

Give Me Some Slack: Efficient Network Measurements

Ran Ben Basat

Department of Computer Science, Technion
sran@cs.technion.ac.il

Gil Einziger

Nokia Bell Labs
gil.einziger@nokia.com

Roy Friedman

Department of Computer Science, Technion
roy@cs.technion.ac.il

Abstract

Many networking applications require timely access to recent network measurements, which can be captured using a sliding window model. Maintaining such measurements is a challenging task due to the fast line speed and scarcity of fast memory in routers. In this work, we study the impact of allowing *slack* in the window size on the asymptotic requirements of sliding window problems. That is, the algorithm can dynamically adjust the window size between W and $W(1+\tau)$ where τ is a small positive parameter. We demonstrate this model's attractiveness by showing that it enables efficient algorithms to problems such as MAXIMUM and GENERAL-SUMMING that require $\Omega(W)$ bits even for constant factor approximations in the exact sliding window model. Additionally, for problems that admit sub-linear approximation algorithms such as BASIC-SUMMING and COUNT-DISTINCT, the slack model enables a further asymptotic improvement.

The main focus of the paper is on the widely studied BASIC-SUMMING problem of computing the sum of the last W integers from $\{0, 1, \dots, R\}$ in a stream. While it is known that $\Omega(W \log R)$ bits are needed in the exact window model, we show that approximate windows allow an *exponential* space reduction for constant τ .

Specifically, for $\tau = \Theta(1)$, we present a space lower bound of $\Omega(\log(RW))$ bits. Additionally, we show an $\Omega(\log(W/\epsilon))$ lower bound for $RW\epsilon$ additive approximations and a $\Omega(\log(W/\epsilon) + \log \log R)$ bits lower bound for $(1 + \epsilon)$ multiplicative approximations. Our work is the first to study this problem in the exact and additive approximation settings. For all settings, we provide memory optimal algorithms that operate in *worst case* constant time. This strictly improves on the work of [14] for $(1 + \epsilon)$ -multiplicative approximation that requires $O(\epsilon^{-1} \log(RW) \log \log(RW))$ space and performs updates in $O(\log(RW))$ worst case time. Finally, we show asymptotic improvements for the COUNT-DISTINCT, GENERAL-SUMMING and MAXIMUM problems.

2012 ACM Subject Classification Theory of computation \rightarrow Streaming, sublinear and near linear time algorithms

Keywords and phrases Streaming, Network Measurements, Statistics, Lower Bounds

Digital Object Identifier 10.4230/LIPIcs.MFCS.2018.34

Related Version A full version of the paper is available at [5], <https://arxiv.org/abs/1703.01166>.



© Ran Ben-Basat, Gil Einziger, and Roy Friedman;
licensed under Creative Commons License CC-BY


43rd International Symposium on Mathematical Foundations of Computer Science (MFCS 2018).

Editors: Igor Potapov, Paul Spirakis, and James Worrell; Article No. 34; pp. 34:1–34:16

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



■ **Figure 1** We need to answer each query with respect to a τ -slack window that must include the last W items, but may or may not consider a suffix of the previous $W\tau$ elements.

1 Introduction

Network algorithms in diverse areas such as traffic engineering, load balancing and quality of service [2, 9, 21, 24, 31] rely on timely link measurements. In such applications recent data is often more relevant than older data, motivating the notions of *aging* and *sliding window* [6, 11, 15, 25, 27]. For example, a sudden decrease in the average packet size on a link may indicate a SYN attack [26]. Additionally, a load balancer may benefit from knowing the current utilization of a link to avoid congestion [2].

While conceptually simple, conveying the necessary information to network algorithms is a difficult challenge due to current memory technology limitations. Specifically, DRAM memory is abundant but too slow to cope with the line rate while SRAM memory is fast enough but has a limited capacity [10, 13, 29]. Online decisions are therefore realized through space efficient data structures [7, 8, 16, 17, 4, 23, 28, 30] that store measurement statistics in a concise manner. For example, [16, 28] utilize probabilistic counters that only require $O(\log \log N)$ bits to approximately represent numbers up to N . Others conserve space using variable sized counter encoding [17, 23] and monitoring only the frequent elements [6].

BASIC-SUMMING is one of the most basic textbook examples of such approximated sliding window stream processing problems [14]. In this problem, one is required to keep track of the sum of the last W elements, when all elements are non-negative integers in the range $\{0, 1, \dots, R\}$. The work in [14] provides a $(1+\epsilon)$ -multiplicative approximation of this problem using $O(\frac{1}{\epsilon} \cdot (\log^2 W + \log R \cdot (\log W + \log \log R)))$ bits. The amortized time complexity is $O(\frac{\log R}{\log W})$ and the worst case is $O(\log W + \log R)$. In contrast, we previously showed an $RW\epsilon$ -additive approximation with $\Theta(\frac{1}{\epsilon} + \log W\epsilon)$ bits [3].

Sliding window counters (approximated or accurate) require asymptotically more space than plain stream counters. Such window counters are prohibitively large for networking devices which already optimize the space consumption of plain counters.

This paper explores the concept of *slack*, or approximated sliding window, bridging this gap. Figure 1 illustrates a “window” in this model. Here, each query may select a τ -slack window whose size is between W (the green elements) and $W(1+\tau)$ (the green plus yellow elements). The goal is to compute the sum with respect to this chosen window.

Slack windows were also considered in previous works [14, 27] and we call the problem of maintaining the sum over a slack window SLACK SUMMING. Datar et al. [14] showed that constant slack reduces the required memory from $O(\frac{1}{\epsilon} \cdot (\log^2 W + \log R \cdot (\log W + \log \log R)))$ to $O(\epsilon^{-1} \log(RW) \log \log(RW))$. For τ -slack windows they provide a $(1+\epsilon)$ -multiplicative approximation using $O(\epsilon^{-1} \log(RW)(\log \log(RW) + \log \tau^{-1}))$ bits.

■ **Table 1** Comparison of BASIC-SUMMING algorithms. Our contributions are in bold. All algorithms process elements in constant time except for the rightmost column where both update in $O(\log(RW))$ time. We present matching lower bounds to all our algorithms.

	Exact Sum	Additive Error	Multiplicative Error	
$\tau = \Theta(1)$	$\Theta(\log(RW))$	$\Theta(\log(W/\epsilon))$	$\Theta(\log(W/\epsilon) + \log \log R)$	$O(\epsilon^{-1} \log(RW) \log \log(RW))$ [14]
Exact Window	$\Theta(W \log R)$	$\Theta(\epsilon^{-1} + \log W)$ [3]	$O(\epsilon^{-1} \log^2(RW))$ [22]	$O(\epsilon^{-1} \log RW \log(W \log R))$ [14]

Our Contributions

This paper studies the space and time complexity reductions that can be attained by allowing *slack* – an error in the window size. Our results demonstrate exponentially smaller and asymptotically faster data structures compared to various problems over exact windows. We start with deriving lower bounds for three variants of the BASIC-SUMMING problem – when computing an exact sum over a slack window, or when combined with an additive and a multiplicative error in the sum. We present algorithms that are based on dividing the stream into $W\tau$ -sized blocks. Our algorithms sum the elements within each block and represent each block’s sum in a cyclic array of size τ^{-1} . We use multiple compression techniques during different stages to drive down the space complexity. The resulting algorithms are space optimal, substantially simpler than previous work, and reduce update time to $O(1)$.

For exact SLACK SUMMING, we present a lower bound of $\Omega(\tau^{-1} \log(RW\tau))$ bits. For $(1+\epsilon)$ multiplicative approximations we prove an $\Omega(\log(W/\epsilon) + \tau^{-1} (\log(\tau/\epsilon) + \log \log(RW)))$ bits bound when $\tau = \Omega\left(\frac{1}{\log RW}\right)$. We show that $\Omega(\tau^{-1} \log[1 + \tau/\epsilon] + \log(W/\epsilon))$ bits are required for $RW\epsilon$ additive approximations.

Next, we introduce algorithms for the SLACK SUMMING problem, which asymptotically reduce the required memory compared to the sliding window model. For the exact and additive error versions of the problem, we provide memory optimal algorithms. In the multiplicative error setting, we provide an $O(\tau^{-1} (\log \epsilon^{-1} + \log \log(RW\tau)) + \log(RW))$ space algorithm. This is asymptotically optimal when $\tau = \Omega(\log^{-1} W)$ and $R = \text{poly}(W)$. It also asymptotically improves [14] when $\tau^{-1} = o(\epsilon^{-1} \log(RW))$. We further provide an asymptotically optimal solution for constant τ , even when $R = W^{\omega(1)}$. All our algorithms are deterministic and operate in worst case constant time. In contrast, the algorithm of [14] works in $O(\log RW)$ worst case time.

To exemplify our results, consider monitoring the average bandwidth (in bytes per second) passed through a router in a 24 hours window, i.e., $W \triangleq 86400$ seconds. Assuming we use a 100GbE fiber transceiver, our stream values are bounded by $R \approx 2^{34}$ bytes. If we are willing to withstand an error of $\epsilon = 2^{-20}$ (i.e., about 16KBps), the work of [3] provides an additive approximation over the sliding window and requires about 120KB. In contrast, using a 10 minutes slack ($\tau \triangleq \frac{1}{144}$), our algorithm for **exact** SLACK SUMMING requires only 800 bytes, 99% less than approximate summing over exact sliding window. For the same slack size, the algorithm of [14] requires more space than our **exact** algorithm even for a large 3% error. Further, if we also allow the same additive error ($\epsilon = 2^{-20}$), we provide an algorithm that requires only 240 bytes - a reduction of more than 99.8% !

Table 1 compares our results for the important case of constant slack with [14]. As depicted, our *exact* algorithm is faster and more space efficient than the multiplicative approximation of [14]. Comparing our multiplicative approximation algorithm to that of [14], we present *exponential* space reductions in the dependencies on ϵ^{-1} and R , with an asymptotic reduction in W as well. We also improve the update time from $O(\log(RW))$ to $O(1)$.

Finally, we apply the slack window approach to multiple streaming problems, including MAXIMUM, GENERAL-SUMMING, COUNT-DISTINCT and STANDARD-DEVIATION. We show that, while some of these problems cannot be approximated on an exact window in sub-linear space (e.g. maximum and general sum), we can easily do so for slack windows. In the count distinct problem, a constant slack yields an asymptotic space reduction over [11, 19].

2 Preliminaries

For $\ell \in \mathbb{N}$, we denote $[\ell] \triangleq \{0, 1, \dots, \ell\}$. We consider a stream of data elements x_1, x_2, \dots, x_t , where at each step a new element $x_i \in [R]$ is added to S . A W -sized window contains only the last W elements: $x_{t-W+1} \dots x_t$. We say that \mathcal{F} is a τ -slack W -sized window if there exists $c \in [W\tau - 1]$ such that $\mathcal{F} = x_{t-(W+c)+1} \dots x_t$. For simplicity, we assume that τ^{-1} and $W\tau$ are integers. Unless explicitly specified, the base of all logs is 2.

Algorithms for the SLACK SUMMING problem are required to support two operations:

1. UPDATE(x_t) Process a new element $x_t \in [R]$.
2. OUTPUT () Return a pair $\langle \hat{S}, c \rangle$ such that $c \in \mathbb{N}$ is the slack size and \hat{S} is an estimation of the last $W + c$ elements sum, i.e., $S \triangleq \sum_{k=t-(W+c)+1}^t x_k$.

We consider three types of algorithms for SLACK SUMMING:

1. **Exact algorithms:** an algorithm \mathbb{A} solves (W, τ) -EXACT SUMMING if its OUTPUT returns $\langle \hat{S}, c \rangle$ that satisfies $0 \leq c < W\tau$ and $\hat{S} = S$.
2. **Additive algorithms:** we say that \mathbb{A} solves (W, τ, ϵ) -ADDITIVE SUMMING if its OUTPUT function returns $\langle \hat{S}, c \rangle$ that satisfies $0 \leq c < W\tau$ and $|S - \hat{S}| < RW\epsilon$.
3. **Multiplicative algorithms:** \mathbb{A} solves (W, τ, ϵ) -MULTIPLICATIVE SUMMING if its OUTPUT returns $\langle \hat{S}, c \rangle$ satisfying $0 \leq c < W\tau$ and $\frac{S}{1+\epsilon} < \hat{S} \leq S$ if $S > 0$, and $\hat{S} = 0$ otherwise.

3 Lower Bounds

In this section, we analyze the space required for solving the SLACK SUMMING problems. Intuitively, our bounds are derived by constructing a set of inputs that any algorithm must distinguish to meet the required guarantees. There are two tricks that we frequently use in these lower bounds. The first is setting the input such that the slack consists only of zeros, and thus the algorithm must return the desired approximation of the remaining window. The next is using a “cycle argument” – consider two inputs x and $x \cdot y$ for $x, y \in \{0, 1, \dots, R\}^*$. If both lead to the same memory configuration, so do such xy^k for any $k \in \mathbb{N}$. Thus, if there is a k such that no single answer approximates x and xy^k well, then x and xy had to lead to separate memory configurations in the first place.

3.1 (W, τ) -Exact Summing

We start by proving lower bounds on the memory required for exact SLACK SUMMING.

► **Lemma 1.** *Any deterministic algorithm \mathbb{A} that solves the (W, τ) -EXACT SUMMING problem must use at least $\lceil \log(RW(W+1)/2+1) \rceil \geq \lfloor \log(RW^2) \rfloor$ bits.*

Proof. Consider the following language

$$L_{E_1} \triangleq \{0^{W\tau+i}\sigma R^{W-i-1}0^j \mid i, j \in [W-1], i \geq j, \sigma \in ([R] \setminus \{0\})\} \cup \{0^{W+W\tau}\}.$$

That is, L_{E_1} contains a word with $W + W\tau$ consecutive zeros and the rest of the words in L_{E_1} are composed of these components in this order:

- $W\tau + i$ zeros for some $i \in [W - 1]$.
- a non zero symbol σ .
- $W - i - 1$ repetitions of the maximal symbol (R).
- j zeros for some $j \in [i]$.

Our lower bound stems from the observation that every word in L_{E_1} must lead to a different state. The language size is: $|L_{E_1}| = 1 + \sum_{i=0}^{W-1} R(i+1) = 1 + RW(W+1)/2$. Therefore, the number of required bits is at least: $\lceil \log |L_{E_1}| \rceil > (\log(RW^2) - 1)$. Further, this number is an integer and therefore at least $\lfloor \log(RW^2) \rfloor$ bits are required.

First, notice that the word composed of $W + W\tau$ zeros requires a unique configuration as \mathbb{A} must return 0 after processing that word. In contrast, it must not return 0 after processing any other word as there is at least a single R within the last W elements.

Let $w_1, w_2 \in L_{E_1}$ be two different words that are not all-zeros. We need to show that w_1 and w_2 require different memory configuration.

By definition of L_{E_1} , $w_1 = 0^{W\tau+i_1}\sigma_1 R^{W-i_1-1}0^{j_1}$ and $w_2 = 0^{W\tau+i_2}\sigma_2 R^{W-i_2-1}0^{j_2}$. Observe that the last W elements of w_1, w_2 are $0^{i_1-j_1}\sigma_1 R^{W-i_1-1}0^{j_1}$ and $0^{i_2-j_2}\sigma_2 R^{W-i_2-1}0^{j_2}$ respectively and that both are preceded with at least $W\tau$ zeros. If $i_1 \neq i_2$ or $\sigma_1 \neq \sigma_2$, then $\sigma_1 + R \cdot (W - i_1 - 1) \neq \sigma_2 + R \cdot (W - i_2 - 1)$ and thus \mathbb{A} cannot return the same count for both, regardless of the slack, as it is all zeros in both w_1 and w_2 .

Next, assume that $i_1 = i_2$, $\sigma_1 = \sigma_2$ and that without loss of generality $j_1 < j_2$. This means that both w_1 and w_2 have the same count.

Since $j_1 < j_2$, w_1 is a strict prefix of w_2 , i.e., $w_2 = w_1 \cdot 0^{j_2-j_1}$. Assume by contradiction that after processing w_1, w_2 \mathbb{A} reaches the same memory configuration. Since \mathbb{A} is deterministic, this means that it must reach the same configuration after seeing $w_1 \cdot 0^{z(j_2-j_1)}$ for any integer z . By choosing $z = W(1 + \tau)$, we get that the algorithm reaches this configuration once again while the entire window consists of zeros. This is a contradiction since $\sigma_1, \sigma_2 \neq 0$, and the algorithm cannot answer both w_1 and $w_1 \cdot 0^{z(j_2-j_1)}$ correctly. ◀

We now use Lemma 1 to show the following lower bound on (W, τ) -EXACT SUMMING algorithms:


► **Theorem 2.** *Any deterministic algorithm \mathbb{A} that solves the (W, τ) -EXACT SUMMING problem must use at least $\max \{ \lfloor \log(RW^2) \rfloor, \lceil \lceil \tau^{-1}/2 \rceil \log(RW\tau + 1) \rceil \}$ bits.*

Proof. Lemma 1 shows a $\lfloor \log(RW^2) \rfloor$ bound. We proceed with showing a lower bound $\lceil \lceil \tau^{-1}/2 \rceil \log(RW\tau + 1) \rceil$ bits. Consider the following languages:

$$L_{E_2} \triangleq \{0^{W\tau+i}\sigma R^{W\tau-i-1} \mid i \in [W\tau - 1], \sigma \in [R]\},$$

$$\overline{L_{E_2}} \triangleq \{w_1 \cdot w_2 \cdots w_{\lceil \tau^{-1}/2 \rceil} \mid \forall i : w_i \in L_{E_2}\}.$$

Notice that $|\overline{L_{E_2}}| = (RW\tau + 1)^{\lceil \tau^{-1}/2 \rceil}$ since each of the words in L_{E_2} has a distinct sum of literals, and each number in $\{0, 1, \dots, RW\tau\}$ is the sum of a word. We show that each input in $\overline{L_{E_2}}$ must be mapped into a distinct memory configuration. Let $S_1 \triangleq w_{1,1} \cdot w_{2,1} \cdots w_{\lceil \tau^{-1}/2 \rceil, 1}$, $S_2 \triangleq w_{1,2} \cdot w_{2,2} \cdots w_{\lceil \tau^{-1}/2 \rceil, 2}$ be two distinct inputs in $\overline{L_{E_2}}$ such that $\forall i : w_{i,1}, w_{i,2} \in L_{E_2}$. Denote $\chi \triangleq \max \{i \in [\lceil \tau^{-1}/2 \rceil] \mid w_{i,1} \neq w_{i,2}\}$ – the last place in which S_1 differs from S_2 ; also, denote $w_{\chi,1} \triangleq 0^{W\tau}a$, $w_{\chi,2} \triangleq 0^{W\tau}b$. Consider the sequences $S_1^* = S_1 \cdot 0^{2W\tau(\chi-1/2)}$ and $S_2^* = S_2 \cdot 0^{2W\tau(\chi-1/2)}$. Notice that the last W elements windows for S_1^*, S_2^* are $a \cdot w_{\chi+1,1} \cdots w_{\lceil \tau^{-1}/2 \rceil, 1} \cdot 0^{2W\tau(\chi-1/2)}$ and $b \cdot w_{\chi+1,2} \cdots w_{\lceil \tau^{-1}/2 \rceil, 2} \cdot 0^{2W\tau(\chi-1/2)}$ respectively, and that the preceding $W\tau$ elements of both are all zeros. An illustration of the setting appears in Figure 2. By our choice of χ , we have that the sum of the last W elements of S_1^*



■ **Figure 2** An illustration of the $\lceil \tau^{-1}/2 \rceil \log(RW\tau + 1)$ lower bound setting. If we assume that after seeing $w_{1,1} \cdot w_{2,1} \cdots w_{\lceil \tau^{-1}/2 \rceil, 1}$ we reach the same configuration as after processing $w_{1,2} \cdot w_{2,2} \cdots w_{\lceil \tau^{-1}/2 \rceil, 2}$, then we provide a wrong answer for at least one of S_1^*, S_2^* .

and S_2^* is different, and since the slack is all zeros, no answer is correct on both. Finally, note that this implies that S_1, S_2 had to reach different configurations, as otherwise \mathbb{A} would reach the same configuration after processing the additional $2W\tau(\chi - 1/2)$ zeros. ◀

3.2 (W, τ, ϵ) -Additive Summing

Next, Theorem 3 shows a lower bound for additive approximations of SLACK SUMMING.

► **Theorem 3.** For $\epsilon < 1/4$, any deterministic algorithm \mathbb{A} that solves the (W, τ, ϵ) -ADDITIVE SUMMING problem requires $\max\{\log(W/\epsilon) - O(1), \lceil \lceil \tau^{-1}/2 \rceil \log \lceil \tau/2\epsilon + 1 \rceil\}$ bits.

Before we prove Theorem 3, we start with a simpler lower bound.

► **Lemma 4.** Let $\epsilon < 1/4$. Any deterministic algorithm that solves the (W, τ, ϵ) -ADDITIVE SUMMING problem must use at least $\log(W/\epsilon) - O(1)$ bits.

Proof. Denote by $\text{rep}(x) \triangleq (x \bmod R) \cdot R^{\lfloor x/R \rfloor}$ a sequence in $\{\sigma R^* \mid \sigma \in [R]\}$ whose sum is x . Next, consider the following languages:

$$L_{A_1} \triangleq \{\text{rep}(k \cdot 2RW\epsilon) \mid k \in [\lceil 1/4\epsilon \rceil] \setminus \{0\}\}; \quad \overline{L_{A_1}} \triangleq 0^{W+W\tau} \cdot L_{A_1} \cdot \{0^q \mid q \in [\lceil W/2 \rceil]\}.$$

First, notice that $|L_{A_1}| = \lceil 1/4\epsilon \rceil$ and that all words in L_{A_1} have length of at most $W/2$. This means that $|\overline{L_{A_1}}| = \lceil 1/4\epsilon \rceil \lceil W/2 + 1 \rceil > \lceil W/8\epsilon \rceil$.

We now show that every word in $\overline{L_{A_1}}$ must have a dedicated memory configuration, thereby implying a $\lceil \log \lceil W/8\epsilon \rceil \rceil$ bits bound. Let $w_1 = 0^{W+W\tau} \cdot x_1 \cdot 0^{q_1}$ and $w_2 = 0^{W+W\tau} \cdot x_2 \cdot 0^{q_2}$ be two distinct words in $\overline{L_{A_1}}$ such that $x_1, x_2 \in L_{A_1}$ and $q_1, q_2 \in \lceil W/2 \rceil$. If $x_1 \neq x_2$, then their most recent W elements differ by more than $2RW\epsilon$ and there is no output that is correct for both. Note that the slack of both w_1 and w_2 is all zeros. Hence, w_1 and w_2 require different memory configurations.

Assume that $x_1 = x_2$ and that by contradiction both w_1 and w_2 reached the same memory configuration. Since $w_1 \neq w_2$ and $x_1 = x_2$, then $q_1 \neq q_2$ and without loss of generality $q_1 < q_2$. This implies that w_1 is a prefix of w_2 so that $w_2 = w_1 \cdot 0^{q_2 - q_1}$. Thus, \mathbb{A} enters the shared configuration after reading w_1 and revisits it after reading $0^{q_2 - q_1}$. \mathbb{A} is a deterministic algorithm and therefore it reaches the same configuration also for the following word: $w_1 \cdot 0^{(W+W\tau)(q_2 - q_1)}$. In that word, the last $W + W\tau$ elements are all zeros while the sum of the last W elements in w_1 is at least $2RW\epsilon$. Hence, there is no return value that is correct for both w_1 and $w_1 \cdot 0^{(W+W\tau)(q_2 - q_1)}$. ◀

We are now ready to prove Theorem 3. The theorem says that for $\epsilon < 1/4$, any deterministic algorithm \mathbb{A} that solves the (W, τ, ϵ) -ADDITIVE SUMMING problem requires $\max\{\log(W/\epsilon) - O(1), \lceil \lceil \tau^{-1}/2 \rceil \log \lceil \tau/2\epsilon + 1 \rceil\}$ bits.

Proof. Lemma 4 shows that \mathbb{A} must use at least $\log(W/\epsilon) - O(1)$ bits. Given $x \in [RW\tau]$, we denote by $\text{rep}(x) \triangleq 0^{2W\tau - \lfloor x/R \rfloor - 1} \cdot (x \bmod R) \cdot R^{\lfloor x/R \rfloor}$ a sequence of the following form: $\{0^{W\tau+i} \sigma R^{W\tau-i-1} \mid i \in [W\tau - 1], \sigma \in [R]\}$ whose sum is x . We consider the following languages:

$$L_{A_2} \triangleq \{\text{rep}(k \cdot 2RW\epsilon) \mid k \in [\lceil \tau/2\epsilon \rceil]\}; \quad \overline{L_{A_2}} \triangleq \{w_1 \cdot w_2 \cdots w_{\lceil \tau-1/2 \rceil} \mid \forall i : w_i \in L_{A_2}\}.$$

Our goal is to show that no two words in $\overline{L_{A_2}}$ have the same memory configuration. Let $S_1, S_2 \in \overline{L_{A_2}}$ so that $S_1 \neq S_2$. Denote $S_1 \triangleq w_{1,1} \cdot w_{2,1} \cdots w_{\lceil \tau-1/2 \rceil, 1}$ and $S_2 \triangleq w_{1,2} \cdot w_{2,2} \cdots w_{\lceil \tau-1/2 \rceil, 2}$, while $\forall i : w_{i,1}, w_{i,2} \in L_{A_2}$. We denote $\chi \triangleq \max\{i \mid w_{i,1} \neq w_{i,2}\}$ – the last place in which S_1 differs from S_2 .

Next, consider the following sequences: $S_1^* = S_1 \cdot 0^{2W\tau(\chi-1/2)}$ and $S_2^* = S_2 \cdot 0^{2W\tau(\chi-1/2)}$. The last $W + W\tau$ elements in S_1^* are $w_{\chi,1} \cdots w_{\lceil \tau-1/2 \rceil, 1} \cdot 0^{2W(\chi-1/2)\tau}$ and in S_2^* $w_{\chi,2} \cdots w_{\lceil \tau-1/2 \rceil, 2} \cdot 0^{2W(\chi-1/2)\tau}$. Additionally, the $W\tau$ elements slack in both S_1^* and S_2^* are all zeros. Now, since the sum of $w_{\chi,1}$ and $w_{\chi,2}$ must differ by at least $2RW\epsilon$, no number can approximate both with less than $RW\epsilon$ error. \blacktriangleleft

3.3 (W, τ, ϵ) -Multiplicative Summing

In this section, we show lower bounds for multiplicative approximations of SLACK SUMMING. We start with Lemma 5, whose proof appears in the full version of this paper [5].

► **Lemma 5.** *For $\epsilon < 1/4$, any deterministic algorithm \mathbb{A} for the (W, τ, ϵ) -MULTIPLICATIVE SUMMING problem requires at least $\log(W/\epsilon) + \log \log(RW\epsilon) - O(1)$ memory bits.*

To extend our multiplicative lower bound, we use the following fact:

► **Fact 1.** *For any $x \neq 1, y \in \mathbb{R}$, the sequence $\{c_i\}_{i=1}^\infty$, defined as $c_n \triangleq \begin{cases} 1 & n = 1 \\ x \cdot c_{n-1} + y & \text{otherwise} \end{cases}$ can be represented using a closed form as $c_n = x^{n-1} + y \cdot \frac{x^n - 1}{x - 1}$.*

Next, let $k \in \mathbb{N}$ and $\psi, \epsilon \in \mathbb{R}$, such that $\psi \geq 2, \epsilon > 0, k \geq 1$; consider the integer sequence

$$a_{n,k} \triangleq \begin{cases} 1 & n = 1 \\ \left\lceil (1 + \epsilon) \left(a_{n-1,k} + \sum_{i=1}^{k-1} \psi^i \right) \right\rceil & \text{otherwise.} \end{cases}$$

Using the fact above, we show the following lemma:

► **Lemma 6.** *For every integer $n \geq 1$ we have $a_{n,k} \leq 4\epsilon^{-1}(1 + \epsilon)^{n+1}\psi^{k-1}$.*

Proof. To apply Fact 1, we define an upper bounding sequence $\{b_{i,k}\}_{i=1}^\infty$ as follows:

$$b_{n,k} \triangleq \begin{cases} 1 & n = 1 \\ (1 + \epsilon) \left(b_{n-1,k} + \sum_{i=1}^{k-1} \psi^i \right) + 1 & \text{otherwise.} \end{cases}$$

Thus, we can rewrite the n 'th element of the sequence as:

$$b_{n,k} = (1 + \epsilon)^{n-1} + \frac{(1+\epsilon)^n - 1}{(1+\epsilon) - 1} \left((1 + \epsilon) \sum_{i=1}^{k-1} \psi^i + 1 \right).$$

We can now use this representation to derive an upper bound of $b_{n,k}$:

$$\begin{aligned} b_{n,k} &= (1 + \epsilon)^{n-1} + \left((1 + \epsilon) \sum_{i=1}^{k-1} \psi^i + 1 \right) \frac{(1+\epsilon)^n - 1}{(1+\epsilon) - 1} \\ &\leq (1 + \epsilon)^{n-1} + \left((1 + \epsilon) 2\psi^{k-1} \right) \frac{(1 + \epsilon)^n - 1}{\epsilon} \leq 4\epsilon^{-1}(1 + \epsilon)^{n+1}\psi^{k-1}. \end{aligned}$$

Finally, since $a_{n,k} \leq b_{n,k}$ for any n, k , we conclude that $a_{n,k} \leq 4\epsilon^{-1}(1 + \epsilon)^{n+1}\psi^{k-1}$. \blacktriangleleft

We now define the integer set I_k as $I_k \triangleq \{a_{n,k} \mid a_{n,k} \leq \psi^k\}$, and proceed to bound $|I_k|$.

► **Lemma 7.** *For any $k \geq 1$ we have $|I_k| \geq \epsilon^{-1} \ln(\psi\epsilon/4) - 1$.*

Proof. Clearly, the cardinality of I_k is the largest n for which $a_{n,k} \leq \psi^k$. According to Lemma 6, we have that $a_{n,k} \leq 4\epsilon^{-1}(1+\epsilon)^{n+1}\psi^{k-1}$, and thus:

$$\begin{aligned} |I_k| &= \arg \max \{n \mid 4\epsilon^{-1}(1+\epsilon)^{n+1}\psi^{k-1} \leq \psi^k\} \\ &\geq \log_{1+\epsilon}(\psi\epsilon/4) - 1 = \frac{\ln(\psi\epsilon/4)}{\ln(1+\epsilon)} - 1 \geq \epsilon^{-1} \ln(\psi\epsilon/4) - 1. \quad \blacktriangleleft \end{aligned}$$

We proceed with a stronger lower bound for non-constant τ values.

► **Lemma 8.** *For $\frac{1}{2\log(RW)-8} \leq \tau \leq 1$, any deterministic algorithm \mathbb{A} that solves (W, τ, ϵ) -MULTIPLICATIVE SUMMING requires at least $\Omega(\tau^{-1}(\log(\tau/\epsilon) + \log \log(RW)))$ bits.*

Proof. We use $\text{rep}(x) \triangleq (x \bmod R) \cdot R^{\lfloor x/R \rfloor}$ to denote a sequence in $\{\sigma R^* \mid \sigma \in [R]\}$ that has a sum of x . For an integer set I_k , we denote $\text{rep}(I_k) \triangleq \{\text{rep}(x) \mid x \in I_k\}$. We now choose the value of ψ to be $\psi \triangleq \lceil \tau^{-1/2} \sqrt{RW/8} \rceil$; notice that $\psi \geq 2$ as required. Next, consider:

$$\begin{aligned} \overline{L_{M,2}} &\triangleq 0^W \cdot 0^{W\tau} \cdot \text{rep}(I_{\lceil \tau^{-1/2} \rceil}) \cdot 0^{W\tau} \cdot \text{rep}(I_{\lceil \tau^{-1/2} \rceil - 1}) \cdots 0^{W\tau} \cdot \text{rep}(I_1) \\ &= \{0^W \cdot w_1 \cdot w_2 \cdots w_{\lceil \tau^{-1/2} \rceil} \mid \forall i : w_i \in \{0^{W\tau} \cdot \text{rep}(x) \mid x \in I_{\lceil \tau^{-1/2} \rceil + 1 - i}\}\}. \end{aligned}$$

That is, every word in the $\overline{L_{M,2}}$ language consists of a concatenation of words $w_1, \dots, w_{\lceil \tau^{-1/2} \rceil}$, such that every w_i starts with $W\tau$ zeros followed by a string representing an integer in $I_{\lceil \tau^{-1/2} \rceil + 1 - i}$, which is defined above. According to Lemma 7 we have that

$$\begin{aligned} \log(|\overline{L_{M,2}}|) &\geq \log\left(\left(\epsilon^{-1} \ln(\psi\epsilon/4) - 1\right)^{\lceil \tau^{-1/2} \rceil}\right) \\ &= \lceil \tau^{-1/2} \rceil (\log \epsilon^{-1} + \log \log(\psi\epsilon) - O(1)) \\ &= \Omega\left(\tau^{-1} \left(\log \epsilon^{-1} + \log \log\left(\lceil \tau^{-1/2} \rceil \sqrt{RW/8} \cdot \epsilon\right)\right)\right) \\ &= \Omega\left(\tau^{-1} \left(\log \epsilon^{-1} + \log\left(\frac{\log(RW/8)}{\lceil \tau^{-1/2} \rceil} + \log \epsilon\right)\right)\right) \\ &= \Omega(\tau^{-1}(\log(\tau/\epsilon) + \log \log(RW))). \end{aligned}$$

Next, we show that every two words in $\overline{L_{M,2}}$ must reach different memory configurations, thereby implying a $\Omega(\log(|\overline{L_{M,2}}|))$ bits lower bound. Let $S_1 \neq S_2 \in \overline{L_{M,2}}$ such that $S_1 = 0^W \cdot w_{1,1} \cdots w_{\lceil \tau^{-1/2} \rceil, 1}$, $S_2 = 0^W \cdot w_{1,2} \cdots w_{\lceil \tau^{-1/2} \rceil, 2}$, and $\forall i \in \{1, \dots, \lceil \tau^{-1/2} \rceil\} j \in \{1, 2\} : w_{i,j} \in \{0^{W\tau} \cdot \text{rep}(x) \mid x \in I_{\lceil \tau^{-1/2} \rceil + 1 - i}\}$. We next assume by contradiction that S_1 and S_2 leads \mathbb{A} to the same memory configuration. Let $\chi \in \{1, \dots, \lceil \tau^{-1/2} \rceil\}$ such that $w_{\chi,1} \neq w_{\chi,2}$. Since \mathbb{A} reaches an identical configuration after reading S_1, S_2 , and as it is deterministic, \mathbb{A} must reach the same configuration when processing $S_1 \cdot 0^{2W\tau(\chi-1/2)}$ and $S_2 \cdot 0^{2W\tau(\chi-1/2)}$. Next, observe that for every $k \in \{1, \dots, \lceil \tau^{-1/2} \rceil\}$, the representation length of any of its words is bounded by $\lceil \psi^k/R \rceil$. Thus, the length of a word in

$$\{w_1 \cdot w_2 \cdots w_{\lceil \tau^{-1/2} \rceil} \mid \forall i : w_i \in \{0^{W\tau} \cdot \text{rep}(x) \mid x \in I_{\lceil \tau^{-1/2} \rceil + 1 - i}\}\} \text{ is at most}$$

$$\begin{aligned} \sum_{k=1}^{\lceil \tau^{-1}/2 \rceil} \lceil W\tau + \psi^k/R \rceil &\leq \lceil \tau^{-1}/2 \rceil (W\tau + 1) + 2\psi^{\lceil \tau^{-1}/2 \rceil}/R \\ &= \lceil \tau^{-1}/2 \rceil (W\tau + 1) + 2W/8 \leq 3W/4 + \lceil \tau^{-1}/2 \rceil + W\tau \leq W + W\tau. \end{aligned}$$

Now, since every word $w_{i,j}$ starts with a sequence of $W\tau$ zeros, the slack size chosen by the algorithm is irrelevant and the sums the algorithm must estimate are $\sum_{i=\chi}^{\lceil \tau^{-1}/2 \rceil} s(w_{i,1})$ and $\sum_{i=\chi}^{\lceil \tau^{-1}/2 \rceil} s(w_{i,2})$, where $s(w_{i,j})$ is simply the sum of the symbols in $w_{i,j}$. Note that $s(w_{\chi,1})$ and $s(w_{\chi,2})$ are integers in $I_{\lceil \tau^{-1}/2 \rceil + 1 - \chi}$. We assume without loss of generality that $s(w_{\chi,1}) < s(w_{\chi,2})$ (i.e., $s(w_{\chi,1}) < s(w_{\chi,2}) \in I_{\lceil \tau^{-1}/2 \rceil + 1 - \chi}$). Finally, it follows that

$$\sum_{i=\chi}^{\lceil \tau^{-1}/2 \rceil} s(w_{i,1}) \leq s(w_{\chi,1}) + \sum_{i=\chi+1}^{\lceil \tau^{-1}/2 \rceil} \max(I_{\lceil \tau^{-1}/2 \rceil + 1 - i}) \leq s(w_{\chi,1}) + \sum_{k=1}^{\chi-1} \psi^k \leq \frac{s(w_{\chi,2})}{1 + \epsilon},$$

where the last inequality follows from the definition of $I_{\lceil \tau^{-1}/2 \rceil + 1 - \chi}$. Thus, no \widehat{S} value is correct for both $S_1 \cdot 0^{2W\tau(\chi-1/2)}$ and $S_2 \cdot 0^{2W\tau(\chi-1/2)}$. ◀

Finally, we combine Lemma 5 and Lemma 8 to obtain the following lower bound:

► **Theorem 9.** For $\epsilon < 1/4$, $\frac{1}{2 \log(RW) - 8} \leq \tau \leq 1$, any deterministic algorithm for the (W, τ, ϵ) -MULTIPLICATIVE SUMMING problem requires at least

$$\Omega(\log(W/\epsilon) + \tau^{-1}(\log(\tau/\epsilon) + \log \log(RW))) \text{ bits.}$$

4 Upper Bounds

In this section, we introduce solutions for the SLACK SUMMING problems. In general, all our algorithms have a structure that consists of a subset of the following, where “rounding” has a different meaning for the exact, additive and multiplicative variants:

- Round the arriving item.
- Add the item into a counter y and round the counter.
- If a $W\tau$ -sized *block* ends, store it as a compressed representation of y . Sometimes we propagate the compression error to the following block; otherwise, we zero y .
- Use the block values and y to construct an estimation for the sum.

A key idea in our additive and multiplicative algorithms is to introduce rounding errors but maintain the accountability trail so that they do not snowball and exceed the desired guarantees. In the additive algorithm, our *double rounding* technique asymptotically improves over running $1/\tau$ separate plain stream (insertion only) algorithm instances.

4.1 (W, τ) -Exact Summing

We divide the stream into $W\tau$ -sized blocks and sum the number of arriving elements in each block with a $\lceil \log(RW\tau + 1) \rceil$ bits counter. We maintain the sum of the current block in a variable called y , c maintains the number of elements within the current block, and i is the current block number. The variable b is a cyclic buffer of τ^{-1} blocks. Every $W\tau$ steps, we assign the value of y to the oldest block (b_i) and increment i . Intuitively, we “forget” b_i when its block is no longer part of the window. To satisfy queries in constant time, we also maintain the sum of all active counters in a $\lceil \log(RW(1 + \tau) + 1) \rceil$ -bits variable named B . Algorithm 1 provides pseudocode for the described algorithm. We now analyze the memory consumption of Algorithm 1.

Algorithm 1 (W, τ) -EXACT SUMMING Algorithm.

Initialization: $y = 0, b = \bar{0}, B = 0, i = 0, c = 0$.

- 1: **function** UPDATE(x)
- 2: $y \leftarrow y + x$
- 3: $c \leftarrow (c + 1) \bmod W\tau$
- 4: **if** $c = 0$ **then** ▷ End of block
- 5: $B \leftarrow B - b_i + y$
- 6: $b_i \leftarrow y$
- 7: $y \leftarrow 0$
- 8: $i \leftarrow (i + 1) \bmod \tau^{-1}$
- 9: **function** OUTPUT
- 10: **return** $\langle B + y, c \rangle$

► **Theorem 10.** *Algorithm 1 uses $(\tau^{-1} + 1) \lceil \log(RW\tau + 1) \rceil + \log(RW^2) + O(1)$ bits.*

Proof. y takes $\lceil \log(RW\tau + 1) \rceil$ bits; B requires $\lceil \log(RW + 1) \rceil$; i adds $\lceil \log \tau^{-1} \rceil$ bits, while c needs $\lceil \log W\tau \rceil$ bits. Finally, b is a τ^{-1} -sized array of counters, each allocated with $\lceil \log(RW\tau + 1) \rceil$ bits. Overall, it uses $(\tau^{-1} + 1) \lceil \log(RW\tau + 1) \rceil + \log(RW^2) + 4$ bits. ◀

We conclude that Algorithm 1 is asymptotically optimal.

► **Theorem 11.** *Let $\mathcal{B} \triangleq \max \{ \lfloor \log(RW^2) \rfloor, \lceil \lceil \tau^{-1}/2 \rceil \log(RW\tau + 1) \rceil \}$ be the (W, τ) -EXACT SUMMING lower bound of Theorem 2. Algorithm 1 uses at most $\mathcal{B}(4 + o(1))$ memory bits.*


Theorem 11 shows that Algorithm 1 is only x4 larger than the lower bound. In the full version of this paper [5], we show that in some cases we can get considerably closer to the lower bound.

Finally, in the full version of this paper [5] we show that Algorithm 1 is correct.

4.2 (W, τ, ϵ) -Additive Summing

We now show that additional memory savings can be obtained by combining slackness with an additive error. First, we consider the case where $\tau \leq 2\epsilon$. In [3], we proposed an algorithm that sums over (exact) W elements window using the optimal $\Theta(\epsilon^{-1} + \log W)$ bits, with an additive error of $RW\epsilon$. Next, notice that if an algorithm solves (W, τ, ϵ) -ADDITIVE SUMMING, it also solves $(W, \tau, \tau/2)$ -ADDITIVE SUMMING; hence, we can apply Theorem 3 to conclude that it requires $\Omega(\tau^{-1} + \log W) = \Omega(\epsilon^{-1} + \log W)$. Thus, we can run the algorithm from [3] and remain asymptotically memory optimal with no slack at all!

Henceforth, we assume that $\tau > 2\epsilon$; we present an algorithm for the problem using a *2-stage rounding* technique. When a new item arrives, we scale it by R and then round the results to $O(\log \epsilon^{-1})$ bits. As in Section 4.1, we break the stream into non-overlapping blocks of size $W\tau$ and compute the sum of each block separately. However, we now sum the rounded values rather than the exact input, with a $O(\log \frac{W\tau}{\epsilon})$ -bits counter denoted y . Once the block is completed, we *round its sum* such that it is represented with $O(\log \frac{\tau}{\epsilon})$ bits. Note that this second rounding is done for the entire block's sum while we still have the “exact” sum of rounded fractions. Thus, we *propagate* the second rounding error to the following block. An illustration of our algorithm appears in Figure 3. Here, $\text{Round}_v(z)$ refers to rounding a fractional number $z \in [0, 1]$ into the closest number \tilde{z} such that $2^v \cdot \tilde{z} \in \mathbb{N}$. Algorithm 2 provides pseudo code for the algorithm, which uses the following variables:



■ **Figure 3** An illustration of our 2-stage rounding technique. Arriving elements are rounded to $(\lceil \log \epsilon^{-1} \rceil + 1)$ bits. We then sum the rounded fractions of each block and round the resulting sum into $\lceil \log \frac{\tau}{\epsilon} \rceil$ bits. The second rounding error is propagated to the next block.

Algorithm 2 (W, τ, ϵ) -ADDITIVE SUMMING Algorithm.

Initialization: $y = 0, b = 0, B = 0, i = 0, c = 0$.

- 1: **function** UPDATE(x)
- 2: $x' \leftarrow \text{Round}_{v_1}\left(\frac{x}{R}\right)$ ▷ Round $\left(\frac{x}{R}\right)$ such that $x' \cdot 2^{v_1} \in \mathbb{N}$
- 3: $y \leftarrow y + x'$
- 4: $c \leftarrow (c + 1) \bmod W\tau$
- 5: **if** $c = 0$ **then** ▷ End of block
- 6: $B \leftarrow B - b_i$
- 7: $b_i \leftarrow \text{Round}_{v_2}\left(\frac{y}{W\tau}\right)$ ▷ Replace the value for the block that has left the window.
- 8: $B \leftarrow B + b_i$
- 9: $y \leftarrow y - W\tau \cdot b_i$
- 10: $i \leftarrow (i + 1) \bmod \tau^{-1}$
- 11: **function** OUTPUT
- 12: **return** $\langle R \cdot (W\tau \cdot B + y), c \rangle$

1. y - a fixed point variable that uses $\lceil \log W\tau \rceil + 1$ bits to store its integral part and additional $v_1 \triangleq \lceil \log \epsilon^{-1} \rceil + 1$ bits for storing the fractional part.
2. b - a cyclic array that contains τ^{-1} elements, each of which takes $v_2 \triangleq \lceil \log \frac{\tau}{\epsilon} \rceil$ bits.
3. B - keeps the sum of elements in b and is represented using $\log(\tau^{-1} \lceil \log \frac{\tau}{\epsilon} \rceil + 1)$ bits.
4. i - the index variable used for tracking the oldest block in b .
5. c - a variable that keeps the offset within the $W\tau$ sized block.

We now analyze the memory consumption of Algorithm 2.

► **Theorem 12.** *Algorithm 2 uses $\tau^{-1} \log\left(\frac{\tau}{\epsilon}\right) (1 + o(1)) + 2 \log(W/\epsilon)$ bits.*

Proof. y requires $\log\left(\frac{W\tau}{\epsilon}\right) + O(1)$ bits; b requires another $\tau^{-1} \lceil \log\left(\frac{\tau}{\epsilon}\right) \rceil$; B takes additional $\log(\tau^{-1} \lceil \log \frac{\tau}{\epsilon} \rceil + 1)$ bits; i adds $\lceil \log \tau^{-1} \rceil$ bits, while c is represented with $\lceil \log W\tau \rceil$ bits. Overall, the space requirement is $\tau^{-1} \lceil \log\left(\frac{\tau}{\epsilon}\right) \rceil (1 + o(1)) + 2 \log(W/\epsilon)$ bits. ◀

► **Corollary 13.** *Let $\mathcal{B} \triangleq \max\{\log(W/\epsilon) - O(1), \lceil \tau^{-1}/2 \rceil \log \lfloor \tau/2\epsilon + 1 \rfloor\}$ be the (W, τ, ϵ) -ADDITIVE SUMMING space lower bound of Theorem 3, then Algorithm 2 uses $\mathcal{B} \cdot (4 + o(1))$ bits.*

Finally, Theorem 14 shows that Algorithm 2 is correct. The proof is deferred to the full version of this paper [5].

► **Theorem 14.** *Algorithm 2 solves the (W, τ, ϵ) -ADDITIVE SUMMING problem.*

Algorithm 3 (W, τ, ϵ) -MULTIPLICATIVE SUMMING Algorithm.

Initialization: $y = 0, b = \bar{0}, B = 0, i = 0, c = 0$.

- 1: **function** UPDATE(x)
- 2: $y \leftarrow y + x$
- 3: $c \leftarrow (c + 1) \bmod W\tau$
- 4: **if** $c = 0$ **then** ▷ End of block
- 5: $\rho \leftarrow \lfloor \log_{(1+\epsilon/2)} y \rfloor$ ▷ If $y = 0$ we use $\rho = -\infty$ and $(1 + \epsilon/2)^\rho = 0$
- 6: $B \leftarrow B - ((1 + \epsilon/2)^{b_i})_{\downarrow} + ((1 + \epsilon/2)^\rho)_{\downarrow}$
- 7: $b_i \leftarrow \rho$
- 8: $y \leftarrow 0$
- 9: $i \leftarrow (i + 1) \bmod \tau^{-1}$
- 10: **function** OUTPUT
- 11: **return** $\langle B + y, c \rangle$

4.3 (W, τ, ϵ) -Multiplicative Summing

In this section, we present Algorithm 3 that provides a $(1 + \epsilon)$ multiplicative approximation of the SLACK SUMMING problem. Compared to Algorithm 1, we achieve a space reduction by representing each sum of $W\tau$ elements using $O(\log \log(RW\tau) + \log \epsilon^{-1})$ bits. Specifically, when a block ends, if its sum was y , we store $\rho = \lfloor \log_{(1+\epsilon/2)} y \rfloor$ (we allow a value of $-\infty$ for ρ if $y = 0$). To achieve $O(1)$ OUTPUT, we also store an approximate window sum B , which is now a *fixed point* fractional variable with $O(\log RW)$ bits for its integral part and additional $O(\log \epsilon^{-1})$ bits for storing a fraction. To update B 's value for a new ρ , we *round down* the value of $(1 + \epsilon)^\rho$. Specifically, for a real number x , we denote $(x)_{\downarrow} \triangleq \lfloor x \cdot k \rfloor / k$, for $k \triangleq \lceil \frac{4}{\epsilon} \rceil$. Our pseudo code appears in Algorithm 3. The algorithm requires $O(\tau^{-1} (\log \log(RW\tau) + \log \epsilon^{-1}) + \log RW)$ bits of space and is memory optimal when $R = W^{O(1)}$ and $\tau = \Omega\left(\frac{1}{\log RW}\right)$.

The full analysis of Algorithm 3 is deferred to the full version of this paper [5]. Next, we present an alternative (W, τ, ϵ) -MULTIPLICATIVE SUMMING algorithm that achieves optimal space consumption for $\tau = \Theta(1)$, regardless of the value of R .

Improved (W, τ, ϵ) -Multiplicative Summing for $\tau = \Theta(1)$

Algorithm 4 is more space efficient than Algorithm 3 but has a query time of $O(\tau^{-1})$. For $\tau = \Theta(1)$, Algorithm 4 is memory optimal *and* supports constant time queries even if $R = W^{\omega(1)}$; for this case, Algorithm 3 requires $\Omega(\log R)$ bits which is sub optimal.

Intuitively, we shave the $\Omega(\log R)$ bits from the space requirement of Algorithm 3 using an approximate representation for our y variable and by not keeping the B variable that allowed $O(1)$ time queries regardless of the value of τ . To avoid using $\Omega(\log R)$ bits in y , we use a *fixed point* representation in which $O(\log \epsilon^{-1} + \log \log(RW\tau))$ bits are allocated for its integral part and another $O(\log W\tau)$ for the fractional part. The goal of y is still to approximate the sum of the elements within a block, but now we aim for the sum to be approximately $(1 + \epsilon/3)^y$. Whenever a block ends, we store only the integral part of y in our cyclic array b to save space. When queried, we compute an estimate for the sum using all of the values in b , which makes our query procedure take $O(\log \tau^{-1})$ time. To use the fixed point structure of y , we use the operator $(\cdot)_{\Downarrow}$ that rounds a real number x into $(x)_{\Downarrow} \triangleq \lfloor x \cdot W\tau \rfloor / W\tau$. We denote $\log_{(1+\epsilon/3)}(0) = -\infty$, $(-\infty)_{\Downarrow} = -\infty$, $\lfloor -\infty \rfloor = -\infty$ and $(1 + \epsilon/3)^{-\infty} = 0$.

Algorithm 4 (W, τ, ϵ) -MULTIPLICATIVE SUMMING Algorithm for $\tau = \Theta(1)$.

```

Initialization:  $y = -\infty, b = \bar{0}, i = 0, c = 0$ .
1: function UPDATE( $x$ )
2:    $y \leftarrow (\log_{(1+\epsilon/3)}(x + (1 + \epsilon/3)^y))_{\underline{\mathbb{D}}}$ 
3:    $c \leftarrow (c + 1) \bmod W\tau$ 
4:   if  $c = 0$  then ▷ End of block
5:      $b_i \leftarrow \lfloor y \rfloor$ 
6:      $y \leftarrow -\infty$ 
7:      $i \leftarrow (i + 1) \bmod \tau^{-1}$ 
8: function OUTPUT
9:   return  $\langle (1 + \epsilon/3)^y + \sum_{i=0}^{\tau^{-1}-1} (1 + \epsilon/3)^{b_i}, c \rangle$ 

```

In the full version of this paper [5], we prove the following theorem:

► **Theorem 15.** *For $\tau = \Theta(1)$, Algorithm 4 processes elements and answers queries in $O(1)$ time, uses $O(\log(W/\epsilon) + \log \log R)$ bits, and is asymptotically optimal.*

4.4 The Mean of a Slack Window

For some applications there is value in knowing the *mean* of a slack window. For example, a load balancer may be interested in the average transmission throughput. In exact windows, the sum and the mean can be derived from each other as the window size is constant. In slack windows, the window size changes but our algorithms also return the current slack offset $0 \leq c < W\tau$. That is, by dividing \hat{S} by $W + c$ we get an estimation of the mean (we assume that stream size is larger than W). Specifically, Algorithm 1 provides the exact mean; Algorithm 2 approximates it with $R\epsilon$ additive error, while Algorithm 3 yields a $(1 + \epsilon)$ multiplicative approximation.

5 Other Measurements over Slack Windows

We now explore the benefits of the slack model for other problems.

Maximum. While maintaining the maximum of a sliding window can be useful for applications such as anomaly detection [26, 21], tracking it over an exact window is often infeasible. Specifically, any algorithms for a maximum over an (exact) window must use $\Omega(W \log(R/W))$ bits [14]. The following theorem shows that we can get a much more efficient algorithm for slack windows. The proof appears in the full version of this paper [5]. Observe the the following bounds match for τ values that are not too small ($\tau = R^{\Omega(1)-1}$).

► **Theorem 16.** *Tracking the maximum over a slack window deterministically requires $O(\tau^{-1} \log R)$ and $\Omega(\tau^{-1} \log R\tau)$ bits.*

Standard-Deviation. Building on the ability of our summing algorithms to provide the size of the slack window that they approximate, we can compute standard deviations over slack windows. Intuitively, the standard deviation of the window can be expressed as

$$\sigma_W \triangleq \sqrt{\frac{\sum_{x \in \overline{W}} (x - m_{\overline{W}})^2}{|\overline{W}| - 1}} = \sqrt{\frac{\sum_{x \in \overline{W}} x^2 - 2m_{\overline{W}} \sum_{x \in \overline{W}} x + W \cdot m_{\overline{W}}^2}{|\overline{W}| - 1}} = \sqrt{\frac{\sum_{x \in \overline{W}} x^2 - \overline{W} \cdot m_{\overline{W}}^2}{|\overline{W}| - 1}},$$

there \overline{W} is the slack window and $m_{\overline{W}}$ is its mean. We can then use two slack summing instances to track $\sum_{x \in \overline{W}} x^2$ and $m_{\overline{W}} = |\overline{W}|^{-1} \sum_{x \in \overline{W}} x$. This gives us an algorithm that computes the exact standard deviation over slack windows using $O(\tau^{-1} \log(RW\tau))$ space. Similarly, by using approximate rather than exact summing solutions we can compute a $(1 + \epsilon)$ multiplicative approximation for the standard deviation using $O(\tau^{-1}(\log \epsilon^{-1} + \log \log(RW\tau)) + \log W)$ bits, or an $R\epsilon$ -additive approximation using $O(\tau^{-1} \log(\frac{\tau}{\epsilon}) + \log W)$ space. We expand on this further in the full version of this paper [5].

General-Summing. GENERAL-SUMMING is similar to BASIC-SUMMING, except that the integers can be in the range $\{-R, \dots, R\}$. That is, we now allow for negative elements as well. Datar et al. [14] proved that General Sum requires $\Omega(W)$ bits, even for $R = 1$ and constant factor approximation. In contrast, our exact summing algorithm from section 4.1 trivially generalizes to GENERAL-SUMMING and allows exact solution over slack windows.

Count-Distinct. Estimating the number of **distinct** elements in a stream is another useful metric. In networking, the packet header is used to identify different flows, and it is useful to know how many distinct of them are currently active. A sudden spike in the number of active flows is often an indication of a threat to the network. It may indicate the propagation of a worm or virus, port scans that are used to detect vulnerabilities in the system and even *Distributed Denial of Service (DDoS)* attacks [12, 18, 20].

Here, we have studied the memory reduction that can be obtained by following a similar flow to our summing algorithms – we break the stream into $W\tau$ sized blocks and run the state of the art approximation algorithm on each block separately. Luckily, count distinct algorithms are *mergable* [1]. That is, we can merge the summaries for each block to obtain an estimation of the number of distinct items in the union of the blocks. In the full version of this paper [5], we show that this approach yields an algorithm with superior space and query time compared to the state of the art algorithms for counting distinct elements over sliding windows [11, 19]. Formally, we prove the following theorem.

► **Theorem 17.** *For $\tau = \Theta(1)$ and any fixed $m > 0$, there exists an algorithm that uses $O(m)$ space, performs updates in constant time and answers queries in time $O(m)$, such that the result approximates a window whose size is in $[W, W(1 + \tau)]$; the resulting estimation is asymptotically unbiased and has a standard deviation of $\sigma = O(\frac{1}{\sqrt{m}})$. State of the art approaches for exact windows [11, 19] require $O(m \log(W/m))$ space and $O(m \log(W/m))$ time per query for a similar standard deviation.*

6 Discussion

In this work we have explored the slack window model for multiple streaming problems. We have shown that it enables asymptotic space and time improvements. Particularly, introducing slack enables logarithmic space exact algorithms for certain problems such as MAXIMUM and GENERAL-SUMMING. In contrast, these problems do not admit sub-linear space *approximations* in the exact window model. Even in problems that do have sub-linear space approximations such as STANDARD-DEVIATION and COUNT-DISTINCT, adding slack asymptotically improves the space requirement and allows for constant time updates.

Much of our work has focused on the classic BASIC-SUMMING problem. Based on our findings, we argue that allowing a slack in the window size is an attractive approximation axis as it enables greater space reductions compared to an error in the sum. As an example, for a fixed ϵ value, computing a $(1 + \epsilon)$ -multiplicative approximation requires $\Omega(\log(RW) \log W)$

space [14]. Conversely, a $(1 + \tau)$ multiplicative error in the window size, for a constant τ , allows summing using $\Theta(\log(RW))$ bits – same as in summing W elements without sliding windows! Given that for exact windows randomized algorithms have the same asymptotic complexity as deterministic ones [3, 14], we expect randomization to have limited benefits for slack windows as well.

References

- 1 Pankaj K. Agarwal, Graham Cormode, Zengfeng Huang, Jeff Phillips, Zhewei Wei, and Ke Yi. Mergeable summaries. In *ACM PODS*, 2012.
- 2 Mohammad Alizadeh, Tom Edsall, Sarang Dharmapurikar, Ramanan Vaidyanathan, Kevin Chu, Andy Fingerhut, Vinh The Lam, Francis Matus, Rong Pan, Navindra Yadav, and George Varghese. Conga: Distributed congestion-aware load balancing for datacenters. In *ACM SIGCOMM 2014 Conference, SIGCOMM'14, Chicago, IL, USA, August 17-22, 2014*, ACM SIGCOMM 2014, 2014. doi:10.1145/2619239.2626316.
- 3 Ran Ben Basat, Gil Einziger, Roy Friedman, and Yaron Kassner. Efficient Summing over Sliding Windows. In *SWAT*, 2016.
- 4 Ran Ben Basat, Gil Einziger, Roy Friedman, Marcelo Caggiani Luizelli, and Erez Waisbard. Constant time updates in hierarchical heavy hitters. In *ACM SIGCOMM*, 2017.
- 5 R. Ben Basat, G. Einziger, and R. Friedman. Give Me Some Slack: Efficient Network Measurements. *ArXiv e-prints*, 2018. arXiv:1703.01166.
- 6 Ran Ben-Basat, Gil Einziger, Roy Friedman, and Yaron Kassner. Heavy hitters in streams and sliding windows. In *IEEE INFOCOM*, 2016.
- 7 Ran Ben-Basat, Gil Einziger, Roy Friedman, and Yaron Kassner. Optimal elephant flow detection. In *IEEE INFOCOM*, 2017.
- 8 Ran Ben-Basat, Gil Einziger, Roy Friedman, and Yaron Kassner. Randomized admission policy for efficient top-k and frequency estimation. In *IEEE INFOCOM*, 2017.
- 9 Theophilus Benson, Ashok Anand, Aditya Akella, and Ming Zhang. Microte: Fine grained traffic engineering for data centers. In *ACM CoNEXT*, 2011.
- 10 Flavio Bonomi, Michael Mitzenmacher, Rina Panigrahy, Sushil Singh, and George Varghese. An improved construction for counting bloom filters. In *ESA*, 2006.
- 11 Y. Chabchoub and G. Hebrail. Sliding hyperloglog: Estimating cardinality in a data stream over a sliding window. In *2010 IEEE ICDM Workshops*, 2010.
- 12 Varun Chandola, Arindam Banerjee, and Vipin Kumar. Anomaly detection: A survey. In *ACM CSUR*, 2007.
- 13 Min Chen and Shigang Chen. Counter tree: A scalable counter architecture for per-flow traffic measurement. In *IEEE ICNP*, 2015.
- 14 Mayur Datar, Aristides Gionis, Piotr Indyk, and Rajeev Motwani. Maintaining stream statistics over sliding windows. *SIAM Journal of Computing*, 2002.
- 15 G. Einziger and R. Friedman. TinyLFU: A highly efficient cache admission policy. In *PDP 2014*, 2014.
- 16 Gil Einziger, Benny Fellman, and Yaron Kassner. Independent counter estimation buckets. In *IEEE INFOCOM*, 2015.
- 17 Gil Einziger and Roy Friedman. Counting with TinyTable: Every Bit Counts! In *ICDCN 2016*, 2016.
- 18 Cristian Estan, George Varghese, and Mike Fisk. Bitmap algorithms for counting active flows on high speed links. In *ACM IMC*, 2003.
- 19 Éric Fusy and Frédéric Giroire. Estimating the number of active flows in a data stream over a sliding window. In *ANALCO*, 2007.

- 20 Sumit Ganguly, Minos Garofalakis, Rajeev Rastogi, and Krishan Sabnani. Streaming algorithms for robust, real-time detection of ddos attacks. In *27th IEEE International Conference on Distributed Computing Systems (ICDCS 2007), June 25-29, 2007, Toronto, Ontario, Canada*, ICDCS, 2007.
- 21 Pedro Garcia-Teodoro, Jesús E. Díaz-Verdejo, Gabriel Maciá-Fernández, and E. Vázquez. Anomaly-based network intrusion detection: Techniques, systems and challenges. *Computers and Security*, 2009.
- 22 Phillip B. Gibbons and Srikanta Tirthapura. Distributed streams algorithms for sliding windows. In *SPAA*, 2002.
- 23 Nan Hua, Bill Lin, Jun (Jim) Xu, and Haiquan (Chuck) Zhao. Brick: A novel exact active statistics counter architecture. In *ACM/IEEE ANCS*, 2008.
- 24 Abdul Kabbani, Mohammad Alizadeh, Masato Yasuda, Rong Pan, and Balaji Prabhakar. Af-qcn: Approximate fairness with quantized congestion notification for multi-tenanted data centers. In *IEEE HOTI*, 2010.
- 25 Yang Liu, Wenji Chen, and Yong Guan. Near-optimal approximate membership query over time-decaying windows. In *IEEE INFOCOM*, 2013.
- 26 B. Mukherjee, L.T. Heberlein, and K.N. Levitt. Network intrusion detection. *Network, IEEE*, 1994.
- 27 Moni Naor and Eylon Yogev. Sliding bloom filters. In *ISAAC*. Springer, 2013.
- 28 Erez Tsidon, Iddo Hanniel, and Isaac Keslassy. Estimators also need shared values to grow together. In *IEEE INFOCOM*, 2012.
- 29 Hao Wang, H. Zhao, Bill Lin, and Jun Xu. Dram-based statistics counter array architecture with performance guarantee. *IEEE/ACM Transactions on Networking*, 2012.
- 30 Li Yang, Wu Hao, Pan Tian, Dai Huichen, Lu Jianyuan, and Liu Bin. Case: Cache-assisted stretchable estimator for high speed per-flow measurement. In *IEEE INFOCOM*, 2016.
- 31 L. Ying, R. Srikant, and X. Kang. The power of slightly more than one sample in randomized load balancing. In *IEEE INFOCOM*, 2015.