Time-Decayed Correlated Aggregates over Data Streams *

Graham Cormode

AT&T Labs–Research

graham@research.att.com

Srikanta Tirthapura   Bojian Xu

ECE Dept., Iowa State University

{snt,bojianxu}@iastate.edu

Abstract

Data stream analysis frequently relies on identifying correlations and posing conditional queries on the data after it has been seen. *Correlated aggregates* form an important example of such queries, which ask for an aggregation over one dimension of stream elements which satisfy a predicate on another dimension. Since recent events are typically more important than older ones, *time decay* should also be applied to downweight less significant values. We present space-efficient algorithms as well as space lower bounds for the time-decayed correlated sum, a problem at the heart of many related aggregations. By considering different fundamental classes of decay functions, we separate cases where efficient relative error or additive error is possible, from other cases where linear space is necessary to approximate. In particular, we show that no efficient algorithms are possible for the popular sliding window and exponential decay models, resolving an open problem. The results are surprising, since efficient approximations are known for other data stream problems under these decay models. This is a step towards better understanding which sophisticated queries can be answered on massive streams using limited memory and computation.

Keywords: data stream, time decay, correlated, aggregate, sum

1 Introduction

Many applications such as Internet monitoring, information systems auditing, and phone call quality analysis involve monitoring massive data streams in real time. These streams arrive at high rates, and are too large to be stored in secondary storage, let alone in main memory. An example stream of VoIP call data records (CDRs) may have the call start time, end time, packet loss rate, along with identifiers such as source and destination phone numbers. This stream can consist of billions of items per day. The challenge is to collect sufficient summary information about these streams in a single pass to allow subsequent post hoc analysis.

There has been much research on estimating aggregates along a single dimension of a stream, such as the median, frequency moments, entropy, etc. However, most streams consist of multi-dimensional data. It is imperative to compute more complex multi-dimensional aggregates, especially those that can "slice and dice" the data across some dimen-

sions before performing an aggregation, possibly along a different dimension. In this paper, we consider such *correlated aggregates*, which are a powerful class of queries for manipulating multi-dimensional data. These were motivated in the traditional OLAP model [3], and subsequently for streaming data [1, 8]. For example, consider the query on a VoIP CDR stream: "what is the average packet loss rate for calls within the last 24 hours that were less than 1 minute long"? This query involves a selection along the dimensions of call duration and call start time, and aggregation along the third dimension of packet loss rate. Queries of this form are useful in identifying the extent to which low call quality (high packet loss) causes customers to hang up. Another example is: "what is the average packet loss rate for calls started within the last 24 hours with duration greater than the median call length (within the last 24 hours)?", which gives a statistic to monitor overall quality for "long" calls. Such queries cannot be answered by existing streaming systems with guaranteed accuracy, unless they explicitly store all data for the last 24 hours, which is infeasible.

In this work, we present algorithms and lower bounds for approximating time-decayed correlated aggregates on a data stream. These queries can be captured by three main aspects: selection along one dimension (say x-dimension) and aggregation along a second dimension (say y-dimension) using time-decayed weights defined via a third (time) dimension. The time-decay arises from the fact that in most streams, recent data is naturally more important than older data, and in computing an aggregate, we should give a greater weight to more recent data. In the examples above, the time decay arises in the form of a sliding window of a certain duration (24 hours) over the data stream. More generally, we consider arbitrary time-decay functions which return a weight for each element as a non-increasing function of its age—the time elapsed since the element was generated. Importantly, the nature of the time-decay function will determine the extent to which the aggregate can be approximated.

We focus on the *time-decayed correlated sum* (henceforth referred to as DCS), which is a fundamental aggregate, interesting in its own right, and to which other aggregates can be reduced. An exact computation of the correlated sum requires multiple passes through the stream, even with no time-decay where all elements are weighted equally. Since

we can afford only a single pass over the stream, we will aim for approximate answers with accuracy guarantees. In this paper, we present the first streaming algorithms for estimating the DCS of a stream using limited memory, with such guarantees. Prior work on correlated aggregates either did not have accuracy guarantees on the results [8] or else did not allow time-decay [1]. We first define the stream model and the problem more precisely, and then present our results.

1.1 Problem Formulation. We consider a stream $R = e_1, e_2, \ldots, e_n$. Each element e_i is a (v_i, w_i, t_i) tuple, where $v_i \in [m]$ is a value or key from an ordered domain; positive integer w_i is the initial weight; and t_i is the timestamp at which the element was created or observed, also assumed to be a positive integer. For example, in a stream of VoIP call records, there is one stream element per call, where t_i is the time the call was placed, v_i is the duration of the call, and w_i the packet loss rate.

In the synchronous model, the arrivals are in timestamp order: $t_i \leq t_{i+1}$, for all i. In the strictly synchronous version, there is exactly one arrival at each time step, so that $t_i = i$. We also consider the general asynchronous streams model [12, 2, 5], where the order of receipt of stream elements is not necessarily the same as the increasing order of their timestamps. Therefore, it is possible that $t_i > t_j$ while $i < j$. Such asynchrony is inevitable in many applications: in the VoIP call example, CDRs maybe collected at different switches, and in sending this data to a central server, the interleaving of many (possibly synchronous) streams could result in an asynchronous stream.

The aggregates will be time-decayed, i.e. elements with earlier timestamps will be weighted lower than elements with more recent timestamps. The exact model of decay is specified by the user through a *time-decay* function.

DEFINITION 1.1. *A function $f(x)$, $x \geq 0$, is called a time-decay function (or just a decay function) if: (1) $0 \leq f(x) \leq 1$ for all $x > 0$; (2) if $x_1 \leq x_2$, then $f(x_1) \geq f(x_2)$.*

At time $t \geq t_i$ the *age* of $e_i = (v_i, w_i, t_i)$ is defined as $t - t_i$, and the *decayed weight* of e_i is $w_i \cdot f(t - t_i)$.

Time-Decayed Correlated Sum. The query for the time-decayed correlated sum over stream R under a prespecified decay-function f is posed at time t, provides a parameter $\tau \geq 0$, and asks for S_τ^f, defined as follows:

$$S_\tau^f = \sum_{e_i \in R | v_i \geq \tau} w_i \cdot f(t - t_i)$$

A correlated aggregate query could be: "What is the average packet loss rate for all calls which started in the last 24 hours, and were more than 30 minutes in length?". This query can be split into two sub-queries: The first sub-query finds the number of stream elements (v_i, w_i, t_i) which satisfy $v_i > 30$, and $t_i > t - 24$ where t is the current time in hours. The second sub-query finds the sum of w_is for all elements

(v_i, w_i, t_i) such that $v_i > 30$ and $t_i > t - 24$. The average is the ratio of the two answers.

Other correlated aggregates can also be reduced to the sum:

- The time decayed *relative frequency* of a value v is given by $(S_v^f - S_{v+1}^f)/S_0^f$.
- The *sum of decayed weights* of elements in the range $[l, r]$ is $S_l^f - S_{r+1}^f$.
- The *decayed frequency* of range $[l, r]$ is $(S_l^f - S_{r+1}^f)/S_0^f$.
- The *time decayed ϕ-heavy hitters* is found by a binary search over ranges from the universe $[m]$ to find all the v's, such that the time decayed relative frequency of v is at least ϕ.
- The time decayed correlated ϕ-quantile is found by a binary search over the universe $[m]$ to find the largest v, such that $(S_0^f - S_v^f)/S_0^f \leq \phi$.

Time-Decayed Correlated Count. An important special case of DCS is the *time-decayed correlated count* (DCC), where all the weights w_i are assumed to be 1. The correlated count C_τ^f is therefore: $C_\tau^f = \sum_{i:v_i \geq \tau} f(t - t_i)$.

1.2 Decay Functions Classes. We define classes of decay functions, which cover popular time decays from prior work.

Converging decay. A decay function $f(x)$ is a converging decay function if $f(x+1)/f(x)$ is non-decreasing with x. Intuitively, the relative weights of elements with different timestamps under a converging decay function get closer to each other as time goes by. As pointed out by Cohen and Strauss [4], this is an intuitive property of a time-decay function in several applications. Many popular decay functions, such as polynomial decay: $f(x) = (x+1)^{-a}$ where $a > 0$, are converging decay functions.

Exponential decay. Given a constant $\alpha > 0$, the exponential decay function is defined as $f(x) = 2^{-\alpha x}$. Other exponential decays can be written in this form, since $a^{-\lambda x} = 2^{-\lambda \log_2(a)x}$. As $f(x+1)/f(x)$ is a constant, exponential decay qualifies as a converging decay function.

Finite decay. A decay function f is defined to be a *finite decay function* with age limit N, if there exists $N \geq 0$ such that for $x > N$, $f(x) = 0$, and for $x \leq N$, $f(x) > 0$. Examples of finite decay include (1) sliding window decay: $f(x) = 1$ for $x \leq N$ and 0 otherwise, where the age limit N is the window size. (2) Chordal decay [4]: $f(x) = 1 - x/N$ for $0 \leq x \leq N$ and 0 otherwise, with an age limit of $N - 1$. Obviously, no finite decay function is a converging decay function, since $f(N+1)/f(N) = 0$ while $f(N)/f(N-1) > 0$.

1.3 Contributions. Our main result is that there exist small space algorithms for approximating DCS over an arbitrary decay function f with a small *additive* error. But,

the space cost of approximating DCS with a small *relative* error depends strongly on the nature of the decay function— this is possible on some classes of functions using small space, while for other classes, including sliding window and exponential decay, this is provably impossible in sublinear space. More specifically, we show:

1. For *any* decay function f, there is a randomized algorithm for approximating DCS with additive error which uses space logarithmic in the size of the stream. This significantly improves on previous work [8], which presented heuristics only for sliding window decay. (§3.1)

2. On the other hand, for any *finite* decay function, we show that approximating DCS with a small *relative* error needs space linear in the size of the elements within the sliding window. Because sliding window decay is a finite decay function, the above two results resolves the open problem posed in [1], of the space complexity of approximating the correlated sum under sliding window decay. (§4.1)

3. For any sub-exponential converging decay function, there is an algorithm for approximating DCS to within a small relative error using space logarithmic in the stream size, and logarithmic in the "rate" of the decay function. (§3.2)

4. For any exponential decay function, we show that the space complexity of approximating DCS with a small relative error is linear in the stream size, in the worst case. This may be surprising, since there are simple and efficient solutions for maintaining exponentially decayed sum exactly in the non-correlated case. (§4.2)

We evaluate our techniques over real and synthetic data in §5, and observe that they can effectively summarize massive streams in tens of kilobytes.

2 Prior Work

Concepts of correlated aggregation in the (non-streaming) OLAP context appear in [3]. The first work to propose correlated aggregation for streams was Gehrke *et al.* [8]. They assumed that data was locally uniform to give heuristics for computing the non-decayed correlated sum where the threshold (τ) is either an extrema (min, max) or the mean of the all the received values (v_i's). For the sliding window setting, they simply partition the window into fixed-length intervals, and make similar uniformity assumptions for each interval. None of these approaches provide any strong guarantee on the answer quality. Subsequently, Ananthakrishna *et al.* [1] presented summaries that estimate the non-decayed correlated sum with *additive error* guarantees. The problem of tracking sliding window based correlated sums with quality guarantees was given as an open problem in [1]. We show that this relative error guarantees are not possible while using small space, whereas additive guarantees can be obtained.

Xu *et al.* [12] proposed the concept of asynchronous streams. They gave a randomized algorithm to approximate the sum and the median over sliding windows. Busch and Tirthapura [2] later gave a deterministic algorithm for the sum. Cormode *et al.* [6, 5] gave algorithms for general time decay based aggregates over asynchronous streams. By defining timestamps appropriately, *non-decayed* correlated sum can be reduced to the sum of elements within a sliding window over an asynchronous stream. As a result, relative error bounds follow from bounds in [6, 5]. But these methods do not extend to accurately estimating DCS or DCC.

Datar *et al.* [7] presented a bucket-based technique called *exponential histograms* for sliding windows on synchronous streams. This approximates counts and related aggregates, such as sum and ℓ_p norms. Gibbons and Tirthapura [9] improved the worst-case performance for counts using a data structure called *wave*. Going beyond sliding windows, Cohen and Strauss [4] formalized time-decayed data aggregation, and provided strong motivating examples for non-sliding window decay. All these works emphasized the time decay issue, but did not consider the problems of correlated aggregate computation.

3 Upper Bounds

In this section, we present algorithms for approximating DCS over a stream R. The main results are: (1) For an arbitrary decay function f, there is a small space streaming algorithm to approximate S_τ^f with a small additive error. (2) For any *converging* decay function f, there is a small space streaming algorithm to approximate S_τ^f with relative error.

3.1 Additive Error. A predicate $P(v,w)$ is a 0-1 function of v and w. The time-decayed selectivity Q of a predicate $P(v,w)$ on a stream R of (v,w,t) tuples is defined as

$$Q = \frac{\sum_{(v,w,t)\in R} P(v,w) \cdot w \cdot f(c-t)}{\sum_{(v,w,t)\in R} w \cdot f(c-t)}$$

The decayed sum S is defined as $S = \sum_{(v,w,t)\in R} w \cdot f(c-t)$

Note that $S = S_0^f$. We use the following results on time-decayed selectivity estimation from [6] in our algorithm for approximating DCS with a small additive error.

THEOREM 3.1. (THEOREMS 4.1, 4.2, 4.3 FROM [6])
Given $0 < \varepsilon < 1$ and probability $0 < \delta < 1$, there exists a small space sketch of size $O(\frac{1}{\varepsilon^2} \cdot \log \frac{1}{\delta} \cdot \log M)$ that can be computed in one pass from stream R, where M is an upper bound on S. For any decay function f given at query time: (1) the sketch can return an estimate $\widehat{S}$ for S such that $\Pr[|\widehat{S} - S| \le \varepsilon S] \ge 1 - \delta$. (2) Given predicate $P(v,w)$, the sketch gives an estimate $\widehat{Q}$ for the decayed selectivity Q, such that $\Pr[|\widehat{Q} - Q| \le \varepsilon] \ge 1 - \delta$.

This result allows DCS to be additively approximated:

THEOREM 3.2. *For an arbitrary decay function f, there exists a small space sketch of R that can be computed in one pass over the stream. At any time instant, given a threshold τ, the sketch can return $\widehat{S}_\tau^f$, such that $|\widehat{S}_\tau^f - S_\tau^f| \leq \varepsilon S_0^f$ with probability at least $1 - \delta$. The space complexity of the sketch is $O(\frac{1}{\varepsilon^2} \log \frac{1}{\delta} \cdot \log M)$, where M is an upper bound on S_0^f.*

Proof. We run the sketch algorithm in [6] on stream R, with approximation error $\varepsilon/3$ and failure probability $\delta/2$. Let this sketch be denoted by $\mathscr{K}$. To simplify the notation, assume f is fixed, and let $\widehat{S}_\tau, S_\tau$ denote $\widehat{S}_\tau^f, S_\tau^f$ respectively.

Given τ at query time, we define a predicate P for the selectivity estimation as: $P(v, w) = 1$, if $v \geq \tau$, and $P(v, w) = 0$ otherwise. The selectivity of P is $Q = S_\tau/S$. Then $\mathscr{K}$ can return estimates $\widehat{Q}$ of Q and $\widehat{S}$ of S such that

(3.1) $\Pr[|\widehat{Q} - Q| > \varepsilon/3] \leq 1 - \delta/2$

(3.2) $\Pr[|\widehat{S} - S| > \varepsilon S/3] \leq 1 - \delta/2$

Our estimate $\widehat{S}_\tau$ is given by $\widehat{S}_\tau = \widehat{S} \cdot \widehat{Q}$. From (3.1) and (3.2), and using the union bound on probabilities, we get that the following events are both true, with probability at least $1 - \delta$.

(3.3) $Q - \varepsilon/3 \quad \leq \widehat{Q} \quad \leq Q + \varepsilon/3$

(3.4) $S(1 - \varepsilon/3) \quad \leq \widehat{S} \quad \leq S(1 + \varepsilon/3)$

Using the above, and using $Q = S_\tau/S$, we get

$$
\begin{aligned}
\widehat{S}_\tau &\leq \left(\frac{S_\tau}{S} + \varepsilon/3 \right) \cdot S \cdot (1 + \varepsilon/3) \\
&= S_\tau + \frac{S_\tau \varepsilon}{3} + S \left(\frac{\varepsilon}{3} + \frac{\varepsilon^2}{9} \right) \leq S_\tau + \varepsilon S
\end{aligned}
$$

In the last step of the above inequality, we have used the fact $S_\tau \leq S$ and $\varepsilon < 1$. Similarly, we get that if (3.3) and (3.4) are true, then, $\widehat{S}_\tau \geq S_\tau - \varepsilon S$, thus completing the proof that $\mathscr{K}$ can (with high probability) provide an estimate $\widehat{S}_\tau^f$ such that $|\widehat{S}_\tau^f - S_\tau^f| \leq \varepsilon S_0^f$ ∎

An important feature of this approach, made possible due to the flexibility of the sketch in Theorem 3.1, is that it allows the decay function f to be specified at query time, i.e. after the stream R has been seen. This allows for a variety of decay models to be applied in the analysis of the stream after the fact. Further, since the sketch is designed to handle asynchronous arrivals, the timestamps can be arbitrary and arrivals do not need to be in timestamp order.

3.2   Relative Error. In this section, we present a small space sketch that can be maintained over a stream R with the following properties. For an arbitrary *converging* decay function f which is known beforehand, and a parameter τ which is provided at query time, the sketch can return an

estimate $\widehat{S}_\tau^f$ which is within a small relative error of S_τ^f. The space complexity of the sketch depends on f.

The idea behind the sketch is to maintain multiple data structures each of which solves the undecayed correlated sum, and partition stream elements across different data structures, depending on their timestamps, following the approach of the Weight-Based Merging Histogram (WBMH), due to Cohen and Strauss [4]. In the rest of this section, we first give high level intuition, followed by a formal description of the sketch, and a correctness proof. Finally, we describe enhancements that allow faster insertion of stream elements into the sketch.

3.2.1   Intuition. We first describe the weight-based merging histogram. The histogram partitions the stream elements into buckets based on their age. Given a decay function f, and parameter ε_1, the sequence $b_i, i \geq 0$ is defined as follows: $b_0 = 0$, and for $i > 0$, b_i is defined as the largest integer such that $f(b_i - 1) \geq \frac{f(b_{i-1})}{1 + \varepsilon_1}$.

For simplicity, we first describe the algorithm for the case of a (strictly) synchronous stream, where the timestamp of a stream element is just its position in the stream. We later discuss the extension to asynchronous streams. Let G_i denote the interval $[b_i, b_{i+1})$ so that $|G_i| = b_{i+1} - b_i$. Once the decay function $f()$ is given, the G_is are fixed and do not change with time. The elements of the stream are grouped into regions based on their age. For $i \geq 0$, region i contains all stream elements whose age lies in interval G_i.

For any i, we have $f(b_i) < \frac{f(b_0)}{(1 + \varepsilon_1)^k}$, and thus we get $i < \log_{1 + \varepsilon_1} \left(\frac{f(0)}{f(b_i)} \right)$. Since the age of an element cannot be more than n, $b_i \leq n$. Thus we get that the total number of regions is no more than $\beta = \lceil \log_{1 + \varepsilon_1} \left(\frac{f(0)}{f(n)} \right) \rceil$. From the definition of the b_is, we also have the following fact.

FACT 3.1. *Suppose two stream elements have ages a_1 and a_2 so that a_1 and a_2 fall within the same region. Then,*

$$
\frac{1}{1 + \varepsilon_1} \leq \frac{f(a_1)}{f(a_2)} \leq 1 + \varepsilon_1
$$

The data structure maintains a set of *buckets*. Each bucket groups together stream elements whose timestamps fall in a particular range, and maintains a small space summary of these elements. We say that the bucket is "responsible" for this range of timestamps (or equivalently, a range of ages).

Suppose that the goal was to maintain S_0^f, just the time-decayed sum of all stream elements. If the current time c is such that $c \mod b_1 = 0$, then a new bucket is created for handling future elements. The algorithm ensures that the number of buckets does not grow too large through the following rule: if two adjacent buckets are such that the age ranges that they are responsible for are both contained within the same region, then the two buckets are merged into

Figure 1: Weight-based merging histograms.

a single bucket. The count within the resulting bucket is equal to the sum of the counts of the two buckets, and the resulting bucket is responsible for the union of the ranges of timestamps the two buckets were responsible for (see Figures 1(b) and 1(c)).

Due to the merging, there can be at most 2β buckets: one bucket completely contained within each region, and one bucket straddling each boundary between two regions. From Fact 3.1, the weights of all elements contained within a single bucket are close to each other, and since f is a converging decay function, this remains true as the ages of the elements increase. Consequently, WBMH can approximate S_0^f with ε_1 relative error by treating all elements in each bucket as if they shared the smallest timestamp in the range, and scaling the corresponding weight by the total count.

However, this does not solve the more general DCS problem, since it does not allow filtering out elements whose values are smaller than τ. We extend the above data structure to the DCS problem by embedding within each bucket a data structure that can answer the (undecayed) correlated sum of all elements that were inserted into this bucket. This data structure can be any of the algorithms that can estimate the sum of elements within a sliding window on asynchronous streams, including [12, 5, 2]: values of the elements are treated as timestamps, and a window size $m - \tau$ is supplied at query time (where m is an upper bound on the value).

These observations yield our new algorithm for approximating S_τ^f. We replace the simple count for each bucket in the WBMH with a small space sketch, from any of [12, 5, 2]. We will not assume a particular sketch for maintaining the information within a bucket. Instead, our algorithm will work with any sketch that satisfies the following properties—we call such a sketch a "bucket sketch". Let ε_2 denote the accuracy parameter for such a bucket sketch.

1. The bucket sketch must concisely summarize a stream of positive integers using space polylogarithmic in the stream size. Given parameter $\tau \geq 0$ at query time, the sketch must return an estimate of the number of stream elements greater than or equal to τ, such that relative error of the estimate is within ε_2.
2. It must be possible to merge two bucket sketches easily into a single sketch. More precisely, suppose that S_1 is the sketch for a set of elements R_1 and S_2 is the sketch for a set of elements R_2, then it must be possible to

merge together S_1 and S_2 to get a single sketch denoted by $S = S_1 \cup S_2$, such that S retains Property 1 for the set of elements $R_1 \cup R_2$.

The analysis of the sketch proposed in [12] explicitly shows that the above properties hold. Likewise, the sketch designed in [5] also has the necessary properties, since it is built on top of the q-digest summary [11] which are themselves mergable. The different sketches have slightly different time and space complexities; we state and analyze our algorithm in terms of a generic bucket sketch, and subsequently describe the cost depending on the choice of sketch.

3.2.2 Formal Description and Correctness. Recall that ε is the required bound on the relative error. Our algorithm combines two data structures: bucket sketches, with accuracy parameter $\varepsilon_2 = \varepsilon/2$; and the WBMH with accuracy parameter $\varepsilon_1 = \varepsilon/2$. The initialization is shown in the SET-BOUNDARIES procedure (Figure 2), which creates the regions G_i by selecting $b_0, \ldots, b_\beta$. For simplicity of presentation, we have assumed that the maximum stream length n is known beforehand, but this is not necessary — the b_i's can be generated incrementally, i.e., b_i does not need to be generated until element ages exceeding b_{i-1} have been observed.

Figure 3 shows the PROCESSELEMENT procedure for handling a new stream element. Whenever the current time t satisfies $t \mod b_1 = 0$, we create a new bucket to summarize the elements with timestamps from t to $t + b_1 - 1$ and seal the last bucket which was created at time $t - b_1$. The procedure FINDREGIONS(t) returns the set of regions that contain buckets to be merged at time t. In the next section we present novel methods to implement this requirement efficiently. Figure 4 shows the procedure RETURNAPPROXIMATION which generates the answer for a query for S_τ^f. For each bucket, we multiply the common decayed weight with the sliding windowed count using τ as the left boundary of the window, then return the summation of the products over all the buckets as the estimate for S_τ^f.

THEOREM 3.3. *If f is a converging decay function, for any τ given at any time t, the algorithm specified in Figure 2, 3 and 4 can return $\widehat{S_\tau^f}$, such that $(1 - \varepsilon)S_\tau^f \leq \widehat{S_\tau^f} \leq (1 + \varepsilon)S_\tau^f$.*

Proof. For the special converging decay function where $f(x) \equiv 1$ (no decay), then WBMH has only one region and one bucket. So the algorithm reduces to a single bucket

Algorithm 3.1: SETBOUNDARIES(ε)

---

comment: create $G_0, G_1, \ldots, G_\beta$ using $\varepsilon_1 = \varepsilon/2$
$b_0 \leftarrow 0$;
for $1 \leq i \leq \beta$
  do $b_i \leftarrow \max_x \{x | (1 + \frac{\varepsilon}{2}) f(x-1) \geq f(b_{i-1})\}$
$j \leftarrow -1$;
comment: index of the active bucket for new elements

Figure 2: SETBOUNDARIES routine to initialize regions.

sketch. This sketch can directly provide an $\varepsilon_2 = \varepsilon/2$ relative error guarantee for the estimate of S_τ^f.

The broader case is where $f(x+1)/f(x)$ is non-decreasing with x. Let $\{B_1, \ldots, B_k\}$ be the set of buckets at query time t. Let $R_i \subseteq R$ be the substream that is aggregated into B_i, $1 \leq i \leq k$. Since every stream element is aggregated into exactly one bucket at any time t, the R_is partition R: $\bigcup_{i=1}^k R_i = R$ and $R_i \cap R_j = \emptyset$ if $i \neq j$. Note that merging two buckets just creates a new bucket over the union of the two underlying substreams. Let $S_{\tau,i}^f = \sum_{v_j \in R_i, v_j \geq \tau} w_j f(t-j)$, be the time decayed correlated sum over the substream R_i, $1 \leq i \leq k$, so $S_\tau^f = \sum_{i=1}^k S_{\tau,i}^f$. Now we consider the accuracy of the approximation for $S_{\tau,i}^f$ using bucket sketch B_i, for each i in turn.

Note that at query time t, the common decayed weight of B_i is $f(t - F_{B_i})$. Let w_t^v be the true decayed weight of any element v aggregated in S_i, then due to the setting of the regions in WBMH, we have $\frac{1}{1+\varepsilon_1} w_t^v \leq f(t - F_{S_i}) \leq w_t^v$ (Fact 3.1). Let $|\{v \in R_i | v \geq \tau\}| = Q_i$, then summing over all elements in the bucket i we have:

$$\frac{1}{1+\varepsilon_1} S_{\tau,i}^f = \frac{1}{1+\varepsilon_1} \sum_{v \in R_i} w_t^v$$
$$\leq Q_i \cdot f(t - F_{S_i}) \leq \sum_{v \in R_i} w_t^v = S_{\tau,i}^f.$$

Further, bucket sketch S_i can return $\widehat{Q}_i$ such that [12, 5]
$$(1 - \varepsilon_2) Q_i \leq \widehat{Q}_i \leq (1 + \varepsilon_2) Q_i.$$
Combined with the above inequality, we have
$$\frac{1 - \varepsilon_2}{1 + \varepsilon_1} S_{\tau,i}^f \leq \widehat{Q}_i \cdot f(t - F_{S_i}) \leq (1 + \varepsilon_2) S_{\tau,i}^f.$$
Adding up all the $S_{\tau,i}^t$ over $i = 1, 2, \ldots, k$, we get
$$\frac{1 - \varepsilon_2}{1 + \varepsilon_1} \sum_{i=1}^k S_{\tau,i}^f \leq \sum_{i=1}^k \widehat{Q}_i \cdot f(t - F_{S_i}) \leq (1 + \varepsilon_2) \sum_{i=1}^k S_{\tau,i}^f.$$
Using the fact that $0 < \varepsilon < 1$ and $\varepsilon_1 = \varepsilon_2 = \varepsilon/2$, along with
$$S_\tau^f = \sum_{i=1}^k S_{\tau,i}^f \quad \text{and} \quad \widehat{S}_\tau^f = \sum_{i=1}^k \widehat{Q}_i \cdot f(t - F_{S_i})$$
we conclude that $(1 - \varepsilon) S_\tau^f \leq \widehat{S}_\tau^f \leq (1 + \varepsilon) S_\tau^f$. ∎

Algorithm 3.2: PROCESSELEMENT((v_i, w_i, i))

---

if $i \mod b_1 = 0$
  then $\begin{cases} j \leftarrow j + 1 \\ \text{Initialize a new bucket sketch } B_j \text{ with accuracy } \varepsilon/2 \\ F_{B_j} \leftarrow i \\ L_{B_j} \leftarrow i + b_1 - 1 \\ \text{\textbf{comment:} Set timestamp range covered by } B_j \end{cases}$
Insert (v_i, w_i) into B_j
for each $g \in$ FINDREGIONS(i)
comment: Set of regions with buckets to be merged at time i
  do $\begin{cases} b_{min} \leftarrow \min_t \{t | t \in G_g\} \\ b_{max} \leftarrow \max_t \{t | t \in G_g\} \\ \text{Find buckets } B' \text{ and } B'', \text{ such that} \\ \quad b_{min} \leq \ (t - L_{B'}) < (t - F_{B'}) \\ \quad\quad < \ (t - L_{B''}) < (t - F_{B''}) \leq b_{max} \\ \text{\textbf{comment:} Find buckets covered by } G_g \\ B \leftarrow B' \cup B'' \\ \text{\textbf{comment:} merge two buckets} \\ F_B \leftarrow F_{B''} \\ L_B \leftarrow L_{B'} \\ \text{Drop } B' \text{ and } B'' \end{cases}$

Figure 3: PROCESSELEMENT routine to handle updates

3.2.3 Fast Bucket Merging. At every time tick the histogram maintenance algorithm needs to merge buckets that are covered by a single region. In the synchronous stream case, this occurs with every new element arrival. The naive solution is to pass over all buckets and merge any pair falling in the same region on every update. This procedure can severely reduce the speed of stream processing. In this section, we present an algorithm, which directly returns the set of regions that have buckets to be merged at each time t.

DEFINITION 3.1. (BUCKET B'S CAPACITY $|B|$) $|B| = L_B - F_B + 1$, where L_B and F_B are the largest and smallest timestamps that are covered by B (see Figure 3).

Recall that no pair of buckets overlap in the time ranges that they cover. Therefore we have $|B' \cup B''| = |B'| + |B''|$, where B' and B'' are any two buckets in the histogram and $\cup$ is the merging operation on buckets B' and B''. Now consider the simple case where all boundaries are powers of two (i.e. $b_1 = 1, b_2 = 2, b_3 = 4$ and so on). Here, all capacities are also powers of two, and the merging of buckets has a very regular structure: whenever two buckets fit exactly into a region, they are merged. It turns out that the same concept generalizes to arbitrary patterns of growing regions. With the help of Figure 1, we can visualize the buckets traveling through the regions along the age axis, being merged when necessary. For region G_i, let I_i be the capacity of a bucket entering G_i. More formally,

Algorithm 3.3: RETURNAPPROXIMATION(τ)

Let the set of buckets be: $\{B_1, B_2, \ldots, B_k\}$
comment: for some k, $1 \le k \le 2\beta$;
$s \leftarrow 0$;
for $1 \le i \le k$
$\quad$ **do** $\begin{cases} \text{Let } \widehat{Q}_i \text{ be result for } B_i \text{ using } \tau \text{ as window size} \\ s \leftarrow s + \widehat{Q}_i \cdot f(t - F_{B_i}); \end{cases}$
comment: Approx sum of element weights in B_i with $v_i \ge \tau$
return $(\widehat{S}_\tau^f = s)$;

Figure 4: RETURNAPPROXIMATION routine to estimate S_τ^f

DEFINITION 3.2. (REGION i'S CAPACITY I_i) *Define* $I_0 = 1$. *For* $0 < i < \beta$, *let* $I_i = |S|$, *where S is any bucket such that $t - F_S = b_i$ for some value of t.*

In the next lemma, we show that for any specific i, there is a fixed value of I_i: it does not vary over time, and can be easily computed as a function of the region sizes.

LEMMA 3.1. *For* $0 < i < \beta$, $I_i = \lfloor |G_{i-1}|/I_{i-1} \rfloor \cdot I_{i-1}$

Proof. The lemma is proved by induction. For the base case, since the capacity of the new bucket created in G_0 is exactly equal to $|G_0|$, merging cannot happen in G_0. Thus immediately $I_1 = |G_0| = \lfloor |G_0|/I_0 \rfloor \cdot I_0$. For the inductive step, suppose the claim is true for some i. Then, for region i, all buckets entering G_i have the same (constant) size I_i. Exactly $\lfloor |G_i|/I_i \rfloor$ such buckets of size I_i can be merged together within G_i before the "leading edge" of the merged bucket crosses into G_{i+1}. After the bucket of size $\lfloor |G_i|/I_i \rfloor \cdot I_i$ is formed, no further buckets of size I_i can be merged with it in region G_i, so it crosses into G_{i+1}. This procedure repeats, and since $|G_i|$ and I_i are constants, I_{i+1} is fixed as $\lfloor |G_i|/I_i \rfloor \cdot I_i$. This completes the induction. ∎

In the next lemma, we show that given I_i we can compute the times at which G_i has buckets to be merged.

LEMMA 3.2. *For* $0 \le i < \beta$, *the times at which G_i has buckets to be merged is given by $\{b_i + (k \lfloor |G_i|/I_i \rfloor + j)I_i\}$ for integers $2 \le j \le \lfloor |G_i|/I_i \rfloor$ and $k \ge 0$.*

Proof. The new bucket created in G_0 has capacity equal to $|G_0|$, so G_0 does not have any buckets to be merged at any time. For $i > 0$, if $\lfloor |G_i|/I_i \rfloor < 2$, then G_i will not have the chance to have two buckets of size I_i to be merged at any time. Now we consider the case where $\lfloor |G_i|/I_i \rfloor \ge 2$ and $i > 0$. G_i obtains its first whole incoming bucket at time $t = b_i + I_i$. Note that within G_i at most $\lfloor |G_i|/I_i \rfloor$ buckets that enter G_i can be merged together. Thus, (1) at time $t = b_i + 2I_i, b_i + 3I_i, \ldots, b_i + \lfloor |G_i|/I_i \rfloor \cdot I_i$, buckets can be

Algorithm 3.4: INITIALIZEFINDREGIONS()

Initialize hash table T
$I_0 \leftarrow 1$;
for $1 \le i \le \beta - 1$
$\quad$ **do** $I_i \leftarrow \lfloor |G_{i-1}|/I_{i-1} \rfloor I_{i-1}$
comment: From Lemma 3.1
for $1 \le i \le \beta - 1$
$\quad$ **do if** $\lfloor |G_i|/I_i \rfloor \ge 2$
$\quad$ **then** Insert $(i, b_i + 2I_i)$ into hash table T
comment: Compute when G_i first has mergable buckets

Figure 5: Routine to initialize hash table with merging times

merged within G_i; (2) This sequence of merging operations repeat every $\lfloor |G_i|/I_i \rfloor \cdot I_i$ clock ticks, meaning G_i has buckets to be merged at times $\{b_i + (k \lfloor |G_i|/I_i \rfloor + j)I_i\}$ for integers $2 \le j \le \lfloor |G_i|/I_i \rfloor$ and $k \ge 0$. ∎

Lemma 3.2 provides a way for any region to directly compute the sequence of time points at which there are buckets to be merged. Based on this observation, we present an algorithm to return the set of regions that have buckets to be merged at a given time t.

Algorithm for Fast Bucket Merging. The algorithm for tracking which buckets should be merged makes use of a hash table T to store the set of buckets to be merged at timestamp t. More precisely, the table cell corresponding to time t is a set of (i, t) pairs, such that region G_i has buckets to be merged at time t. Figure 5 shows procedure INITIALIZEFINDREGIONS() which first computes I_i using Lemma 3.1. It then uses Lemma 3.2 to fill in the earliest time at which region G_i will have buckets to be merged. At time t, FINDREGIONS(t) (Figure 6) retrieves the set of buckets to merge, and deletes them from the hash table. Then, for each returned region, we compute its next merging time using Lemma 3.2 and store the results into the corresponding hash table cells for the future lookup.

3.2.4 Time and Space Complexity. The space complexity includes the space cost for the buckets in the histogram and the hash table. The space to represent each bucket depends on the choice of the bucket sketch.

THEOREM 3.4. *The space complexity of the algorithm in Figure 2, 3 and 4 is $O(\beta(\mathscr{Z} + \log n))$ bits, where*

1. $\beta = \left\lceil \log_{1+\varepsilon/2}(f(0)/f(n)) \right\rceil$

2. $\mathscr{Z} = O\left(\frac{1}{\varepsilon^2} \log \frac{\beta}{\delta} \log n \log m\right)$ *using the sketch of [12].*

3. $\mathscr{Z} = O\left(\frac{1}{\varepsilon} \log m \log \left(\frac{\varepsilon n}{\log n}\right)\right)$ *using the sketch of [5].*

Algorithm 3.5: FINDREGIONS(t)

$M \leftarrow \emptyset$
for each $(i,t) \in T$
  comment: Region G has buckets to be merged at time t
  do $\begin{cases} M \leftarrow M \cup \{i\} \\ \textbf{if } (t - b_i)/I_i \mod \lfloor |G_i|/I_i \rfloor = 0 \\ \quad \textbf{then } t' \leftarrow t + 2I_i \\ \quad \textbf{else } t' \leftarrow t + I_i; \\ \textbf{comment: } \text{Find when } G_i \text{ next has mergable buckets} \\ \text{Insert } (i,t') \text{ into hash table } T \end{cases}$
return (M)
comment: set of regions with buckets to be merged at time t

Figure 6: FINDREGIONS(t) finds mergable regions at time t

Proof. The number of buckets used is at most 2β. For the randomized sketch designed in [12], in order to have a δ failure probability bound, by the union bound, we need to set the failure probability for each bucket to be δ/β, so we get $\mathscr{Z} = O\left(\frac{1}{\varepsilon^2} \log \frac{\beta}{\delta} \log n \log m\right)$ (Lemma 11 in [12]). For the deterministic sketch designed in [5], $\mathscr{Z} = O\left(\frac{1}{\varepsilon} \log m \log \left(\frac{\varepsilon n}{\log n}\right)\right)$ (§3.1 in [5]). The size of the hash table can be set to $O(\beta)$ cells, because each of the β regions occupies at most one cell. Each cell uses $O(\log n)$ bits of space to store the region's index and merging time. So all together, the total space cost is $O(\beta(\mathscr{Z} + \log n))$.

THEOREM 3.5. *The (amortized) time complexity of the algorithm per update is linear in the size of the bucket sketch data structure used.*

Proof. The cost of the algorithm is dominated by the cost of merging bucket sketches together when necessary. Inserting a new element into the sketch takes time sublinear in the size of the bucket sketch. Updating the hash table has to be done once for every merge that occurs, and takes constant time. The merge of two bucket sketches can be carried out in time linear in the size of the bucket sketch data structure [12, 5]. So the time is determined by the (amortized) number of merges per clock tick.

The number of merge operations over the course of algorithm can be bounded in terms of the number of updates (for synchronous streams, where there is one arrival per clock tick). Observe that for $\varepsilon < 1$, the set of regions generated will mean that $\lfloor |G_i|/I_i \rfloor \leq 2$ for all i. This is seen by contradiction: suppose that $\lfloor |G_i|/I_i \rfloor > 2$. Then we could have merged two of the buckets of capacity I_i in the preceding region: since $|G_{i-1}| > 2|G_i|/3$ (by choice of ε), $|G_i|/I_i > 2$ implies $|G_{i-1}|/I_i \geq 2$. From this, we see that the bucket capacities must be powers of two, since I_i must be either I_{i-1} or $2I_{i-1}$. By a standard charging argument, each

merge can be charged back to the corresponding insertion of a new stream element. The consequence is that the amortized number of merges per clock tick is bounded by a constant. This implies the stated time bound. ∎

Space dependence on decay function f. As shown in Theorem 3.4, the space complexity depends crucially on decay function f, since it determines the number of regions (implicitly the number of buckets). We show the consequence for various broad classes of decay function:

- For exponential decay functions $f(x) = 2^{-\alpha x}$, $\alpha > 0$, we have $\beta = \alpha n \log_{1+\varepsilon/2} 2$ and therefore the space complexity is $O\left(\frac{n}{\varepsilon^2} \log^2 n\right)$. This means that this algorithm needs space linear in the input size.

- For polynomial decay functions $f(x) = (x+1)^{-a}$, $a > 0$, since $\beta = a \log_{1+\varepsilon/2} n$, we obtain a small space complexity $O\left(\frac{1}{\varepsilon^2} \log^2 n \log m \log \frac{\beta}{\delta}\right)$ using the sketch of [12], and $O(\frac{1}{\varepsilon} \log n \log m \log(\varepsilon n/\log n) + \log^2 n)$ using the sketch of [5];

- In the case of no decay ($f(x) \equiv 1$), the region G_0 is infinitely large, so the algorithm maintains only one bucket, giving space cost $O(\mathscr{Z} + \log n)$.

Intuitively the algorithm can approximate S_τ^f with a relative error bound using small space if f decays more slowly than the exponential decay. Further, the space decreases the "slower" that f decays, the limiting case being that of no decay. We complement this observation with the result that the DCS problem under exponential decay requires linear space in order to provide relative error guarantees.

Asynchronous Streams. So far our discussion of the algorithm for relative error has focused on the case of synchronous streams, where the elements arrive in order of timestamps. In an asynchronous setting, a new element (v_1, w_1, t_1) may have timestamp $t_1 < t$ where t is the current time. But this can easily be handled by the algorithm described above: the new element is just directly inserted into the earlier bucket which is responsible for timestamp t_1. The accuracy and space guarantees do not alter, although the time cost is affected since the correct bucket must be found for each new arrival, and buckets to merge determined.

4 Lower Bounds

This section shows large space lower bounds for finite decay or (super) exponential decay for DCC on synchronous streams. Since DCC is a special case of DCS, these lower bounds also apply to DCS on asynchronous streams.

4.1 Finite Decay. Finite decay, defined in § 1.2, captures the case when after some age N, the decayed weight is zero.

THEOREM 4.1. *For any finite decay function f with age limit N, any streaming algorithm (deterministic or randomized) that can provide $\widehat{C}_\tau^f$ such that $|\widehat{C}_\tau^f - C_\tau^f| < \varepsilon C_\tau^f$ for τ given at query time must store $\Omega(N)$ bits.*

Proof. The bound follows from the hardness of finding the maximum element within a sliding window on a stream of integers. Tracking the maximum within a sliding window of size N over a data stream needs $\Omega(N\log(m/N))$ bits of space, where m is the size of the universe from which the stream elements are drawn (§7.4 of [7]).

We argue that if we could approximate $\widehat{C}_\tau^f$, where f has age limit N, we could also find the maximum of the last N elements in R. Let α denote the value of the maximum element in the last N elements of the stream. By definition, the decayed weights of the N most recent elements are positive, while all older elements have weight zero. Note that C_τ^f is a monotonically decreasing function of τ, so $C_\alpha^f > 0$ (and $C_\tau^f > C_\alpha^f$ for any $\tau < \alpha$) while $C_\tau^f = 0$ for $\tau > \alpha$. If C_τ^f can be approximated with relative error, then we can distinguish the cases $C_\tau^f > 0$ and $C_\tau^f = 0$. By repeatedly querying different values of τ for C_τ^f, we find a value τ^* such that $C_{\tau^*}^f > 0$ and $C_{\tau^*+1}^f = 0$. Then τ^* must be α, the maximum element of the last N elements. ∎

Since sliding window is a special case of finite decay, this shows that approximating C_τ^f (a problem identified in [1]) cannot be solved with relative error in sublinear space.

4.2 Exponential Decay. Exponential decay functions $f(x) = 2^{-\alpha x}$, $\alpha > 0$ are widely used in non-correlated time decayed steaming data aggregation. It is easy to maintain simple sums and counts under such decay efficiently [4]. However, in this section we will show that it is *not* possible to approximate C_τ^f with relative error guarantees using small space if m (the size of the universe) is large and f is exponential decay. This remains true for other classes of decay that are "faster" than exponential decay. We first present two natural approaches to approximate C_τ^f under an exponential decay function f, and analyze their space cost to show that each stores large amounts of information.

Algorithm I. Since tracking sums under exponential decay can be performed efficiently using a single counter, we can just track the decayed count of elements for each $v \in [m]$—denote this as W_v^f. Then C_τ^f can be estimated as $\sum_{v \geq \tau} W_v^f$. To ensure an accurate answer, each W_v^f must be tracked with sufficiently many bits of precision. One way to do this is to maintain the timestamps of the last $\lceil \frac{1}{\alpha}\log_2 \frac{1}{\varepsilon}\rceil$ elements in the substream $R_v = \{v_i \in R | v_i = v\}$. From these, one can compute W_v^f with relative error ε, and hence C_τ^f with the same relative error. Each timestamp is $O(\log n)$ bits, so the total space cost is $O(m\log n\lceil\frac{1}{\alpha}\log\frac{1}{\varepsilon}\rceil)$ bits. ∎

Algorithm II. The second algorithm tries to reduce the dependence on m by observing that for some close values of τ, the value of C_τ^f may be quite similar, so there is potential for "compression". As $f(x) = 2^{-\alpha x}$, $\alpha > 0$, we can write:
$$C_\tau^f = \sum_{v_i \geq \tau} 2^{\alpha(i-t)} = 2^{-\alpha t}\sum_{v_i \geq \tau} 2^{\alpha i},$$
where t is the query time. We reduce approximating C_τ^f with a relative error bound to a counting problem over an asynchronous stream with sliding window queries. We create a new stream R' in this model by treating each stream element as an item with timestamp set to its value v_i and with weight $2^{\alpha i}$. The query C_τ^f at time t can be interpreted as a sliding window query on the derived stream R' at time m with width $m - \tau$. The answer to this query is $\sum_{v_i \geq \tau} 2^{\alpha i}$; by the above equation, scaling this by $2^{\alpha t}$ approximates C_τ^f.

The derived stream R' can be summarized by sketches such as those in [12, 2]. These answer the sliding window query with relative error ε, implying relative error for C_τ^f. But the cost of these sketches applied here is $O(\frac{\alpha n}{\varepsilon}\log^2 m)$ bits: in the reduction, the number of copies of each stream element increases exponentially, and the space cost of the sketches depends logarithmically on this quantity. ∎

Hardness of Exponential Decay. Algorithm I is a conceptually simple approach, which stores information for each possible value in the domain. Algorithm II uses summaries that are compact in their original setting, but when applied to the DCC problem, their space must increase to give an accurate answer for any τ. The core reason for the high space cost of both algorithms is the fact that as τ varies between 0 and m, the value of C_τ^f can vary over an exponentially large range, and a large data structure is required to track so many different values. This is made precise by the next theorem, which shows that the space cost of Algorithm I is close to optimal. We go on to provide a small space sketch with a weakened guarantee in § 4.4, by limiting the range of values of C_τ^f for which an accurate answer is required.

THEOREM 4.2. *For an exponential decay function $f(x) = 2^{-\alpha x}$, $\alpha > 0$ and $\varepsilon \leq 1/2$, any algorithm (deterministic or randomized) that provides $\widehat{C}_\tau^f$ over a stream of size $n = \Theta(m)$, such that $|\widehat{C}_\tau^f - C_\tau^f| < \varepsilon C_\tau^f$ for τ given at query time must store $\Omega(m\log\frac{n}{m})$ bits, where m is the universe size.*

Proof. The proof uses a reduction from the INDEX problem in two-party communication complexity [10]. In the INDEX problem, the first player holds a binary string b of length N, and the second holds an index $i \in [N]$. The first player is allowed to send a single message to the second, who must then output the value of $b[i]$ (the ith bit of string b). Since no communication is allowed from the second player to the first, the size of the message must be $\Omega(N)$ bits, even allowing the protocol a constant probability of failure [10].

(a) Setting the Intervals over a Stream   (b) Mapping from binary string b to intervals

Figure 7: Creating a stream for the lower bound proof using $p = 1$

We show that a sublinear streaming data structure to approximate DCC under exponential decay would allow a sublinear communication protocol for INDEX. Given a binary string b of length mp, we construct an instance of a stream. Here m is the size of the domain of the stream values, and $p \geq 1$ is an integer parameter set later. The n positions in stream R are divided into m intervals: $I_0, I_1, \ldots, I_{m-1}$, as shown in Figure 7(a). Let $\ell = 2\lceil \frac{1}{\alpha} \rceil$; each interval has $2^p \ell$ positions, so that the length of R is $\Omega(m2^p\ell)$.

Every position in the stream is set to 0 by default; the construction places one non-zero element in each interval at a position that is a multiple of ℓ (shaded in Figure 7(a)). We interpret the binary string b as an integer b. Let b_P be that value represented in base $P = 2^p$ (so $b = \sum_i P^i b_P[i] = \sum_j 2^j b[j]$). In interval I_i, we place an element with value i at position $b_P[i]\ell$, shown in Figure 7(b) for $p = 1$. We write

$$C_\tau^f = \sum_{i=\tau}^m f((Pi + b_P[i])\ell)$$

$$\leq \sum_{i=\tau+1}^m f(Pi\ell) + f((P\tau + b_P[\tau])\ell)$$

$$< \sum_{i=1}^\infty f(P(\tau+1)\ell + (i-1)\ell) + f((P\tau + b_P[\tau])\ell)$$

$$\leq \frac{1}{3}f(P(\tau+1)\ell - \ell) + f(P\tau\ell + b_P[\tau]\ell)$$

$$= \frac{1}{3}f(P\tau\ell + (P-1)\ell) + f(P\tau\ell + b_P[\tau]\ell)$$

$$\leq \frac{4}{3}f(P\tau\ell + b_P[\tau]\ell)$$

which follows since $f(x) = 2^{-\alpha x}$ and $\ell = 2\lceil \frac{1}{\alpha} \rceil$. Thus,

(4.5) $\qquad f(P\tau\ell + b_P[\tau]\ell) < C_\tau^f < \frac{4}{3}f(P\tau\ell + b_P[\tau]\ell)$

Denote $b_P[\tau] = j$, where $0 \leq j \leq P-1$ (j is a digit in base P). If C_τ^f can be approximated within a relative error $\varepsilon = \frac{1}{2}$, then j can be retrieved by the data stream algorithm: the approximation of C_τ^f, $\hat{C}$, satisfies
$$\frac{1}{2}f(P\tau\ell + b_P[\tau]\ell) < \hat{C} < 2f(P\tau\ell + b_P[\tau]\ell)$$
Meanwhile, observe that for any $k > j$, $f(P\tau\ell + j\ell)/f(P\tau\ell + k\ell) \geq 4$. As a result, we can distinguish the case of $b_P[\tau] = j$ and $b_P[\tau] = k$ (the case $k < j$ is symmetrical).

To complete the proof, we observe that if we had an algorithm to approximate C_τ^f using small space, the first player could execute the algorithm on the stream derived from their binary string and send the memory contents of the algorithm as a message to the second player. The above analysis allows them to determine the value of C_τ^f for $\tau = \lfloor \frac{i}{p} \rfloor$, from which they can recover $b_P[\tau]$ and hence $b[i]$. The communication lower bound for INDEX is $\Omega(mp)$ bits, which implies that the data structure must also be $\Omega(mp)$ bits. The stream length is $n = O(\frac{m2^p}{\alpha})$, so fixing n sets p, and bounds the space by $\Omega(m\log \frac{n}{m})$ for constant α. We conclude by observing that since the communication lower bound allows randomization, this space lower bound holds for randomized stream algorithms. ∎

4.3 Super-exponential Decay. Theorem 4.2 applies to decay functions f that decay faster than exponential decay. Examples of such decay functions include: (1) *polyexponential decay* [4]: $f(x) = (x+1)^k 2^{-\alpha x}/k!$ where $k > 0$, and $\alpha > 0$ are constants. (2) super-exponential decay: $f(x) = 2^{-\alpha x^\beta}$, where $\alpha > 0$ and $\beta > 1$. We can show:

THEOREM 4.3. *A decay function $f(x)$ is (super)-exponential, if there exist constants $\sigma > 1$ and $c \geq 0$, such that for every $x \geq c$, $f(x)/f(x+1) \geq \sigma$. Any algorithm that can provide $\hat{C}_\tau^f$ for super-exponential f over a stream of size $n = \Theta(m)$, such that $|\hat{C}_\tau^f - C_\tau^f| < \varepsilon C_\tau^f$ must use $\Omega(m)$ bits of space.*

Proof. The argument is based on the proof of Theorem 4.2. When $n \geq m \cdot \lceil \log_\sigma 2 \rceil + c$, we divide the substream from the position $(c+1)$ to the position $(4m \cdot \lceil \log_\sigma 2 \rceil + c)$ into m intervals based on $\ell = 2\lceil \log_\sigma 2 \rceil$ and $p = 1$. By using the construction from Theorem 4.2, the result follows. ∎

4.4 Finite (Super) Exponential Decay. As noted above, the lower bound proof relies on distinguishing a sequence of exponentially decreasing possible values of the DCC. In practical situations, it often suffices to return an answer of zero when the true answer is less than some specified bound μ. This creates a "finite" version of exponential decay.

DEFINITION 4.1. *A decay function f is a finite exponential*

(a) Accuracy for world cup data set

(b) Accuracy for synthetic data set

(c) Throughput versus space on both data sets

Figure 8: Throughput and accuracy with sliding window decay, additive error.

decay function if $f(x) = 2^{-\alpha x}$, $\alpha > 0$ *when* $2^{-\alpha x} \geq \mu$ *for* $0 < \mu < 1$; *and* $f(x) = 0$, *otherwise.*

Since finite exponential decay is a finite decay, the lower bound of Theorem 4.1 implies that $\Omega(\frac{1}{\alpha} \log \frac{1}{\mu})$ space is needed to approximate C_τ^f for such an f. A simple algorithm for C_τ^f simply stores all stream elements with non-zero decayed weight. The space used for a synchronous stream is $O(\frac{1}{\alpha} \log m \log \frac{1}{\mu})$ bits, which is (nearly) optimal (treating $\log m$ as a small constant). This approach extends to the finite versions of super-exponential decay.

4.5   Sub-exponential decay. For any decay function $f(x)$, where $f(x) > 0$ and $\lim_{x \to \infty} f(x) = 0$, we can always find m positions (timestamps) in the stream: $0 \leq x_1 < x_2 < \ldots < x_m$, such that for every i, $1 < i \leq m$, we have $f(t - x_{i-1})/f(t - x_i) \leq \frac{1}{2}$. Thus, it is natural to analyze what happens when we apply the construction from the lower bound in Theorem 4.2 to streams under such functions. Certainly, the same style of argument constructs a stream that forces a large data structure. But, if we fix some m and set $p = 1$, the stream has to be truly enormous to imply a large space lower bound: e.g., for the polynomial decay function $f(x) = (x+1)^{-a}$, $a > 0$, we need $n \geq 2^{m/\alpha}$ to force $\Omega(m)$ space. This is in agreement with the upper bounds in §3.2 which gave algorithms which depend logarithmically on n: for such truly huge values of n, this leads to a requirement of $\log 2^{m/\alpha} = \Omega(m)$, so there is no contradiction.

5   Experiments

We present results from an experimental evaluation of the algorithms on two data sets. The first was web traffic logs from the 1998 World Cup on June 19th (the 'worldcup' data set) from `http://ita.ee.lbl.gov/`. Each stream element was a tuple (v, w, t), where v was the client id, w the packet size modulo 100, and t the timestamp. The dataset had 33695769 elements. The second was a synthetically generated data set (the 'synthetic' data set). The size of the

synthetic data is the same as the worldcup data set. Here, the timestamp of an element is a random number chosen uniformly from the range $[1, \max_t]$ where $\max_t = 898293600$ is the maximum timestamp in the world cup data set. The value v is chosen uniformly from the range $[1, \max_v]$, where $\max_v = 1823218$ is the maximum value in the worldcup data set. The weight is chosen similarly, i.e. uniformly from the range $[1, \max_w]$ where $\max_w = 99$ is the maximum weight in the world cup data.

We implemented our algorithms using C++/STL and all experiments were performed on a SUSE Linux Laptop with 1GB memory. Both input streams were asynchronous, and elements do not arrive in timestamp order.

Additive Error. We implemented the algorithm for additive error (§3.1) using the sketch in [12] as the basis. On the sketch, queries were made for the correlated sum S_τ^f where f was the sliding decay function with window size $4.5 \cdot 10^7$ for the synthetic data, and 3600 for the worldcup data. We tried a range of values of the threshold τ, from the 5 percent quantile (5th percentile) of the values of stream elements to the 95 percent quantile. We analyzed the accuracy of the estimates returned by the sketch, for a given space budget.

Figures 8(a) and 8(b) show the observed additive error as a function of the space used by the algorithm for different values of τ. The space cost is measured in the number of nodes, where each node is the space required to store a single stream element (v, w, t), which takes a constant number of bytes. This cost can be compared to the naive method which stores all input elements (nearly 34 million nodes). The observed error is usually significantly smaller than the guarantee provided by theory. The theoretical guarantee holds irrespective of the value of τ or the window size. Note that the additive error decreased as the square root of the space cost, as expected. Figure 8(c) shows the throughput, which is defined as the number of stream elements processed per second, as a function of the space used. From the results, the trend is for the throughput to decrease slowly as the space increases. Across a wide range of values for the space, the
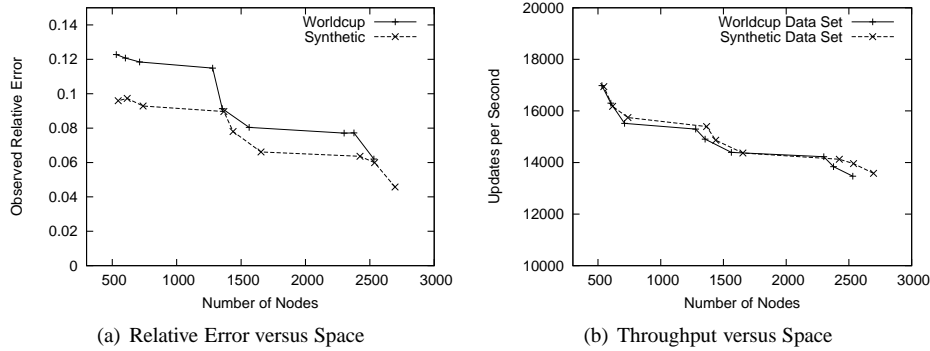
(a) Relative Error versus Space

(b) Throughput versus Space

Figure 9: Performance of Relative Error Algorithm, with Polynomial Decay.

throughput is between 250K and 350K updates per second.

Relative Error. We performed similar experiments to test our algorithms for relative error, based on the polynomial decay function $f(x) = 1/(x+1)^{1.5}$. The thresholds are the same as in the additive error algorithm. The results are shown in Figure 9. In general, the space cost for a given error for polynomial decay was much smaller than the algorithm for sliding windows (Figure 9(a)). This greater space efficiency comes at some cost: we have to fix the decay function *apriori*—the additive error result allows the decay function to be specified at query time. The throughput for the relative error algorithm is also appreciably lower than the additive error algorithm (Figure 9(b)), by over an order of magnitude. This is partly due to the greater time complexity of the relative error algorithm caused by the periodic bucket merging operations which access every node in the merged buckets, and partly because our implementation is not fully tuned.

6 Concluding Remarks

Our results shed light on the problem of computing correlated sums over time-decayed streams. The upper bounds are quite strong, since they apply to asynchronous streams with arbitrary timestamps. It is also possible to extend these results to a distributed streaming model, since the summarizing data structures used can naturally be computed over distributed data, and merged together to give a summary of the union of the streams. The lower bounds are similarly strong, since they apply to the most restricted model, for computing DCC where there is exactly one arrival per time unit.

The correlated sum is at the heart of many correlated aggregates, but there are other natural correlated computations to consider which do not follow immediately from DCS. Some we expect to be hard in general: correlated maximum $\max_{v_i > \tau} w_i f(t - t_i)$ has a linear space lower bound under finite decay functions, since this lower bound follows from the uncorrelated case. Other analysis tasks seem feasible but

challenging: for example, to output a good set of cluster centers for those points with $v_i > \tau$, weighted by $w_i f(t - t_i)$. It will be of interest to understand exactly which such correlated aggregations are possible in a streaming setting.

Acknowledgments. We thank Divesh Srivastava for helpful discussions, and Kewei Tu for useful pointers on §3.2.3.

References

[1] R. Ananthakrishna, A. Das, J. Gehrke, F. Korn, S. Muthukrishnan, and D. Srivastava. Efficient approximation of correlated sums on data streams. *IEEE Transactions on Knowledge and Data Engineering*, 15(3):569–572, 2003.

[2] C. Busch and S. Tirthapura. A deterministic algorithm for summarizing asynchronous streams over a sliding window. In *STACS*, 2007.

[3] D. Chatziantoniou and K. A. Ross. Querying multiple features of groups in relational databases. In *VLDB*, 1996.

[4] E. Cohen and M. Strauss. Maintaining time-decaying stream aggregates. In *PODS*, 2003.

[5] G. Cormode, F. Korn, and S. Tirthapura. Time-decaying aggregates in out-of-order streams. In *ICDE*, 2008.

[6] G. Cormode, S. Tirthapura, and B. Xu. Time-decaying sketches for sensor data aggregation. In *PODC*, 2007.

[7] M. Datar, A. Gionis, P. Indyk, and R. Motwani. Maintaining stream statistics over sliding windows. *SIAM Journal on Computing*, 31(6):1794–1813, 2002.

[8] J. Gehrke, F. Korn, and D. Srivastava. On computing correlated aggregates over continual data streams. In *SIGMOD*, 2001.

[9] P. Gibbons and S. Tirthapura. Distributed streams algorithms for sliding windows. In *SPAA*, 2002.

[10] E. Kushilevitz and N. Nisan. *Communication Complexity*. Cambridge University Press, 1997.

[11] N. Shrivastava, C. Buragohain, D. Agrawal, and S. Suri. Medians and beyond: new aggregation techniques for sensor networks. In *SenSys*, 2004.

[12] B. Xu, S. Tirthapura, and C. Busch. Sketching asynchronous data streams over sliding windows. *Distributed Computing*, 20(5):359–374, 2008.