize the universe: Dynamic maintenance of quantiles. In *VLDB’02: Proceedings of the 28th International Conference on Very Large Databases*, pages 454–465. Elsevier, 2002.
- [12] M. Greenwald and S. Khanna. Space-efficient online computation of quantile summaries. *ACM SIGMOD Record*, 30(2):58–66, 2001.
- [13] M. B. Greenwald and S. Khanna. Quantiles and equi-depth histograms over streams. In *Data Stream Management*, pages 45–86. Springer, 2016.
- [14] X. Guo, K. C. Lichtendahl, and Y. Grushka-Cockayne. *Quantile Forecasts of Product Life Cycles Using Exponential Smoothing*. Harvard Business School, 2018.
- [15] W. Hoeffding. Probability inequalities for sums of bounded random variables. In *The Collected Works of Wassily Hoeffding*, pages 409–426. Springer, 1994.
- [16] P. Indyk. Algorithms for dynamic geometric problems over data streams. In *Proceedings of the thirty-sixth annual ACM symposium on Theory of computing*, pages 373–380, 2004.
- [17] N. Ivkin, E. Liberty, K. Lang, Z. Karnin, and V. Braverman. Streaming quantiles algorithms with small space and update time. *arXiv preprint arXiv:1907.00236*, 2019.
- [18] R. Jayaram and D. P. Woodruff. Data streams with bounded deletions. In *Proceedings of the 37th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems*, pages 341–354, 2018.
- [19] Z. Karnin, K. Lang, and E. Liberty. Optimal quantile approximation in streams. In *2016 IEEE 57th annual symposium on foundations of computer science (focs)*, pages 71–78. IEEE, 2016.
- [20] G. S. Manku, S. Rajagopalan, and B. G. Lindsay. Approximate medians and other quantiles in one pass and with limited memory. *ACM SIGMOD Record*, 27(2):426–435, 1998.
- [21] G. S. Manku, S. Rajagopalan, and B. G. Lindsay. Random sampling techniques for space efficient online computation of order statistics of large datasets. *ACM SIGMOD Record*, 28(2):251–262, 1999.
- [22] D. Mituzas. Page view statistics for wikimedia projects, 2013.
- [23] J. I. Munro and M. S. Paterson. Selection and sorting with limited storage. *Theoretical computer science*, 12(3):315–323, 1980.
- [24] V. Poosala, P. J. Haas, Y. E. Ioannidis, and E. J. Shekita. Improved histograms for selectivity estimation of range predicates. *ACM Sigmod Record*, 25(2):294–305, 1996.
- [25] V. Poosala, Y. E. Ioannidis, et al. Estimation of query-result distribution and its application in parallel-join load balancing. In *VLDB*, volume 96, pages 3–6, 1996.
- [26] N. Shrivastava, C. Buragohain, D. Agrawal, and S. Suri. Medians and beyond: new aggregation techniques for sensor networks. In *Proceedings of the 2nd international conference on Embedded networked sensor systems*, pages 239–249, 2004.
- [27] L. Wang, G. Luo, K. Yi, and G. Cormode. Quantiles over data streams: an experimental study. In *Proceedings of the 2013 ACM SIGMOD International Conference on Management of Data*, pages 737–748, 2013.
- [28] G. K. Zipf. *Human behavior and the principle of least effort: An introduction to human ecology*. Ravenio Books, 2016.



## APPENDIX

### A. DISCARD OPERATIONS INTRODUCE NO ERROR TO THE TOTAL ERROR

PROPOSITION 1. *Discarding matched inserted and deleted elements within a compactor at level  $h$  during a compaction cycle does not introduce any error to the sketch's total error.*

PROOF. By Induction.

*Base case:* X and -X both are in *compactor*[1]. Assume the input stream inserts X and then deletes X, incurring no compaction process in between. Now, if *compactor*[1] undergoes a compaction, removing X and -X from *compactor*[1], does not introduce any error since X is intended to be deleted.

*Induction hypothesis:* Assume *compactor*[ $h$ ] includes X and -X and removing both X and -X from *compactor*[ $h$ ] during its compaction does not introduce any additional errors.

*Induction Step:* *Compactor*[ $h + 1$ ] contains X and -X. Removing X and -X from *compactor*[ $h + 1$ ] does not introduce any additional errors. The input data must have consisted of  $2 * 2^h$  of insertions and  $2 * 2^h$  deletions, and X must have been inserted at least once and deleted at least once. Note the X in *compactor*[ $h+1$ ] is the push result of a (X,Y) or (Y,X) pair in *compactor*[ $h$ ] and Y may not be the same as X. The error introduced from the push operations is accounted separately. Since X at *compactor*[ $h+1$ ] implies  $2w_h$  occurrence of X and X at *compactor*[ $h$ ] implies  $w_h$  occurrence of X, having X and -X in *compactor*[ $h + 1$ ] is logically equivalent to having 2 pairs (X, X) and (-X,-X) in *compactor*[ $h$ ], and after reorganization we can regroup to two pairs of (-X, X) in *compactor*[ $h$ ], as shown in Figure 9. Since removing the (X, -X) pair in *compactor*[ $h$ ] introduces no additional error by induction hypothesis, removing of X and -X in *compactor*[ $h + 1$ ] also introduces no additional error.

*Conclusion:* By the principle of induction, discard matched pair in the same *compactor*[ $h$ ] does not introduce error.  $\square$

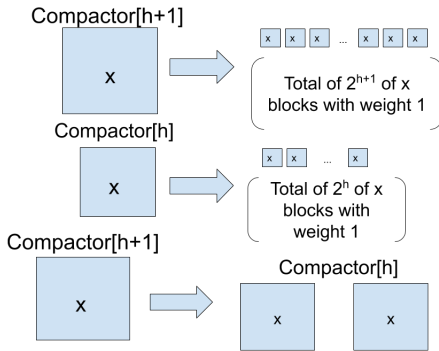


Figure 9: X in *compactor*[ $h+1$ ] is logically equivalent to having two X in *compactor*[ $h$ ], since they represent equal numbers of X blocks with weight 1.

### B. BOUND ON THE NUMBER OF COMPACTIONS

Let  $m_h$  denote the number of compactions occurs for compaction with height  $h$ . Since each compaction is performed on  $k_h$  elements with weight  $w_h$ :  $m_h \leq \frac{n}{k_h w_h} = \frac{n}{c^{H-h} k 2^{h-1}} = \frac{n}{c^{H-h} k 2^{(H-h)+h-1-(H-h)}} = \frac{n}{k 2^{(H-h)+h-1} (\frac{2}{c})^{H-h}}$

$$m_h \leq \frac{n}{k 2^{H-1}} \left(\frac{2}{c}\right)^{H-h}$$

By substituting in equation 2:

$$m_h \leq 2 \frac{1+r}{1-r} \left(\frac{2}{c}\right)^{H-h}$$

### C. EXPANDING BOTTOM SUMMATION

We want to derive the upper bound for the failure probability using Hoeffding's inequality, and we can derive an upper bound for the bottom summation in the equation. Let  $r$  denote the ratio between I and D,  $r = \frac{D}{I} \leq (1 - 1/\alpha)$ .  $\sum_{h=1}^H \sum_{i=1}^{m_h} w_h^2 = \sum_{h=1}^H m_h w_h^2 \leq \frac{1+r}{1-r} \sum_{h=1}^H 2 \left(\frac{2}{c}\right)^{H-h} 2^{2h-2} = \frac{1+r}{1-r} \sum_{h=1}^H \frac{2^{H-h}}{c^{H-h}} 2^{2h-1} = \frac{1+r}{1-r} \left(\frac{2}{c}\right)^{H-1} \sum_{h=1}^H 2^h * c^{h-1} = 2 \frac{1+r}{1-r} \left(\frac{2}{c}\right)^{H-1} \sum_{h=1}^H (2c)^{h-1}$  since it is a geometric sequence,  $2 \frac{1+r}{1-r} \left(\frac{2}{c}\right)^{H-1} \frac{(2c)^H - 1}{2c-1} \leq 2 \frac{1+r}{1-r} \left(\frac{2}{c}\right)^{H-1} \frac{(2c)^H}{2c-1} = 2 \frac{1+r}{1-r} 2^{2H} c / (2c-1)$  since  $H \leq \log_2 n / ck + 2$ :

$$\sum_{h=1}^H \sum_{i=1}^{m_h} w_h^2 \leq \frac{1+r}{1-r} \frac{16n^2}{c(2c-1)k^2}$$

### D. KLL $^\pm$ WITH COMPACTORS AND SAMPLERS.

Let  $m_{h,type}$  denote the total number of sample operations performed on a *type* sampler with height  $h$  (i.e, over  $2^{h-1}$  items). Since the sample operation is performed after processing at least  $2^{h-1}$  items,  $m_{h,insert} \leq \frac{I}{2^{h-1}}$  and  $m_{h,delete} \leq \frac{D}{2^{h-1}}$ . Let  $m_{h,n}$  denote the total number of sample operations performed at sampler height  $h$ , such that  $m_{h,n} = m_{h,insert} + m_{h,delete}$ , and  $m_{h,n} \leq \frac{I+D}{2^{h-1}}$ :

$$Err_{H'} = \sum_{h=1}^{H'} \sum_{i=1}^{m_{h,n}} 2^h Y_{h,i}$$

where  $Y_{h,i}$ <sup>9</sup> is a random variable for the  $i$ -th sample operation on sampler with sampler height  $h$  such that the expected error is 0 and is  $|Y_{h,i}| \leq 1$ . We can apply Hoeffding's inequality to bound the total error of the sample operations:

$$P(|Err_{H'}| > \epsilon(I - D)) \leq 2 \exp\left(-\frac{\epsilon^2(I - D)^2}{2Err_{H'}^2}\right)$$

The denominator can be expanded:  $Err_{H'}^2 = \sum_{h=1}^{H'} \sum_{i=1}^{m_{h,n}} 2^{2h} = \sum_{h=1}^{H'} m_{h,n} 2^{2h} \leq \sum_{h=1}^{H'} n 2^{h+1} \leq 4n 2^{H'} = \frac{4n 2^H}{2^{H-H'}} \leq \frac{16n^2}{ck 2^{H-H'}}$ . Substitute the bottom summation into the inequality, we get :

$$P(|Err_{H'}| > \epsilon(I - D)) \leq 2 \exp\left(-\left(\frac{1-r}{1+r}\right)^3 c \epsilon^2 k 2^{H-H'} / 32\right)$$

and hence  $P(Err > 2\epsilon(I - D)) \leq 2 \exp\left(-\left(\frac{1-r}{1+r}\right)^3 c \epsilon^2 k 2^{H-H'} / 32\right) + 2 \exp\left(-\left(\frac{1-r}{1+r}\right)^3 C \epsilon^2 k^2\right)$ . Since  $\frac{1-r}{1+r} \geq \frac{1}{2\alpha-1}$ , taking  $H' = H - O(\log(k))$  and setting  $k = O(\alpha^{1.5} / \epsilon \sqrt{\log(1/\delta)})$  is sufficient.

<sup>9</sup>lemma 2 at page 8 in [19].



## E. FIXING TOPMOST COMPACTOR

Let the topmost  $s$  compactors to have fix capacity  $k$  instead of exponentially decreasing the capacity, then our sketch contains three components: two samplers correspond to level 1 to  $H''$ , compactors with decreasing capacity correspond to level  $H''$  to  $H'$ , and  $s$  topmost compactors with fixed capacity  $k$  correspond to  $H'$  to  $H' + s = H$ . The two sampler and bottom compactor contribute at most  $\epsilon(I - D)$  errors with  $1 - \delta$  probability as long as  $(2\alpha - 1)^{1.5}\epsilon k 2^s \geq c' \sqrt{\log(2/\delta)}$ . Moreover, the topmost compactors contribute at most  $\sum_{H'+1}^H m_h w_h = \sum_{H'+1}^H n/k = sn/k$  errors, such that we need  $s(1+r)I/k \leq \epsilon(1-r)I$ . Setting  $s$  to  $O(\log \log \frac{1}{\delta})$  and  $k$  to  $O(\frac{\alpha^{1.5}}{\epsilon} \log \log (\frac{1}{\delta}))$  will satisfy both requirements. The topmost compactors of fixed capacity become the dominant term in space complexity,  $O(sk)$ , and thus the total space complexity becomes  $O(\frac{\alpha^{1.5}}{\epsilon} \log^2 \log (\frac{1}{\delta}))$

## F. WORST CASE UPDATE TIME

---

**Algorithm 4:** Sweep Compactor Algorithm [17]

---

```

1  $i^* = \operatorname{argmin}(C[h][i] \geq C[h].\theta)$  ;
2 if  $i == \text{None}$  then
3    $i = 0$  ;
4 if  $\operatorname{sign}(C[h][i]) \neq \operatorname{sign}(C[h][i+1])$  then
5   //can't push between an insertion and a deletion
6    $i += 1$  ;
7  $\text{offset} = \operatorname{random}(0,1)$  ;
8  $C[h].\theta = C[h][i+1]$  ;
9  $\text{push}(C[h][i+\text{offset}])$  ;
10  $\text{delete } C[h][i+1]$  and  $C[h][i]$  ;

```

---

[17] proposed many novel optimizations which are orthogonal to our work. For instance, the Sweep Compactor which reduce the worst case update from  $O(k \log k)$  to  $O(\log k)$  can be applied in  $\text{KLL}^\pm$  sketch. The worst case update time occurs when a full compactor need to be compacted and the most expensive operation in the compaction algorithm is the sorting the all items in compactor. [17] improve the worst update time by maintaining a balanced binary search tree to store items instead saving the items in a list, and once compactor is full, it will push only one item into higher compactor (saves two spaces).

Essentially to improve the worst case update time for  $\text{KLL}^\pm$  sketch, we need to apply discard operation on the balanced BST structure, Fortunately, the worst case time complexity for search, insert, delete operations in balanced BST structure are log time to the size of the tree. Therefore, before inserting any item  $e$ , we can break it down into two parts: (i) search for item  $-e$ ; (ii) If  $-e$  is found, then delete  $-e$ . Else, insert  $e$  into the tree. Therefore, We are certain that when the compactor is full and need to be compacted, the tree contains no  $(-e_i, e_i)$  pairs since such pairs are checked and removed at each update. Now, the compaction algorithm Sweep Compactor can be applied to  $\text{KLL}^\pm$ .