

# Thread Cooperation in Multicore Architectures for Frequency Counting Over Multiple Data Streams

Sudipto Das   Shyam Antony   Divyakant Agrawal   Amr El Abbadi  
Department of Computer Science  
University of California, Santa Barbara  
Santa Barbara, CA 93106-5110, USA  
{sudipto, shyam, agrawal, amr}@cs.ucsb.edu

## ABSTRACT

Many real-world data stream analysis applications such as *network monitoring*, *click stream analysis*, and others require combining multiple streams of data arriving from multiple sources. This is referred to as *multi-stream analysis*. To deal with high stream arrival rates, it is desirable that such systems be capable of supporting very high processing throughput. The advent of multicore processors and powerful servers driven by these processors calls for efficient parallel designs that can effectively utilize the parallelism of the multicores, since performance improvement is possible only through effective parallelism. In this paper, we address the problem of parallelizing *multi-stream analysis* in the context of multicore processors. Specifically, we concentrate on parallelizing frequent elements, *top-k*, and frequency counting over multiple streams. We discuss the challenges in designing an efficient parallel system for multi-stream processing. Our evaluation and analysis reveals that traditional “contention” based locking results in excessive overhead and wait, which in turn leads to severe performance degradation in modern multicore architectures. Based on our analysis, we propose a “cooperation” based locking paradigm for efficient parallelization of frequency counting. The proposed “cooperation” based paradigm removes waits associated with synchronization, and allows replacing locks by much cheaper atomic synchronization primitives. Our implementation of the proposed paradigm to parallelize a well known frequency counting algorithm shows the benefits of the proposed “cooperation” based locking paradigm when compared to the traditional “contention” based locking paradigm. In our experiments, the proposed “cooperation” based design outperforms the traditional “contention” based design by a factor of 2 – 5.5X for synthetic zipfian data sets.

## 1. INTRODUCTION

Data stream analysis forms an important class of applications where data is streaming in, and processing has to be done in real time. In traditional designs [23], data was considered to be arriving from a single source, and a plethora of sequential designs for the *single-stream* problem have been proposed [19, 20, 23]. But a number of common data stream applications require streams from

multiple sources to be combined to answer queries, and we refer to this as *multi-stream* analysis. Multi stream analysis arises in many contexts, including large scale web advertising networks, the internet, sensor networks, and real time data analysis [6]. In large scale web-based advertising networks, clicks originating from different internet hosts result in multiple click streams that need to be merged at a single server for analysis and identification of suspicious publishers or IP addresses [21]. In the Internet, network monitors at core routers process streams of packets arriving from different origins for accounting purposes, and detection of malicious activities such as Distributed Denial of Service attacks [29]. In sensor networks, data streams from multiple sensors are typically combined at a single server for analysis, mining, and querying [2, 18]. Similarly, analysis of web server logs, monitoring calls in cellular networks, real time analysis of credit card and ATM transactions etc. require real time multi-stream analysis [4].

A number of high-throughput distributed stream processing systems have been proposed, such as Borealis [1] and Medusa [3]. These systems distribute the load on multiple distributed servers to ensure high throughput. But with the advent of multicore processors and extremely powerful servers driven by these processors, even centralized servers are now capable of delivering high throughput stream processing. For instance, the Sun SPARC Enterprise T5220 server powered by a UltraSPARC T2 processor [25] and supporting 1Gbps network connectivity has 64 hardware thread contexts which can be used for efficient and high throughput stream processing. Even standard workstations are now powered by processors with 4 – 8 cores [16], and provide great potential for parallel processing. Efficient parallel designs are therefore needed to exploit such parallel hardware to support the processing of large numbers of streams with high and varying arrival rates. Note that stream processing systems have considerable data dependencies and shared data structures, and efficient parallel systems require thoughtful and efficient design.

**Multi-Stream Analysis and Multicores:** *Multi-stream analysis* has inherent parallel processing needs since stream elements are arriving from multiple sources at possibly variable rates, and the streams need to be combined for answering queries on the union. In the past when processors consisted of only one execution unit, the most efficient approach was to exploit some form of operating system support such as `poll` or `select` to transform the problem into a *single stream* problem. It is widely acknowledged that these system calls introduce a significant overhead into the system. But in a sequential world, this was a preferred solution. But with the advent of multicore processors [16, 25], each processor now has multiple threads that can execute in parallel, and a system can improve its performance by exploiting the inherent parallelism to process each incoming stream using one or more threads. In this

Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the VLDB copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Very Large Data Base Endowment. To copy otherwise, or to republish, to post on servers or to redistribute to lists, requires a fee and/or special permission from the publisher, ACM.

VLDB '09, August 24-28, 2009, Lyon, France

Copyright 2009 VLDB Endowment, ACM 000-0-00000-000-0/00/00.

paper we investigate the challenges of such a parallel design.

**Parallelizing Frequency Counting.** Efficient parallel designs are the only means of improving performance in a multicore architecture. But, due to data dependencies arising from updates to a common structure, the problem is not “embarrassingly parallel”. Thus, adding more threads will *not* lead to linear speedup. As it turns out, similar to many parallel algorithms and structures, parallelizing frequency counting is non-trivial, and actually it is very hard to eke out parallelism. Our experiments with a parallel frequency counting system designed using the traditional notion of “contention” based locking reveal the high overhead of a parallel system. Frequently, contending threads that share a resource need to *wait* for it if it is being exclusively held by some other thread. This results in wasted CPU cycles when the (hardware) threads do not perform useful work. To minimize the waits and the wasted CPU cycles, we propose a “cooperation” based locking paradigm where the threads “cooperate” (and not “contend”) when sharing resources. Whenever a resource sought by a thread  $T_i$  is not available, then  $T_i$ ’s work is “delegated” to the thread  $T_j$  currently holding the resource, and  $T_i$  can move to its next job.  $T_i$ ’s request will eventually be processed by  $T_j$ . As a result, even though locks and shared resources are still present, *waits* associated with the locks have been eliminated. This allows efficient usage of the inherent parallelism, and results in higher processing throughput. Additionally, the proposed design paradigm allows replacing locks by much cheaper atomic synchronization primitives. In a previous work [8], we demonstrate the effectiveness of the “thread cooperation” paradigm for *intra-operator* parallelism of stream operations, and in this paper, we use this paradigm for *multi-stream analysis*. Concepts similar to “thread cooperation”, such as the escrow transactional model [26], have earlier been used in the field of databases.

In this paper, we parallelize a standard frequency counting algorithm which forms the basis for frequent elements and top- $k$  queries. Frequent elements [5, 19, 20] and top- $k$  [7, 20, 24] queries are an important class of queries for stream analysis applications, and the research community has proposed several algorithms for answering such queries efficiently. A frequent elements query returns all the elements whose frequency of occurrence is above a certain threshold. For example, a query of the form “advertisements that are clicked more than 0.1% of the total clicks” is a frequent elements query. On the other hand, a top- $k$  query returns the  $k$  elements with the highest frequency. Again, a query of the form “top-25 most clicked advertisements” is a top- $k$  query.

#### Contributions:

- We propose and formalize the notion of “thread cooperation” for parallel analysis of multiple data streams. This concept of “thread cooperation” removes *waits* associated with locks, and will find use in various other “lock-based” designs.

- We propose a parallel design of the *Space Saving* algorithm [20], which Cormode et al. [5] have shown to have best performance amongst a number of *counter based* algorithms. Even though our design is based on *Space Saving*, the proposed paradigm can easily be augmented to accommodate other standard frequency counting algorithms such as *Lossy Counting* [19].

- We provide correctness proof sketches of the proposed design, and analyze the performance of the proposed “cooperation based” design. Our experiments show a factor of 2 – 5.5X improvement over traditional “contention” based locking paradigm.

**Organization:** Section 2 provides a survey of related work, and a brief description of the *Space Saving* [20] algorithm which we select to parallelize. Section 3 analyzes some common approaches towards designing a parallel system, and our analysis motivates the need for a new design paradigm. Section 4 formalizes the “thread

cooperation” paradigm and describes the implementation details of the parallel system while providing arguments for correctness. Section 5 experimentally evaluates the proposed “cooperation” based design and compares it with traditional “contention” based design. Section 6 concludes the paper.

## 2. RELATED WORK

### 2.1 Data Stream Processing

A class of frequency counting algorithms, referred to as *counter based* techniques [19, 20], monitor a subset of the stream elements and maintain an approximate frequency count of the elements. The goal is to guarantee high accuracy while having a small memory footprint. Different approaches use different heuristics to determine the set of elements to be monitored, and thus limit the space overhead. For example, in *Lossy Counting* [19], the stream is divided into rounds, and at the end of every round, potentially infrequent elements are deleted. This  $\epsilon$ -approximate algorithm has a space bound of  $O(\frac{1}{\epsilon} \log(\epsilon N))$ , where  $N$  is the length of the stream and  $\epsilon$  is the error bound. The *Space Saving* algorithm [20], on the other hand, uses a different heuristic to limit space, details of which are provided in Section 2.2.

Different solutions have also been suggested for answering top- $k$  queries [7, 24]. Mouratidis et al. [24] suggest the use of geometrical properties to determine the  $k$ -skyband, and use this abstraction to answer top- $k$  queries, whereas Das et al. [7] propose a technique which is capable of answering ad-hoc top- $k$  queries, i.e., the algorithm does not need prior knowledge of the attribute on which the top- $k$  queries have to be answered.

There has also been considerable research in distributed frequent elements [18], distributed top- $k$  monitoring [2], and distributed stream processing systems [1, 3]. At first glance, it might seem that the multiple streams problem can be formulated as distributed streams, and the entire body of relevant literature can be used. But there is a subtle difference between distributed streams and the multiple streams problem we are considering. In distributed stream processing, since the streams are processed on distributed nodes, memory and processing is independent, and the onus is on reducing the communication overhead. On the other hand, in the multi-stream case, the streams are processed on the same machine sharing processing and memory, and the onus is on efficient processing while effectively sharing the processor and memory among the different streams. Additionally, distributed stream processing systems have to deal with fault-tolerance, load balancing and many related issues of distributed systems [1], while these issues are not relevant in multi-stream analysis on a single powerful server.

The advent of modern multicore architectures [16, 25] have opened new frontiers, and their ubiquitous presence calls for algorithms that can efficiently exploit the parallelism inherent to these architectures. Although much research has been done in the stored database arena for exploiting parallelism [12, 13], very little research has focussed on stream operators. Gedik et al. [10] propose the use of Cell processors for parallel windowed stream joins. The proposed technique is targeted to the Cell architecture, and leverages specific features of Cell processors to improve performance. The growing demand for high throughput stream processing calls for designing other efficient parallel algorithms for stream processing, and this is the focus of this paper.

### 2.2 Space Saving Algorithm

We select to parallelize the *Space Saving* algorithm [20] which has been demonstrated to have the best throughput amongst its class of frequency counting algorithms [5]. An interesting property of

**Algorithm 2.1** *Space Saving* algorithm

---

```

1:  $maxCounters \leftarrow 1/\epsilon, numMonitored \leftarrow 0$ 
2: for each element  $(e)$  in the stream do
3:   /* Check if  $(e)$  is already being monitored */
4:   if (LOOKUP( $(e)$ )) then
5:      $IncrementCounter(e)$ 
6:   else
7:     if ( $numMonitored < maxCounters$ ) then
8:        $AddElement(e, 1); numMonitored++$ 
9:     else
10:       $Overwrite(Minimum\ frequency\ element, e)$ 

```

---

**Table 1: Main Operations in Space Saving**


Operation	Description
LOOKUP( $e$ )	Check whether element $e$ is being monitored
$IncrementCounter(e)$	Increase the frequency of $e$
$AddElement(e, freq)$	Add element $e$ with frequency $freq$
$Overwrite(min, e)$	Overwrite the minimum frequency element with $e$

the algorithm is that it is deterministic and provides tight space bounds corresponding to the user specified error bound  $\epsilon$  which determines the number of elements and counters that need to be monitored. *Space Saving* monitors only  $O(\frac{1}{\epsilon})$  counters for providing  $\epsilon$ -approximate answers. Algorithm 2.1 gives an overview of *Space Saving*. The main operations performed by the algorithm have been listed in Table 1. The algorithm monitors a subset of the stream elements which we refer to as the **Monitored Set**. If the element being processed is already being monitored, then its count is incremented (*IncrementCounter*). Otherwise, if the number of monitored elements is less than the maximum bound ( $O(1/\epsilon)$ ), then the element is added to the monitored set (*AddElement*), else the current element *Overwrites* the element with minimum frequency and increments its count by one. Overwriting the minimum frequency element is a heuristic used by the algorithm to limit space. The intuition being that the minimum frequency element is least likely to be of interest to a frequent elements or top- $k$  query. Thus, the space bound of the algorithm is  $\min(O(\frac{1}{\epsilon}), |A|)$ , where  $|A|$  is the size of the alphabet of the stream.

For *Overwriting*, this algorithm needs to track the minimum frequency element. The *Stream Summary* structure [20] is used for this purpose. This structure consists of a *doubly-linked list* of frequency buckets which are sorted by frequency. Each bucket contains a list of elements which has the same frequency as that of the bucket. A nice property of this structure is that it maintains the elements sorted by frequency in  $O(1)$  time per element. This allows answering both frequent elements and top- $k$  queries using the same structure. For every element being processed, the algorithm looks up an element in the *Search Structure* (LOOKUP), and then updates the element in the *Stream Summary* structure. For lookup, the algorithm needs to have an efficient *search structure* that can be integrated with the *Stream Summary* structure, and a hash table is most suited for this purpose. The **Monitored Set** is thus represented by a combination of *Search Structure* and *Stream Summary*. Therefore, LOOKUP has to be supported by the hash table and the rest of the operation in Table 1 must be supported by *Stream Summary*.

### 3. INTUITIVE PARALLEL DESIGNS

In this section, we start by examining several possible intuitive designs for parallelizing frequency counting for *multi-stream analysis*. In these designs, each stream is assigned a thread which processes the elements. Recall that the *Space Saving* algorithm updates a structure (**Monitored Set**) while processing the elements.

**Figure 1: Profiling of the independent design using hierarchical merge. The query rate was set to 1 query every 50,000 elements.**

So the manner in which the threads share the structure determines the design of the system. We explore two straightforward designs, provide a brief experimental evaluation of these techniques, and explore some extensions to the two basic techniques. Our experimental setup uses an Intel quad core processor [16], and synthetically generated data following Zipfian distribution [30] where the zipfian factor  $\alpha$  is varied from 1.5 to 3.0 in steps of 0.5. The parameter  $\alpha$  controls the distribution, smaller  $\alpha$  corresponds to a uniform distribution while higher  $\alpha$  corresponds to a more skewed distribution. More details of the experimental setup appear in Section 5. Due to space limitations, we provide selected results in this section, and additional results can be found in [9].


#### 3.1 Independent Structures

This straightforward design corresponds to the *shared nothing* paradigm, where threads do not share any data or state information, and each thread has its own independent local structure. The idea is to simulate sequential execution, and run multiple copies of the same algorithm executing on different streams. Each stream has a local structure, and therefore there is no need for synchronization. If there are  $n$  different streams, then there are  $n$  different local structures, and these structures need to be merged periodically into a global structure so that queries can be answered from the global structure. Even though each thread processes only a fraction of the entire stream, since the minimum element in the local structure is not guaranteed to be the global minimum element, if the local structure is not large enough, it might lead to higher error. Lemma 3.1 (proof in [9]) provides a lower bound for the size of the local structures in order that the global error bound is preserved.

LEMMA 3.1. *For  $\epsilon$ -approximate answers, the size of every local structure must be at least  $O(1/\epsilon)$ .*

COROLLARY 3.2. *Using independent structures for  $n$  streams, the total space complexity for  $\epsilon$ -approximate answers is  $O(n/\epsilon)$ .*

Therefore, it follows from Corollary 3.2 that the independent design has high space overhead. In addition, the local structures need to be merged periodically to obtain the global structure from where the queries can be answered. Since the *Stream Summary* structure can be used to efficiently answer both frequent elements as well as top- $k$  queries, in order for the merged global structure to retain those properties while remaining  $\epsilon$ -approximate, all elements of every local structure need to be merged. Different heuristics (discussed in our technical report [9]) can be used to reduce the merge overhead, but then either the error increases, or the properties of *Stream Summary* are not preserved. Furthermore, the merge overhead is dependent on the query frequency – the greater the query frequency, the higher the number of merges and higher is the overhead. The merge overhead also increases with any increase in the number of parallel streams and threads, and for smaller values of the error bound  $\epsilon$ . Thus, it is intuitive that even though the frequency counting part of the independent design would scale, the



**Figure 2: Profiling of the shared design using “contention” based locking.**

space overhead and the high merge overhead makes this design inefficient with increasing number of streams and query frequency. Our experimental evaluation confirms this, and Figure 1 provides a profile of the time spent by the *independent* design with 1 query per 50,000 elements processed. Along  $x$ -axis is the number of streams (and hence threads) and along  $y$ -axis is the time spent for an operation as percentage of total time. It is evident from Figure 1 that as the number of threads increase, the merge overhead makes this design inefficient. Refer to our technical report [9] for a detailed evaluation of the *independent* design and details about the merge algorithm used.

### 3.2 Shared Structures

The limitations of the independent design are high space overhead and merge overhead. To reduce the overhead, the threads can share the **Monitored Set**, and this forms the basis for the **Shared design** which we discuss in this subsection. Since multiple threads are accessing the same structure, the threads must synchronize, and this synchronization needs to be done at two levels:

- **Element Level Synchronization:** Multiple threads operating on the same element must be serialized so that there is only one thread operating on the element in the *Stream Summary*. This is achieved by acquiring a lock in the *search structure* on the element being processed. *Element level synchronization therefore serializes all threads processing the same element in the stream.*

- **Bucket Level Synchronization:** Since an increment or overwrite operation needs to move an element from one frequency bucket to another bucket in the *Stream Summary* structure, a thread performing this operation needs to obtain locks on the source and the destination buckets. Since there can be several elements within a frequency bucket, a lock on a bucket prevents other threads from operating on any element belonging to that bucket. *Bucket-level synchronization therefore serializes accesses to a frequency bucket.*

The queries (which are only readers) also need to obtain locks so that the writers are blocked while a reader is processing a bucket. In this design, since a single global structure is used, the space complexity remains  $O(\frac{1}{\epsilon})$  and there is no merge overhead. Thus, the *shared* design saves on space and merge cost. But since multiple locks need to be acquired for every element being processed, and due to multiple levels of synchronization, using the traditional “contention” based locking paradigm results in high contention overhead.

Figure 2 provides a profile of the time spent by the shared design. Along  $x$ -axis is the number of streams (and hence threads) and along  $y$ -axis is the time spent for an operation as percentage of total time. Figure 2 reveals that a majority of the time is spent contending either for synchronization at the *element-level* or at the *bucket-level*. Thus, even though the *shared* design has low space and merge overhead, the overhead due to contention and wait for shared resources makes this design inefficient. Multicore architectures have a large number of CPU cores, and threads waiting for

shared resources waste cycles of idle CPU cores, which could have otherwise done useful work.

### 3.3 Discussion

**Extensions of the straightforward techniques:** There are several other simple schemes that extend the techniques proposed above. One possible approach can be to maintain a combination of local and global counters (i.e., a **Hybrid Structure**) to limit the contention (by maintaining local counters for frequent elements and a global structure for infrequent elements) as well as space overhead (no need to replicate relatively infrequent elements). Analysis reveals that this approach is also not very efficient. For a relatively uniform input, most of the items will hit the global structure (since most items are infrequent), thereby degenerating into the *shared design*. For a relatively skewed input, very few elements will hit the global structure (since most elements are infrequent and hence stored locally to reduce contention), thus reducing to the *independent design*. Another extension is a **hash based scheme** where elements are hashed to threads, i.e., same elements go to same threads. Even though this design will need less synchronization, load balancing among threads will be an issue when the input distribution is skewed, as some threads are overloaded, while others are lightly loaded.

**Characteristics of a Good Design:** Ideally, we would like a design with the following properties:

- *Small memory footprint and no merge overhead* similar to the shared design.
- *Low contention overhead* similar to the independent design.

In the next section, we show how “contention” in locking can be reduced by “thread cooperation” so that shared structures can also be used efficiently.

## 4. THREAD COOPERATION PARADIGM


### 4.1 Formalization

As observed in Section 3.3, instead of threads contending for shared resources, they can rather *cooperate*, thereby boosting each other and in turn improving the overall system performance. In this paradigm, even though locks are not eliminated completely, “waits” associated with locks have been eliminated. The delegated request is “queued”, and the type of queueing depends on the semantics of the request being delegated.

Let us consider a system comprised of threads (hardware or software)<sup>1</sup> and the task to be completed is a set of *requests* that are *associative* and *commutative*. The system is said to *make progress* if the input set of *requests* is being consumed. The following principles formally express the *thread cooperation* model:

**PRINCIPLE 4.1. Request Delegation:** *If thread  $T_i$  is trying to acquire a shared resource  $\mathbb{R}$  and it succeeds in acquiring  $\mathbb{R}$ , it will complete its request using  $\mathbb{R}$ . If it fails to acquire  $\mathbb{R}$ , it will “delegate” its request to the thread  $T_j$  that currently has exclusive access to  $\mathbb{R}$ , and “proceeds” to its next request. All other threads trying to acquire  $\mathbb{R}$  will “delegate” their requests to  $T_j$ . Once  $T_j$  finishes its own request, and before it relinquishes control over  $\mathbb{R}$ , it will check for any pending requests on  $\mathbb{R}$  and will relinquish  $\mathbb{R}$  only when all pending delegated requests have been processed.*

<sup>1</sup>A *hardware thread* corresponds to an execution unit with its own set of registers and local resources and is different from a *software thread*. Multiple *software threads* can be mapped to a *hardware thread* to time-share the resources.



**Figure 3: Overview of the parallel design of *Space Saving* using *thread cooperation*.**

PRINCIPLE 4.2. *Minimal Existence*: Once a thread has acquired a shared resource, it will abstain from acquiring other shared resources.

Note that *requests* are *delegated* in two different scenarios. *First*, when a shared resource is not *available*, and *second*, when a thread has already acquired a shared resource, and needs another shared resource. In either case, the threads rely on *cooperation* to make progress. To guarantee correctness, an implementation must satisfy the following invariant:

INVARIANT 4.1. *Fulfillment Guarantee*: Once a thread has delegated a request, the request is neither lost nor left unfulfilled.

There are multiple benefits of the *cooperation* based approach:

- **Removal of Waits**: Whenever there is contention for a shared resource, the threads follow Principle 4.1 and do not *wait*. This prevents wastage of useful computational resources and allows threads to make progress.
- **Arbitration Overhead Removal**: Since the threads do not contend for shared resources, this paradigm reduces the overhead of arbitration of locks among the contending threads.

The “thread cooperation” model logically separates the requests (tasks to be performed by a thread) from the execution threads. This allows the system to be viewed as a set of threads processing a set of requests, and it does not matter which thread processes which request. This is analogous to how the operating system (OS) schedules jobs on the hardware execution threads. *Delegation* ensures that only the request is “blocked” and the associated thread can detach itself from the request and move forward to process the next request. This is analogous to the OS blocking I/O requests and not the actual CPU that was processing the job. Note that threads sharing resources do need synchronization, but the *waits* associated with synchronization primitives have been eliminated. In addition, as described later, cooperation based synchronization allows replacing locks with atomic operations which are much cheaper synchronization primitives.

## 4.2 Parallelizing Space Saving

Once we have formalized the *thread cooperation* paradigm, we now show how we can use this paradigm for parallelizing *Space Saving*. Figure 3 provides a high level overview of the system design for parallelizing *Space Saving* using the *thread cooperation* paradigm. The boundary conceptually separates the *Stream Summary* and the *Search Structure*, and these structures interact with each other through a well-defined interface, and need not be aware of the implementation details. The parallel stream processing system should guarantee that the following invariant holds:

INVARIANT 4.2. *Boundary Invariant*: If Thread  $T_i$  processing element  $\langle e \rangle$  has crossed the boundary into the *Stream Summary* structure, then it is the only thread active in the *Stream Summary* that is processing the element  $\langle e \rangle$ .

**Table 2: Basic atomic primitives.**

Operation	Description
CAS(addr, expected, new)	Atomically compare the <i>expected</i> value to the stored value at <i>addr</i> , and swap with <i>new</i> if successful
IAF(addr)	Atomically increment the location <i>addr</i> and read the incremented value
SWP(addr, new)	Atomically swap the present value at <i>addr</i> with the passed <i>new</i> value and return the old value

The *Search Structure* should guarantee that Invariant 4.2 holds and this provides the **element level synchronization** as discussed in Section 3.2. As long as the desired properties and the invariants hold, this framework is independent of the choice of the different structures involved and the actual algorithm used to process the stream elements. The “cooperation” based model does not guarantee any ordering of the requests, but this does not impact the correctness of the *Space Saving* algorithm.

LEMMA 4.1. *Order Invariance*: The final outcome of the *Space Saving* algorithm is invariant to the order in which the stream elements are processed.

This follows directly from the definition of the frequency counting problem. Given a stream of  $N$  elements, the *Space Saving* algorithm reports all the elements whose frequency is above  $\epsilon N$ , where  $\epsilon$  is the user specified error bound [20]. Therefore, as long as an element appears more than  $\epsilon N$  times in the input stream, *Space Saving* guarantees that the element will be reported, irrespective of where the element appeared in the stream.

COROLLARY 4.2. *Correctness of the Cooperation based parallel design*: Reordering of requests due to request delegation does not affect the correctness of *Space Saving*.

Corollary 4.2 follows directly from Lemma 4.1. In general, most stream processing algorithms are invariant to the order of processing, and the *thread cooperation* paradigm can be applied to them.

## 4.3 Implementation Details

As depicted in Figure 3, the system can be viewed as two different components which interact with each other through an interface that ensures that Invariant 4.2 holds. Recall that when processing the elements, a thread first accesses the *Search Structure* and this structure directs the thread to the appropriate location in the *Stream Summary*. In the implementation, an element in the hash table (*search structure*) points to the corresponding element in the *Stream Summary* structure, and this element in the *Stream Summary* in turn points to the bucket to which it belongs. Our implementation uses a number of atomic instructions. Refer to Table 2 for a brief description of the atomic operations used in the implementation. Only the CAS operation need to be supported by the underlying hardware, and the rest of the operations can be implemented using the CAS primitive. Note that the atomic operations supported by the hardware are much cheaper than acquiring locks. We now discuss the two different components in Figure 3 in detail, and provide sketches of correctness proof of the methods. Our approach for proving correctness is as follows: *The design is correct if all the invariants are satisfied. Therefore, for each individual operations, if all the invariants are satisfied as a precondition for the operation, the operation is said to be correct if it ensures that all the invariants are satisfied as a postcondition.*

### 4.3.1 Search Structure or Hash table

The hash table lookup is equivalent to the LOOKUP operation in Table 1. Even though there have been proposals for different lock-free hash table implementations (such as Shalev et al. [28]), the hash table used by the frequency counting algorithms has certain special requirements and characteristics. We design a simple concurrent hash table where separate chaining is used to resolve hash collisions, and bucket-level locks are used to synchronize insertions to the same chain. We modify the conventional design of a chained hash table [17] to suit the requirements specific to the application.

**Lookup and insertion:** The hash table is designed with minimal locking: the “readers” in the hash table do not need locks. Locks are needed only to serialize insertions to the same hash bucket. In this design, threads are blocked only when they are trying to insert into the same hash bucket. If a moderately robust hash function (such as *multiplicative hashing* [17]) is used, then the likelihood of two “writers” mapping to the same hash bucket is very rare and hence the associated wait is also small. In order for the “readers” to traverse the list “lock-free”, the following structural invariant must be satisfied at all times:

INVARIANT 4.3. *At all times during execution, a pointer should either be pointing to a valid node or be set to null.*

**Deletion:** Deletion of a hash table entry is done lazily – the entry to be deleted is atomically marked as deleted, and this implies logical removal. Once an entry is removed logically, it can either be physically removed by some other thread or be reused. In this design, the location of the entry is reused for future insertions into the chain, and the following invariant guarantees correctness. In the implementation, deletions are marked by setting the *value* field of the  $\langle key, value \rangle$  pair in the hash table to null.

INVARIANT 4.4. *Once a location has been marked as deleted, the key in the location has been removed from the hash table and any subsequent lookup for that key should return failure.*

INVARIANT 4.5. **Unique key:** *At no point should two valid locations in the hash table contain the same key.*

The invariants listed above correspond to basic properties of the hash table and any concurrent design must satisfy them. In addition, there are certain application specific requirements. To guarantee **element level synchronization**, the most common approach would be to have the hash table ensure this level of synchronization, because every thread processing a stream element first accesses the hash table. Therefore, if thread  $T_i$  has *exclusive access* to element  $e$ , then any thread  $T_j (j \neq i)$  processing  $e$  must wait. Since the elements (which form keys in the hash table) are a point of contention in the hash table, Principle 4.1 can be used to delegate the request and allow  $T_j$  to proceed with other requests while  $T_i$  guarantees to process the *delegated* requests as per Invariant 4.1. Since the requests in a frequency counting algorithm are increments to the count, “queueing” a delegation in the hash table is simple, and can be achieved through a field that stores a pending *count* corresponding to the element. If the pending *count* is 0, this implies that no thread owns the element and there are no pending requests. The first thread that increments the *count* acquires *exclusive access* on the element, and all other threads delegate the request by atomically incrementing the count. The following invariant guarantees **element level synchronization** and hence satisfies Invariant 4.2:

INVARIANT 4.6. *A thread obtains exclusive access to a key in the hash table if the IAF operation on count corresponding to the key returns 1.*




Figure 4: **Concurrent Stream Summary** structure based on the “thread cooperation” paradigm.

**Relinquishing an element:** Once a thread has processed a stream element, it has to relinquish *exclusive access* to the element. But before relinquishing access, it has to ensure that there are no delegated requests so that Invariant 4.1 is satisfied. To relinquish, the thread first performs a CAS of 1 (expected value) with 0 (new value) on the *count* field corresponding to the entry in the hash table. If this succeeds, there are no delegated requests, and the element has been relinquished. Any other thread performing a subsequent IAF for acquiring the element will read *count* as 1, and hence will be able to acquire the element (Invariant 4.6). Failure implies pending delegated requests and the thread will do a SWP with 1 which will return the number of delegated requests, and set *count* to 1. This ensures that the owning thread still has exclusive access to the element and prevents other threads from crossing the boundary. The thread then subtracts 1 from this value (which corresponds to the request for this thread in the earlier round) and crosses the boundary with the subtracted value as increment to the frequency of the element. This ensures that both Invariants 4.1 and 4.2 are satisfied. Note that as a result of delegated requests accumulating inside the hash table, increments to counts of elements are processed in bulk, and we refer to this as **bulk increment**. Due to space limitations, the pseudocode, details of the operations, and correctness proofs of the operations of hash table have been moved to [9].

### 4.3.2 Concurrent Stream Summary

In this section, we discuss the implementation of the *Concurrent Stream Summary* structure based on the proposed *thread cooperation* paradigm. The concurrent analogue of *Stream Summary* is very similar to the original *Stream Summary* structure in [20]. This modified structure is shown in Figure 4. The structure consists of a singly linked list of *frequency buckets*. Each bucket maintains a list of elements which have the same frequency as that represented by the bucket. The buckets are in ascending order of frequency (from left to right), and the elements traverse through this structure of buckets as their frequency changes while the stream is being processed. Each bucket has a *queue* of requests for holding requests delegated to the bucket, and we use a standard “lock-free” queue implementation [27]. As shown in Figure 4, each bucket in the linked list consist of different fields. The *next* field points to the next node in the list. The *inUse* field is used to guarantee **bucket level synchronization** and if set to 1, indicates that the bucket is already owned by a thread. When a thread wants to access a bucket, it performs a CAS of 0 (expected value) with 1 (new value) on the *inUse* field. Success implies that the bucket was available and now *inUse* is set to 1, and the thread performing the operation gets exclusive access to it. On failure, the thread appends the request to the bucket’s queue, and moves on to the next stream element. The *marked* field is used to represent logical presence, and if it is set to true, then the node is considered to have been *logically removed* from the list and marked for *garbage collection*.

**Optimistic Traversal:** Multiple “reader” threads might traverse

the list of frequency buckets, and the list has been designed to allow *optimistic traversal*. In other words, the readers traverse the list without acquiring locks, and abort the traversal if something goes wrong. To ensure correctness and safety, the following invariant should always be satisfied:

INVARIANT 4.7. *At no point of execution should a “reader” read invalid locations (garbage), i.e., all pointers in the list of frequency bucket should point to valid buckets or be set to null.*

**Maintenance of sort order:** The sort order of the buckets should be maintained at all times according to the following invariant:

INVARIANT 4.8. *If  $m$  and  $n$  are two buckets such that  $m.next = n$ , and  $m.marked = false \wedge n.marked = false$ , then:  $m.frequency < n.frequency$ .*

Since the readers traverse “lock-free”, if empty buckets are reused for different frequencies, it could lead to violation of the sort order. The following invariant prevents this:

INVARIANT 4.9. *Once a bucket is assigned a frequency, its frequency can never be changed.*

**Lazy Removal:** If a bucket does not contain any elements and there are no pending requests, its *marked* field is atomically set to **true** and is *logically removed* from the structure. Physical removal is done lazily. Although marked buckets are still present in the structure, any operation on these buckets is not-defined, so any “reader” accessing a marked bucket should abort the read and restart the process. To ensure correctness, the following invariant should hold:

INVARIANT 4.10. *A bucket with marked set to true cannot be reused. It can only eventually be removed from the structure. Any operation on a marked bucket is undefined and should fail.*

Once a bucket is logically removed, physical removal from the structure is done by a thread which acquires the bucket immediately preceding the *marked* bucket. If a sequence of buckets are marked, the entire sequence can be removed by the thread owning the bucket immediately preceding the sequence. This comprises the garbage collection process. During physical removal, only the next pointer of the bucket immediately preceding the sequence of marked buckets is changed. The thread owning the preceding bucket can therefore garbage collect the buckets without acquiring additional locks.

INVARIANT 4.11. *A marked bucket or a sequence of marked buckets can only be physically removed from the structure by a thread that is the owner of the bucket immediately preceding the sequence of buckets. If there are no preceding buckets, then the owner of the bucket can itself physically remove the bucket.*

As “readers” traverse lock free, once a sequence of buckets is removed, there might still be readers traversing these buckets. The structure of the sequence is therefore preserved so that the readers can eventually rejoin the main structure. A mechanism similar to reference counting in Java, or hazard pointers [22] can be used to reclaim a physically removed bucket without resulting in dangling pointers in threads that might still have a pointer to the reclaimed memory. We now explain the individual operations of *Stream Summary* which implements the operations in Table 1.

#### Implementation of AddElement and its Correctness

An *AddElement* request arrives at a bucket either when a new element is being added to the **Monitored set** (corresponding to *AddElement* in Table 1), or an existing element’s count is being incremented (during increment and overwrite). Owing to the *delegation*

---

#### Algorithm 4.1 Processing the AddElement Request

---

```

1: Procedure AddElement(element, bucket)
2: if (element.frequency = bucket.frequency) then
3:   bucket.AddToList(element)
4: else if (element.frequency < bucket.frequency) then
5:   if (bucket = minFreq) then
6:     newBucket := getNewBucket(element.frequency)
7:     newBucket.AddToList(element)
8:     newBucket.next := bucket
9:     minFreq := newBucket
10:  else
11:    DelegateRequestToBucket(minFreq)
12:  else
13:    FindDestBucket(bucket, element)
14: end Procedure AddElement

```

---

of the requests, the *Add* request to a bucket can be of three types: (i) an element with the same frequency as that of the bucket, (ii) an element with frequency greater than that of the bucket (*delegation* of the add request for a not existent bucket immediately following the current bucket), and (iii) an element with frequency less than the present bucket (add request of new element with frequency 1 being added to the **Monitored set** and *delegated* to a minimum frequency bucket whose frequency is greater than 1). Algorithm 4.1 provides a high level overview of how the add request is processed, and this operation subsumes the *AddElement* operation in Table 1. If the addition is to the same bucket, it can be processed right away, while addition to a higher frequency bucket is handled by the *FindDestBucket*. For inserting elements with frequency less than the current minimum frequency, the request is *delegated* to the minimum frequency bucket.

To argue the correctness of Algorithm 4.1, we must show that it does not violate any of the structural invariants. In line 6, when the new node is created, the creator is the only thread with access to the bucket and hence is free to perform any operation on the bucket<sup>2</sup>. The new bucket is first made to point to the old minimum frequency (line 8) and then it is made *public*<sup>3</sup> by the assignment in line 9. Thus Invariant 4.7 is satisfied at all times. Again, the conditional at line 5 ensures that Invariant 4.8 is maintained. Therefore, if *FindDestBucket* does not violate any invariants, then all other invariants are satisfied because this function does not affect those properties. Therefore, this operation is correct.

#### Implementation of FindDestBucket and its Correctness

Finding the next bucket when incrementing the count of an element is an important operation for both *AddElement* and *IncrementCounter*. This function finds the appropriate bucket for inserting an element into the structure. If the current next node is not the destination for the element, then either a new bucket needs to be inserted after the present bucket, or the list of buckets needs to be traversed. List traversal is necessary to handle **bulk increments** resulting from delegations in the hash table (Section 4.3.1). In either case, the design of the structure allows optimistic traversal. If at any point during traversal through the list, the “reader” finds that it is accessing a node that has been marked for garbage collection, then it aborts the present run, and starts the traversal again. From an efficiency perspective, this failure and abort will be rare, as this case would arise when dealing with **bulk increments**, and this is common only for the high frequency elements, which are generally towards the extreme right of the **Concurrent Stream Summary** structure. For the less frequent elements in the middle of the structure, the next node will be the destination in most cases.

<sup>2</sup>Such a node is called *private* and only the creator can access it.

<sup>3</sup>Making *public* implies that other threads can access the bucket and is now treated as a shared resource.

---

**Algorithm 4.2** Processing the *IncrementCounter* Request

---

```
1: Procedure IncrementCounter(element, increment)
2: bucket.RemoveFromList(element)
3: element.frequency += increment
4: FindDestBucket(bucket, element)
5: if (bucket.isEmpty() && bucket.noPendingRequests()) then
6:   gcStatus := bucket.atomicMarkGarbageCollected()
7:   if (gcStatus = true && bucket = minFreq) then
8:     newMinFreq := findNewMinFreqBucket(minFreq)
9:     minFreq := newMinFreq
10: end Procedure IncrementCounter
```

---

The pseudocode and the correctness proof of this operation can be found in [9].

**Implementation of IncrementCounter and its Correctness**

The *IncrementCounter* request (Table 1) increments the count of an element while maintaining the sort order of the elements, and is one of the most frequent operations for the application under consideration. Algorithm 4.2 provides an overview of how the operation is implemented. After an element has been incremented (i.e., removed from the original bucket), if the bucket becomes empty and has no pending requests, then the bucket should be atomically marked for garbage collection.

As above, correctness of the operation is proved by showing it does not violate any of the structural invariants. Each of the constituent operations satisfy the invariants (as argued in the previous sections). If after the increment, the bucket becomes empty, and there are no pending requests, then the thread will atomically mark for garbage collection (line 6) thus satisfying Invariant 4.9. If the minimum frequency bucket is marked for garbage collection, since there are no other buckets preceding this bucket the current thread removes the bucket (line 9) thus satisfying Invariant 4.11. The *findNewMinFreqBucket* is similar to the *FindDestBucket* operation and returns the first valid bucket immediately after the current *minFreq* bucket, and in the process removes any marked nodes if present. Since this find minimum operation is similar to the *FindDestBucket* operation, its correctness is similar to the correctness proof of *FindDestBucket*. Other invariants are satisfied because this operation does not alter these properties.

**Implementation of Overwrite and its Correctness**

*Overwrite* is specific to the *Space Saving* algorithm which uses this operation to limit the number of counters monitored. *Overwrite* implements *Space Saving* algorithm’s heuristic of replacing a minimum frequency element with a new un-monitored incoming element. It selects a candidate element from the minimum frequency bucket, overwrites it with the current element being processed, and then increments the count of the new element. To select a candidate for overwriting, the thread will follow the principle of *Minimal Existence* (Principle 4.2) and will not block on any shared resource. It will start from the first element in the minimum frequency bucket, and to overwrite the element, the corresponding entry in the *Search Structure* should be deleted. This deletion is non-blocking, and failure implies that some other thread is trying to increment that element, and the increment request should be in the request queue of the minimum frequency bucket. The thread then moves on to the next element in the bucket, and the process is repeated. If all elements are busy and none of them could be overwritten, this implies that all these elements have pending increment requests. The overwrite request is thus *deferred* till all the increment requests have been processed. Since a single thread will process all these requests, the processing can be highly optimized as the thread now has more knowledge about the requests. The pseudocode for the operation and its correctness proof can be found in [9].

### 4.3.3 Auto Tuning and Throttling

The crux of the proposed “thread cooperation” model is the use of request *delegation* to remove “waits”. When a request is delegated inside the hash table (*Search Structure*), the counts accumulate, and all delegated requests are processed in bulk. But when a request is delegated in the *concurrent stream summary* structure, each delegated request adds to the queue of requests, and the thread currently owning the bucket has to process all the requests to satisfy Invariant 4.1. Thus, if many requests are delegated to the same bucket, this will lead to an excessive increase in the length of the request queue, and in turn may lead to performance degradation. This happens predominantly for uniform streams where most of the elements are infrequent, and therefore a large number of elements are deleted from the **Monitored Set**. Since all deletions happen at the minimum frequency bucket, this bucket becomes a “hot-spot”. In such a scenario, parallelism is limited by the structure, and additional threads are simply adding to the overhead of the system.

To avoid performance degradation in such a scenario, we propose a mechanism of *auto tuning and throttling* of threads. Whenever a thread delegates a request in the *stream summary*, it monitors the queue size of the bucket to which the request is delegated. If the queue size is above a threshold  $\rho$ , then this bucket is a potential “hot-spot” and continuous additions to its request queue might lead to performance degradation. Therefore, this thread will “throttle” itself and sleep for a small time slice. This can be treated as another form of “thread cooperation” where a thread abstains from overloading an already overloaded bucket. This feedback loop dramatically improves the performance when dealing with uniform streams, but also introduces some overhead due to the necessity of monitoring the queue sizes. In Section 5 we evaluate the benefits as well as overhead resulting from *auto tuning and throttling*.

### 4.3.4 Space Analysis

The parallel implementation of *Space Saving* as described earlier uses a single shared *concurrent stream summary*. But additional space is required to “queue” requests that have been delegated. If an element has been queued, at any instant, there is only one request corresponding to that element. This is equivalent to the element being in the input stream, which would anyway have to be allocated space to be buffered. Thus instead of the element being present in the input buffer, it is encapsulated as a request in the queue. Additional space is required for locks and synchronization, which is again of the same order as the number of counters. Thus, the space complexity for the parallel algorithm is of the same order as that of the sequential algorithm, i.e.,  $O(\frac{1}{\epsilon})$ .

### 4.3.5 Answering Queries

The *Concurrent Stream Summary* structure maintains the elements in a sorted order, so that queries can efficiently traverse this structure to find the appropriate elements. Queries can be processed with high efficiency because the elements of interest will have high frequencies and reside in the rightmost end of the structure, and the low frequency elements will be cluttered in the leftmost end (assuming frequencies increase from left to right). Therefore, as the queries start from the minimum frequency, they can very quickly prune out the low frequency elements and reach the region of interest. Again, queries can be answered without acquiring any locks. Details of how different types of frequent elements and top- $k$  queries are answered can be found in [9].

## 4.4 Generalization of the Framework

In this paper, we explain the application of the proposed framework for parallelizing the *Space Saving* algorithm. One of the rea-

sons for selecting *Space Saving* was that a recent work [5] compared a number of frequency counting algorithms and demonstrated that *Space Saving* outperforms similar algorithms in terms of processing speed and accuracy (expressed as precision and recall). Another reason is that *Space Saving* has a constant space bound of  $O(\frac{1}{\epsilon})$  and this simplifies the algorithm implementation.

The *thread cooperation* based framework is not limited to the *Space Saving* algorithm and can be easily extended to other frequency counting algorithms. It is evident from the description of the *IncrementCounter* operation that the framework can handle arbitrary increments in frequency of the elements. Thus, any frequency counting algorithm in which the frequency of the counters *increases monotonically* (algorithms such as *Lossy Counting* [19]), can be adapted to the framework. For adapting other algorithms, most of the operations (except the *Overwrite* operation) will remain unchanged. For example, the *Lossy Counting* [19] algorithm divides the stream into multiple rounds and at the end of the round removes elements which are infrequent. Therefore, for adaptation into the framework, only the *Overwrite* request in *Space Saving* has to be replaced by a request that removes the minimum frequency bucket at round boundaries.

## 5. EXPERIMENTAL EVALUATION

In this section, we experimentally evaluate the proposed framework. The experiments have been performed on an Intel Core 2 Quad Q6600 processor [16]. This processor has 4 cores, each corresponding to a hardware thread and operating at a clock speed of 2.4GHz, and the cores share a L2 Cache of 4MB. The machine has 3.2 GB available main memory and runs Fedora Core Linux with kernel 2.6.26.6-49.fc8. All algorithms and the framework have been implemented in C++ and compiled using GNU C++ compiler with Level 2 optimization. The data set is synthetically generated and follows zipfian distribution which is very close to realistic data distribution [30]. The zipfian factor  $\alpha$  determines whether the distribution is uniform or skewed. The frequency of the elements in the distribution varies as  $f_i = \frac{N}{i^\alpha \zeta(\alpha)}$  where  $\zeta(\alpha) = \sum_{i=1}^{|A|} \frac{1}{i^\alpha}$  where  $N$  is the length of the stream,  $|A|$  is the size of the alphabet, and  $f_i$  represents the frequency of the  $i^{th}$  frequent element. Smaller values of  $\alpha$  represent lesser skew in the distribution with  $\alpha = 0$  representing uniform distribution. As the value of  $\alpha$  increases, the skew of the data distribution also increases. The data set has a total of 50 million elements and an alphabet of 5 million. GCC built-in atomic primitives were used for performing the atomic operations, and *pthread* library was used for threads and locks. In our first set of experiments, we choose data with  $\alpha$  in the range 1.5 to 3.0. The lower  $\alpha$  values have not been evaluated because the frequent elements and top- $k$  elements are more interesting and meaningful in a skewed distribution, than in a uniform distribution. In the case of a uniform data (such as  $\alpha \leq 1.0$ ) there are a lot of infrequent elements, and so in *Space Saving*, there will be a lot of *Overwrite* requests, thereby making the minimum frequency bucket a “hot spot”. Thus, such a distribution will not benefit from parallelization. For such a scenario, we propose *auto tuning and throttling* (Section 4.3.3), and the effectiveness of the scheme is demonstrated by the experiments in the latter part of this section for data sets with  $\alpha \leq 1$ . In all our experiments, the number of streams is varied from 4 up to 32 in multiples of 2, and each stream has 1 million elements, i.e., in an experiment with 32 streams, the total number of elements processed is 32 million. The performance of algorithms does not vary significantly for smaller values of error bound  $\epsilon$ , while larger values of  $\epsilon$  result in performance degradation. Larger  $\epsilon$  also results in smaller number of buckets, and hence




Figure 5: Execution time for saturated workload.




Figure 6: Speedup of the “cooperation” based design compared to the “contention” based design for saturated workload.

lesser chance for exploiting parallelism. In our experiments, we set  $\epsilon = 0.0001$ . In this section, for a fair evaluation, we compare the *cooperative* design to the *shared* design. The *independent* design is not included in the comparison since its performance is heavily dependent on the query rate and  $\epsilon$ , whereas the remaining two are more stable.

### 5.1 Saturated Workload

In this experiment, we evaluate the maximum throughput of the proposed design by using a *saturated workload* wherein a thread always has some element to process whenever it has finished processing the earlier element, or in other words, a thread is *saturated* with work. Figure 5 provides a comparison of the execution times of the “cooperation” based design (referred to as **Coop** in this graph as well as the rest of the graphs in this section) and the “contention” based design (referred to as **Shared**). Along the  $x$ -axis, we plot the number of streams (and hence the number of threads), and along the  $y$ -axis we plot the execution time in seconds. Different lines in the graph correspond to different values of  $\alpha$ . An ideal system with linear speed-up will have a horizontal line, and the smaller the slope of the graphs, the better it is. As is evident from Figure 5, not only does **Coop** outperform **Shared** in terms of execution time, **Coop** demonstrates better scalability since the slope of the lines corresponding to **Coop** is considerably smaller than that of **Shared**. The benefits of **Coop** with respect to execution time can be seen from Figure 6 which plots the speedup of **Coop** compared to **Shared**. In this figure, the  $x$ -axis represents the varying zipfian  $\alpha$  and the  $y$ -axis corresponds to the speedup. For a particular zipfian  $\alpha$ , the speedup is computed as:  $T_{(i,\alpha)}^{Shared} / T_{(i,\alpha)}^{Coop}$ , where  $i$  represents the number of streams (4, 8, etc.). As is evident from Figure 6, **Coop** outperforms **Shared** by a factor ranging from 2 – 5.5X. This experiment demonstrates the superiority of the performance of the proposed “cooperation” based design.

### 5.2 Unsaturated Workload

In many realistic applications, the workload might not be saturated, i.e., due to variable arrival rates, when a thread is ready to process the next stream element, there is a likelihood that the next element has not yet arrived. In such a case, we say that the read “failed” and the thread needs to wait for the arrival of the next




Figure 7: Execution time for an unsaturated workload.

element when the read “succeeds”. We model this behavior by allowing a read to “fail” with certain probability. In other words, a failure probability of 10% means that with every read, there is a 10% chance that the read will fail. Thus the saturated workload considered in the previous subsection has a failure probability of 0%. That experiment would demonstrate the overhead on the systems at higher arrival rates. When a read “fails”, the thread does not access the **Monitored Set**. This might have two implications: *i*) if the structure is already overloaded, this will lead to a reduction in the load and can be treated as “load-shedding” and this might result in reduced execution times, *ii*) if the structure is lightly loaded, it has the potential to accommodate more concurrency, and hence a failed read implies failed opportunity and this will lead to increased execution time. Thus, for a system that is performing at its peak for a saturated workload, an unsaturated workload will result in an increase in total time of completion due to slower arrival rate. But for a system which is overloaded for a saturated workload, an unsaturated workload might lead to lesser execution time.

Figure 7 plots the execution time of the two designs as the failure probability is varied. Along the  $x$ -axis we plot the number of streams and along the  $y$ -axis, we plot the execution time in seconds. The different lines correspond to different values of zipfian  $\alpha$  and the different designs. The zipfian  $\alpha$  of the input distribution is varied from 1.5 to 3.0 and the experiment is repeated for different probabilities of not finding a stream element when a thread is ready. It is evident from Figure 7 that the behavior of the designs is almost similar to that with the saturated case (Figure 5) except that with increasing failure probability, the graphs for **Coop** and **Shared** are getting closer. This is due to the fact that the low contention overhead in **Coop** allows more concurrency compared to that of **Shared**. In **Shared**, the contention overhead is very high (Figure 2), and if a thread “fails” a read and does not access the structure, it contributes to reducing the contention overhead on the structure, and hence the overall execution time does not increase much. These experiments further strengthen the claim that the “cooperation” based design has low contention overhead and allows higher concurrency, as a result of which decreased packet arrival rate increases execution time which is bound to happen in a system operating at its peak. Refer to our technical report [9] for detailed analysis, and for graphs of speedup corresponding to the execution times reported in this subsection.

### 5.3 Evaluation of Auto Tuning

In Section 4.3.3, we discussed the potential for using *auto tuning and throttling* to deal with “hot-spots” when processing uniform data distributions. In all the previous experiments, we considered distributions for zipfian  $\alpha \geq 1.5$ , and in this section, we evaluate the impact of *auto tuning* for  $\alpha \leq 1.0$ . Due to too many overwrites of minimum frequency element in the case of uniform distributions, the minimum frequency bucket becomes a “hot-spot”. Profiling experiments revealed that without *auto-tuning*, for  $\alpha = 0.5$ , the max-




Figure 8: Comparison of execution times of *shared* and *cooperation* based design with *auto tuning* for saturated workload.

imum sampled queue size grows to about 3 – 4 orders magnitude higher than in the case with  $\alpha \geq 1.5$ . This results in severe performance degradation since a majority of the work is delegated to a single thread, thereby reducing the system to a sequential system with all the overhead of a concurrent system. *Auto tuning and throttling* prevents queues from building up rapidly, and allows better distribution of the work. Our experimental samples using  $\alpha \geq 1.5$  showed maximum queue size of the order of hundreds and average queue size of the order of tens. So, in our experiments, we set the queue size threshold  $\rho$  to 500. Therefore, for skewed data, the threads will hardly be throttled and can perform at their peak rate, while for almost uniform data, threads will be repeatedly throttled incorporating admission control and better load distribution.

In Figure 8, we plot the execution time of the Cooperation based design using *auto tuning* and compare it with **Shared**. Along the  $x$ -axis, we plot the number of streams and along the  $y$ -axis, we plot the execution time in seconds. Different lines correspond to different values of  $\alpha$ , and **Coop AT** represents “cooperation” based design with *auto tuning*. In this experiment, we use a saturated workload. To improve clarity of the graph, we only provide the lines corresponding to  $\alpha = 0.5$  and 1 (our new data sets) and  $\alpha = 1.5$  and 2.5 (representatives from the previously used data sets). As is evident from the figure, **Coop AT** performs reasonably for lower values of  $\alpha$  (almost uniform) where **Coop AT** outperforms **Shared** by a factor of  $\sim 2X$ . For higher values (skewed), its performance is almost similar to that of the “cooperation” based design without *auto tuning*.

*Auto tuning* introduces the overhead of profiling of queue sizes, and when the threads are not throttled (in the skewed case), this is pure overhead compared to **Coop**. But our experiments reveal that this overhead is limited to  $\sim 5\%$  (for example, for  $\alpha = 2.0$  and number of streams as 16, the average execution time for **Coop** is 4.195586s and that of **Coop AT** is 4.299514 resulting in an overhead of  $\sim 3.29\%$ ). Thus for a very small overhead, we have an algorithm that performs well for the entire spectrum of  $\alpha$ . If the system designer knows that the stream to be processed has a high skew, then *auto tuning* can be turned off allowing higher performance, but turning on *auto tuning* only marginally degrades performance while allowing greater flexibility.

### 5.4 Profiling the Cooperation based design

In this section, we provide an in depth analysis of the proposed “cooperation” based parallel design. Figure 9 plots the various profiling results of the design without using *auto tuning* while Figure 10 plots the corresponding results of the design using *auto tuning*. Figures 9(a) and 10(a) plot a breakup of the time spent by the design on various important operations expressed as a percentage of the total time of execution. The  $x$ -axis plots the number of streams, and the  $y$ -axis plots the time spent on operations as a percentage of total time. *Hash* refers to the time spent for the hash table operations, *Stream Summary* refers to the time spent for the *Stream Sum-*

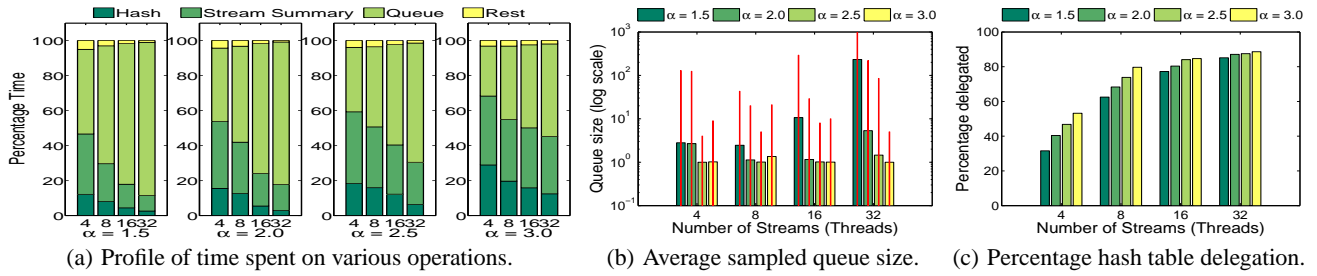


Figure 9: Detailed profile of cooperation based design without auto tuning while processing a saturated workload.

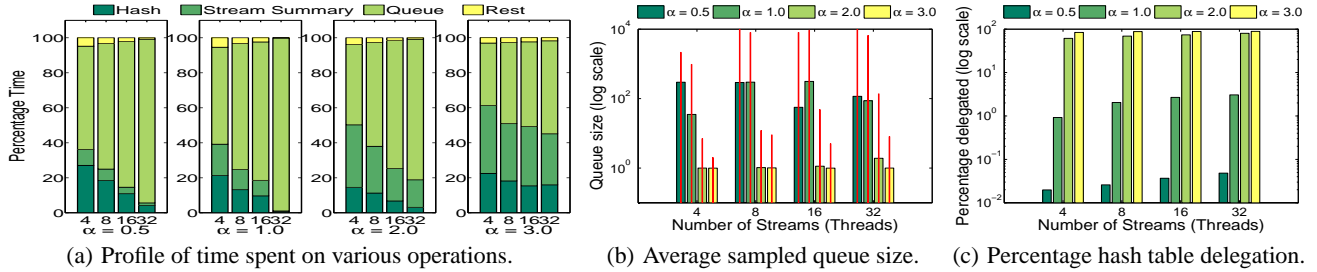


Figure 10: Detailed profile of cooperation based design with auto tuning while processing a saturated workload.

mary operations, *Queue* refers to the sum of time spent in appending the delegated requests to the queues associated with buckets in *Stream Summary* and removing the requests from the queue when processing them. As is evident from Figures 9(a) and 10(a), very little time is spent in the hash table for **element level synchronization** and the *Stream Summary* for **bucket level synchronization** (in contrast to Figure 2). A considerable portion of the time is spent on the operations related to the request queue, and a more scalable queue implementation would thus result in further improvement in performance of the proposed design. Our present queue implementation does not allocate memory wisely, and this is in part responsible for the high overhead of the queue operations. Note that the time for the queue operations is higher for smaller values of zipfian  $\alpha$ . This is because for more uniform distributions, different threads are generally processing different elements, and hence would not be blocked at the hash table, therefore resulting in large number of accesses to the *Stream Summary*. Since the summary structure has limited parallelism due to **bucket level synchronization**, a lot of requests are delegated in this structure. The increase in number of concurrent accesses to the request queues result in the increase in time for queue operations.

The above analysis is further supported by Figures 9(b), 10(b) and Figures 9(c), 10(c) which plot the average sampled queue size and percentage of requests delegated in the hash table respectively. Figures 9(b) and 10(b) plot the number of streams (or threads) along the  $x$ -axis, and the sampled queue size in log scale along the  $y$ -axis. The bars in the graph plot the average of the sampled queue sizes, while the thin lines overlaid on the bars plot the corresponding maximum sampled queue size. As discussed earlier, for smaller values of zipfian  $\alpha$ , the minimum frequency bucket becomes a “hot spot”, and this is reflected by the considerably larger queue sizes for  $\alpha = 1.5$  in Figure 9(b). The queue sizes for  $\alpha = 0.5$  and 1.0 in Figure 10(b) demonstrate the success of *auto tuning* in limiting the queue sizes even for uniform distributions, and explains the improved performance of **Coop AT** as shown in Figure 8. Finally, Figures 9(c) and 10(c) show the successful removal of *waits* associated with **element level synchronization** by plotting the requests delegated in the hash table as a percentage of the total number of re-

quests. The number of streams is plotted along the  $x$ -axis, and the percentage of delegated requests is plotted along the  $y$ -axis. For higher  $\alpha$  (skewed streams) and increasing number of threads, more threads are likely to be processing the same elements (since in a skewed distribution, few elements have dominant presence in the stream). Therefore, more requests are delegated in the hash table due to **element level synchronization**. The number of hash delegations increase with the number of threads and the increase in  $\alpha$ . Note that in spite of the large number of requests delegated in the hash table, neither is the hash table a point of contention, nor is a significant portion of execution time spent synchronization in the hash table (Figures 9(a) and 10(a)). This demonstrates the efficiency of the proposed *cooperation* based design in reducing the *waits* associated with synchronization. Due to space limitations, additional profiling experiments have been moved to [9].

## 6. CONCLUDING REMARKS

In this paper, we consider the problem of analyzing multiple data streams and propose a parallel algorithm for the problem in the context of multicore processors. The proposed design uses the concept of *thread cooperation* to remove *waits* associated with locks. Removal of the *waits* is particularly significant in the context of multicore processors where a waiting thread results in wasted CPU cycles. Additionally, the proposed model allows synchronization using only atomic operations which are much cheaper than locks required for synchronizing in the “contention” based design. The proposed design conceptually segregates the requests from the execution threads, and whenever a shared resource required by the request is not available, only the request is blocked (*delegated*), and the thread can move on to process the next request. This is extremely important in the context of parallelism in multicore processors, and the benefits of the proposed design is evident from the experimental results which show that the proposed “cooperation” based design outperforms the traditional “contention” based design by a factor of 2 – 5.5X over the entire spectrum of uniform to skewed data sets. In spite of *multi-stream* analysis systems being extremely hard to parallelize, our parallel design effectively uses the inherent parallelism in multicores with significantly lower

overhead. In addition to the gains for comparatively skewed data, our design is also efficient for uniform data sets where achieving parallelism is a much harder problem. In this paper, we implement the *Space Saving* algorithm, but the framework is general enough to accommodate other frequency counting algorithms [19].

During the course of implementation of this system, we gained useful experience and learned several important lessons which apply to parallel systems in general. *First*, even though individual lock operations are not very expensive, acquiring and releasing millions of locks per second results in significant overhead, which is in part responsible for the inefficiency of the *shared design*. *Second*, a lot of operations use locks only for mutual exclusion, and do not require the serialization order imposed by locks. For such applications, mutual exclusion can be achieved by using cheaper atomic operations, while *delegation* allows removal of waits. Exploiting these application specific characteristics might result in significant improvement in performance. *Third*, frequent system calls are also expensive, and introduce significant overhead. Memory allocation is one such expensive call, and frequent memory allocation can considerably deteriorate performance. As a result, spending time with a custom memory allocator is often useful.

In the last few years, there has been growing interest in different non-conventional parallel processing architectures such as Graphics processors [11, 14], cell broadband engine [10, 15], etc. These processors were originally designed for different application domains, but recent results have shown promise in the use of these processors for data management operations. The present *cooperation based* design does not differentiate the processing capabilities of the threads, and hence is suitable for *homogeneous* conventional multicore architectures such as chip multiprocessors (CMP) [16, 25]. Heterogeneous processing capability in combination with specialized instructions, such as vector operations, supported by Graphics and Cell processors open new challenges. In the future, we would like to explore the possibility of extension of the proposed paradigm to these other novel architectures. Additionally, experiments in this paper concentrate on a specific type of CMP architecture [16] known as the *fat camp* processors [12]. The *fat camp* processors are characterized by fewer cores where each core has been optimized for extremely efficient single thread performance. Another class of CMPs, referred to as *lean camp* processors [12], are characterized by large number of cores and hardware threads [25], but each hardware thread is not optimized for efficient single thread performance. In the future, we would also like to evaluate the performance of the proposed design on these *lean camp* processors.

## Acknowledgements

We would like to thank the anonymous reviewers, our shepherd Qiong Luo, and Pamela Bhattacharya for their insightful comments that has helped improve this paper. This work is partly supported by NSF Grant IIS-0744539.

## 7. REFERENCES

- [1] D. J. Abadi, Y. Ahmad, M. Balazinska, U. Çetintemel, M. Cherniack, J.-H. Hwang, W. Lindner, A. Maskey, A. Rasin, E. Ryzkina, N. Tatbul, Y. Xing, and S. B. Zdonik. The Design of the Borealis Stream Processing Engine. In *CIDR*, pages 277–289, 2005.
- [2] B. Babcock and C. Olston. Distributed top-k monitoring. In *SIGMOD*, pages 28–39, 2003.
- [3] M. Balazinska, H. Balakrishnan, and M. Stonebraker. Load Management and High Availability in the Medusa Distributed Stream Processing System. In *SIGMOD*, pages 929–930, 2004.
- [4] A. Bulut and A. K. Singh. A unified framework for monitoring data streams in real time. In *ICDE*, pages 44–55, 2005.
- [5] G. Cormode and M. Hadjieleftheriou. Finding frequent items in data streams. *Proc. VLDB Endow.*, 1(2):1530–1541, 2008.
- [6] C. Cranor, T. Johnson, O. Spataschek, and V. Shkapenyuk. Gigascope: A Stream Database for Network Applications. In *SIGMOD*, pages 647–651, 2003.
- [7] G. Das, D. Gunopulos, N. Koudas, and N. Sarkas. Ad-hoc top-k query answering for data streams. In *VLDB*, pages 183–194, 2007.
- [8] S. Das, S. Antony, D. Agrawal, and A. El Abbadi. CoTS: A Scalable Framework for Parallelizing Frequency Counting over Data Streams. In *ICDE*, pages 1323–1326, 2009.
- [9] S. Das, S. Antony, D. Agrawal, and A. El Abbadi. Thread Cooperation in Multicore Architectures for Frequency Counting Over Multiple Data Streams. Technical Report 2009-04, CS, UCSB, 2009. [http://www.cs.ucsb.edu/research/tech\\_reports/](http://www.cs.ucsb.edu/research/tech_reports/).
- [10] B. Gedik, P. S. Yu, and R. Bordawekar. Executing Stream Joins on the Cell Processor. In *VLDB*, pages 363–374, 2007.
- [11] N. K. Govindaraju, N. Raghuvanshi, and D. Manocha. Fast and approximate stream mining of quantiles and frequencies using graphics processors. In *SIGMOD*, pages 611–622, 2005.
- [12] N. Hardavellas, I. Pandis, R. Johnson, N. G. Mancheril, A. Ailamaki, and B. Falsafi. Database Servers on Chip Multiprocessors: Limitations and Opportunities. In *CIDR*, pages 339–350, 2007.
- [13] S. Harizopoulos, D. J. Abadi, S. Madden, and M. Stonebraker. OLTP through the looking glass, and what we found there. In *SIGMOD*, pages 981–992, 2008.
- [14] B. He, K. Yang, R. Fang, M. Lu, N. Govindaraju, Q. Luo, and P. Sander. Relational joins on graphics processors. In *SIGMOD*, pages 511–524, 2008.
- [15] S. Héman, N. Nes, M. Zukowski, and P. Boncz. Vectorized data processing on the cell broadband engine. In *DaMoN*, pages 1–6, 2007.
- [16] Intel® Core™ 2 Quad Processor Overview. <http://www.intel.com/products/processor/core2quad/index.htm>; (Accessed 06/08/2009), 2009.
- [17] D. E. Knuth. *The Art of Computer Programming, Sorting and Searching*, volume 3. Addison-Wesley, Cambridge, MA, 1997.
- [18] A. Manjhi, V. Shkapenyuk, K. Dhamdhere, and C. Olston. Finding (Recently) Frequent Items in Distributed Data Streams. In *ICDE*, pages 767–778, 2005.
- [19] G. S. Manku and R. Motwani. Approximate frequency counts over data streams. In *VLDB*, pages 346–357, 2002.
- [20] A. Metwally, D. Agrawal, and A. El Abbadi. An integrated efficient solution for computing frequent and top-k elements in data streams. *ACM Trans. Database Syst.*, 31(3):1095–1133, 2006.
- [21] A. Metwally, D. Agrawal, A. El Abbadi, and Q. Zheng. On hit inflation techniques and detection in streams of web advertising networks. In *ICDCS*, page 52, 2007.
- [22] M. M. Michael. Safe memory reclamation for dynamic lock-free objects using atomic reads and writes. In *PODC*, pages 21–30, 2002.
- [23] R. Motwani, J. Widom, A. Arasu, B. Babcock, S. Babu, M. Datar, G. Manku, C. Olston, J. Rosenstein, and R. Varma. Query Processing, Resource Management, and Approximation in a Data Stream Management System. In *CIDR*, 2003.
- [24] K. Mouratidis, S. Bakiras, and D. Papadias. Continuous monitoring of top-k queries over sliding windows. In *SIGMOD*, pages 635–646, 2006.
- [25] Sun UltraSPARC T2 Processor – The industry’s first true system on chip. <http://www.sun.com/processors/UltraSPARC-T2/index.xml>; (Accessed 06/08/2009), 2007.
- [26] P. E. O’Neil. The escrow transactional method. *ACM Trans. Database Syst.*, 11(4):405–430, 1986.
- [27] W. N. Scherer III, D. Lea, and M. L. Scott. Scalable synchronous queues. In *PPoPP*, pages 147–156, 2006.
- [28] O. Shalev and N. Shavit. Split-ordered lists: Lock-free extensible hash tables. *J. ACM*, 53(3):379–405, 2006.
- [29] S. Venkataraman, D. Song, P. Gibbons, and A. Blum. New streaming algorithms for fast detection of superspreaders. In *NDSS*, pages 149–166, 2005.
- [30] G. K. Zipf. *Human Behavior and The Principle of Least Effort*. Addison-Wesley, Cambridge, MA, 1949.