# Efficient data streaming multiway aggregation through concurrent algorithmic designs and new abstract data types\*

Vincenzo Gulisano, Yiannis Nikolakopoulos, Daniel Cederman Marina Papatriantafilou and Philippas Tsigas

{vinmas, ioaniko, cederman, ptrianta, tsigas}@chalmers.se

#### Abstract

Data streaming relies on continuous queries to process unbounded streams of data in a real-time fashion. It is commonly demanding in computation capacity, given that the relevant applications involve very large volumes of data. Data structures act as articulation points and maintain the state of data streaming operators, potentially supporting high parallelism and balancing the work between them. Prompted by this fact, in this work we study and analyze parallelization needs of these articulation points, focusing on the problem of streaming multiway aggregation, where large data volumes are received from multiple input streams. The analysis of the parallelization needs, as well as of the use and limitations of existing aggregate designs and their data structures, leads us to identify needs for proper shared objects that can achieve low-latency and high-throughput multiway aggregation. We present the requirements of such objects as abstract data types and we provide efficient lock-free linearizable algorithmic implementations of them, along with new multiway aggregate algorithmic designs that leverage them, supporting both deterministic order-sensitive and order-insensitive aggregate functions. Furthermore, we point out future directions that open through these contributions. The paper includes an extensive experimental study, based on a variety of aggregation continuous queries on two large datasets extracted from SoundCloud, a music social network, and from a Smart Grid network. In all the experiments, the proposed data structures and the enhanced aggregate operators improved the processing performance significantly, up to one order of magnitude, in terms of both throughput and latency, over the commonly-used techniques based on queues.

<sup>\*</sup>A brief announcement about parts of this work has been accepted at the 26th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA), 2014 [4]. The research leading to these results has been partially supported by the European Union Seventh Framework Programme (FP7/2007-2013) through the EXCESS Project (www.excess-project.eu) under grant agreement 611183, through the SysSec Project, under grant agreement 257007, through the FP7-SEC-285477-CRISALIS project, by the collaboration framework of Chalmers Energy Area of Advance and by the Chalmers Center for E-science.

# **1** Introduction

For data intensive computing that can support continuous complex analysis of large volumes data, the data streaming processing paradigm emerged as a more appropriate alternative to the traditional "store-then-process" one. As emphasized in [7], the low-latency and high-throughput requirements of such continuous real-time complex processing of increasingly large data volumes make parallelism a necessity.

In data streaming [28, 1, 6], *continuous queries*, defined as Directed Acyclic Graphs (DAGs) of interconnected operators, are executed by Stream Processing Engines (SPEs) that process incoming data in a real-time fashion, producing results on an on-going basis. A good portion of the research has so far focused on leveraging the processing capacity of clusters of nodes and originally centralized SPEs [1] evolved rapidly to distributed [2] and parallel [10, 16] ones. At the same time, research has also focused on leveraging multi-core CPUs and GPUs architectures, as discussed in [5, 23, 24]

A parallel data streaming application can be seen as a pipeline where data is continuously produced, processed and consumed. In a parallel environment the underlying data structures should provide the means for organizing the data so that the communication and the work imbalance between the concurrent threads performing the computation are minimized while the pipeline parallelism is maximized. Finding the appropriate data structures that fit the needs of an application in a concurrent environment is a key research issue [25, 19]. Defining and providing the data structures that meet the needs of concurrent data streaming applications is a rich issue not addressed in the literature; this is all the more important, given the high performance demands of the relevant applications.

By shedding light on the data structures, we identify new key challenges to improve data streaming aggregation, one of the most common and throughputdemanding monitoring application [9]. In particular, we focus on multiway aggregation, where big volumes of data received from multiple input streams must be merged and sorted in order to be processed deterministically [10]. Sample application scenarios include monitoring applications in the context of social media, where information could be aggregated to study trends, or in the context of real-time pricing applications in Smart Grids, or in the context of adaptive traffic systems.

**Contributions** We study data structures as articulation points between pipeline stages of streaming aggregation. The shared access to the data by the collaborating threads defines new synchronization needs that can be integrated in the functionality provided by the shared data structures. By studying the use and limitations of existing aggregate designs and the data structures they use, we motivate the need for shared data objects appropriate for streaming aggregation. We propose two types of such objects (*T-Gate* and *W-Hive*) and their concurrent and lock-free algorithmic implementations, upon which we build three enhanced multiway aggregate operators that balance the work among concurrent threads and outperform existing

implementations in both order-sensitive and order-insensitive functions. We provide an extensive study using two large datasets extracted from the SoundCloud<sup>1</sup> social media and from a Smart Grid network. For both datasets the enhanced aggregation resulted in large improvements, up to one order of magnitude, both in terms of processing throughput and latency.

Our contributions open up space for new research questions for the role of concurrent data structures in parallel data streaming and are expected to influence significantly the design and implementations of parallel SPEs.

The paper is organized as follows. Section 2 introduces the data streaming processing paradigm and the multiway aggregate operator. Section 3 presents the state of the art implementation of data streaming multiway aggregation and, by rethinking parallelism in this context, discusses how its efficiency can be enhanced by means of concurrent data structures. Sections 4 and 5 present a detailed overview of the algorithmic design and implementation of the enhanced operators and data structures that we propose. In Section 6 we show the liveness and safety properties, namely lock-freedom and linearizability, of the proposed operators and data structures implementations. Section 7 presents the experimental evaluation. We discuss related work in Section 8 and conclude in Section 9.

### 2 Data Streaming and Multiway Aggregation

A stream is defined as an unbounded sequence of tuples  $t_0, t_1, \ldots$  sharing the same schema composed by attributes  $\langle ts, A_1, \ldots, A_n \rangle$ . Given a tuple t, attribute t.ts represents its creation timestamp while  $A_1, \ldots, A_n$  are application-related attributes. Following the data streaming literature (e.g., [2, 10]), we assume that each stream contains timestamp-sorted tuples.

*Continuous queries* (or simply queries) are defined as DAGs of *operators* that consume and produce tuples. Operators are distinguished into *stateless* or *stateful*, depending on whether they keep any state that evolves with the tuples being processed. Due to the unbounded nature of streams, stateful operations are computed over a *sliding window*, defined by parameters *size* and *advance*. Sliding windows can be time-based (e.g., to group tuples received during periods of 5 minutes every 2 minutes) or tuple-based (e.g., to group the last 10 received tuples every 3 incoming tuples). We focus in this paper on time-based sliding windows (or simply windows). We use POSIX notation<sup>2</sup> to specify the periods covered by a window and assume all windows start at time 0. That is, a window with size and advance of 10 and 2 seconds, respectively, will cover periods [0,10), [2,12), [4,14), and so on.

The multiway aggregate operator consumes an arbitrary number of input streams and is defined by its window's *size* and *advance*, by a function F applied to the tuples and by an optional *group-by* parameter K (a subset of the input tuple's at-

<sup>&</sup>lt;sup>1</sup>https://soundcloud.com/

<sup>&</sup>lt;sup>2</sup>Defined as the number of seconds elapsed since Thursday, 1 January 1970

tributes, also referred to as the tuple's key). Functions F can be order-sensitive (e.g., forward only the first received tuple) or order-insensitive (e.g., count the number of tuples) with respect to the processing order of the tuples that contribute to the same window. If K is defined, function F is computed for each distinct value of the group-by parameter. In this case, the operator keeps separate windows not only for different time intervals, but also for different values of K.

We define a *winset* as the set of windows covering the same time interval for different values of K. As an example, suppose an aggregate operator consumes tuples composed by attributes  $\langle ts, meter, consumption \rangle$  (each referring to the consumption reported by a meter at time ts) and computes the average consumption for K = meter and for a window with size and advance of 10 and 2 seconds, respectively. In this case, winsets will hold the windows covering interval [0,10) for each distinct meter, the windows covering interval [2,12) for each distinct meter, and so on.

Scenarios such as parallel-distributed SPEs [10, 3] and replica-based fault tolerant SPEs [2], demand for deterministic aggregation of input tuples. Processing of a multiway aggregation is deterministic if tuples are processed in timestamp order (when F is order-sensitive) or if all the tuples contributing to the same window are processed before producing the result (when F is order-insensitive). To this end, a parallel execution of a multiway aggregation must ensure that tuples are not simply processed in the order they are received [10] (i.e., input tuples from different input streams are not arbitrarily interleaved). To ensure deterministic processing, tuples from multiple input streams need to be *merged* into one sequence and *sorted* in timestamp order [10], an operation we refer to as *S-Merge*. We say that a tuple is *ready* to be processed if at least one tuple with an equal or higher timestamp has been received at each input stream (we refer the reader to Section 3 for an example focusing on the processing of *ready* tuples).

**Definition 1.** Let  $t_i^j$  be the *i*-th tuple received from input *j*.  $t_i^j$  is ready to be processed if  $t_i^j$ .ts  $\leq$  merge<sub>ts</sub>, where merge<sub>ts</sub> = min<sub>k</sub>{max<sub>l</sub>( $t_l^k$ .ts)} is the minimum among the latest (over l) tuple timestamps received from every input k.

Thus, we can formalize deterministic aggregate operators for both order sensitive and insensitive functions as follows:

**Definition 2.** An aggregate operator implementation is deterministic if in every execution it is guaranteed that the output tuples are computed from ready tuples and they are processed in the order imposed by the respective aggregate function F.



Figure 1: Sample query to count the power outages reported by smart meters and sample execution of the aggregate operator for a sliding window with size and advance of 3 and 2 time units, respectively.

# **3** Rethinking aggregation's parallelism: the role of data structures

The multiway aggregate operator is composed by four main stages:

- 1. Add: fetching incoming tuples from each input stream,
- 2. S-Merge: merging and sorting of input streams' tuples,
- 3. Update: updating of the windows a tuple contributes to, and
- 4. Output: forwarding of output tuples.

Figure 1 presents a sample multiway aggregation query used in a Smart Grid application to count the number of power outages reported by each meter over a sliding window with size and advance of 3 and 2 time units, respectively. A mesh network of Smart Meters (SMs) forwards such alarms to a set of Data Concentrators (DCs) [11], which in turn produce the timestamp sorted input streams. Notice that, being a mesh network, messages generated by the same SM could be forwarded by distinct DCs. The input tuples' schema  $\langle ts, meter \rangle$  specifies the time ts at which the alarm forwarded by a given meter has been received by a DC. Tuples produced by A are composed by attributes  $\langle ts, meter, \#alarms \rangle$  and specify the number of alarms generated by each meter for the window starting at time ts. In the example, the aggregate operator has two input streams. Input tuple  $\langle 4, SM_1 \rangle$ is received at the first input stream, while tuple  $\langle 6, SM_1 \rangle$  is received at the second input stream. The figure presents the different steps performed by the operator for a given initial state. Notice that, given Definition 1, tuple  $\langle 4, SM_1 \rangle$  is *ready* to be processed while  $\langle 6, SM_1 \rangle$  is not.

**State of the art** Widely used SPEs such as Borealis [2] or StreamCloud [10] perform multiway aggregation by relying on per-input queues to store incoming tuples. Distinct threads, which we refer to as input threads  $I_t$ , insert tuples to such queues while a dedicated output thread  $O_t$  processes them. Concurrent accesses are synchronized with the help of locks. Figure 2a presents this design, which we refer to as Multi-Queue (MQ). The output thread  $O_t$  peeks the first tuple in each

queue to determine which one is *ready* to be processed. The same thread is also responsible for the *Update* and *Output* operations. Since  $O_t$  is the only thread in charge of updating windows, no locking mechanism is required to access the *winsets*, usually implemented as hash tables to easily support arbitrary numbers of windows and to locate them quickly given the group-by parameter K.

**Parallelization challenges** In existing implementations, *S-Merge* usually relies on simple sorting techniques, whose cost is linear to the number of inputs. Examples include the Input Merger operator [10] or the SUnion operator [2]. In order to prevent such sorting techniques from becoming a bottleneck and allow for the processing of tuples coming from arbitrary number of input streams, the first challenge relies on the (1) parallelization of the *S*-Merge operation. It should be noticed that, an enhanced parallel sorting technique should still ensure correct pick-up of ready tuples. The second challenge relies on the (2) parallelization of the Update stage. To guarantee deterministic processing (as discussed in Section 2), Update cannot be invoked in parallel on tuples contributing to the same window and sharing the same K value (or when no group-by parameter is defined) for order-sensitive functions. This restriction can be relaxed for order-insensitive functions, since the result of a window would not be affected by the order in which concurrent threads update it. Since the aggregate operator defines a single output stream to which tuples are added in timestamp order, we do not take into account the parallelization of the Output function.

**Utilizing concurrent data structures** A core challenge in the parallelization of the pipeline stages of multiway aggregation is the "*balancing act*" [19] of maximizing their concurrency while ensuring consistency and correct synchronization. To this end, the key-enablers that can address such challenge are the data structures, seen as articulation points between such stages, as well as their efficient algorithmic implementations. Efficient implementations of data structures often employ fine-grain synchronization that can avoid the use of waiting or locking. Lock-free data structures have been shown to increase applications' throughput and are part of the Java and C# standard libraries. The correctness of such implementations is commonly shown through *linearizability* [14], which guarantees that, given a history of concurrent operations, there exists a sequential ordering of them, consistent with their real-time ordering and with the sequential semantics of the data structure.

In order to parallelize the *S-Merge* and *Update* stages, we first explored which existing concurrent data structures could be used to sort input tuples at insertion time. In principle, tree-like data structures could provide concurrent logarithmic-time insertion operations. The need for easily extracting such tuples in timestamp order would be better addressed by a concurrent skip list (e.g., the one proposed by [29]) due to its underlying list-like node structure of sorted elements with short-cuts allowing for fast insertion (cf. Sec. 5.1 for more details on skip lists). Never-

Table 1: Methods supported by the data structures

T-Gate		W-Hive	
insertTuple(tuple,	Inserts a tuple from input	updateWindows(tuple,	Updates the windows that
input)	stream in sorted order.	thread)	the tuple contributes to.
getNextReadyTuple	() Returns (once and only	getNextWinSet()	Returns (once and only
	once) the earliest ready tu-		once) the earliest winset to
	ple (cf. definition 1).		which tuples do not con-
			tribute anymore.



Figure 2: Overview of aggregate designs.

theless, a skip list would not differentiate between tuples that are *ready* and tuples that are not. Because of that, checking whether a tuple is *ready* or not would still be penalized by a cost that is linear to the number of inputs, as is the case for the multi-queue implementations. Based on this observation, while leveraging the skip list's multi-level shortcuts mechanism (allowing for a logarithmic find of the insert position in the list), we propose new concurrent shared data object types that better fit the parallelization challenges proper of multiway aggregation. We complement the qualitative estimation of the reasons that motivated the design and implementation of new concurrent data structures by comparing them with a lock-free skip list in Section 7.

# 4 New abstract data types and aggregate designs

This section overviews our enhanced aggregate operators. For all of them, one input thread per stream,  $I_t$ , fetches tuples from its respective input stream while a single output thread,  $O_t$ , forwards output tuples. Figure 2 presents the different designs and describes how operations are assigned to threads. While presenting the different designs, we discuss the data structures needed to maintain tuples and *winsets*, and introduce our concurrent data structures and their APIs (Table 1), with the functionality of each method.

**Tuple Merged List - Single Consumer (TuML**<sub>SC</sub>) This algorithmic design (Fig. 2b) addresses the first parallelization challenge by performing both *Add* and *S-Merge* in parallel. TuML<sub>SC</sub> relies on the Tuple-Gate (T-Gate), a concurrent data structure whose API provides two methods, whose definitions are given in Table 1. Method insertTuple(tuple, input) allows for tuples to be inserted, while being merged and sorted, by multiple input streams in parallel. Method getNext ReadyTuple() guarantees that no tuple is returned for processing before it is *ready*. The T-Gate stores an ordered list of tuples. By keeping track of the latest added tuple from each input stream, the method can quickly check if the first tuple in the list is *ready*. The output thread reads sorted and *ready* tuples from the T-Gate and performs the *Update* and *Output* stages.

**Tuple Merged List - Multiple Consumer (TuML**<sub>MC</sub>) TuML<sub>MC</sub> (Fig. 2c) extends the TuML<sub>SC</sub>, addressing the second parallelization challenge by performing also the *Update* operation in parallel. Multiple update threads,  $U_t$ , get *ready* tuples from T-Gate concurrently by invoking getNextReadyTuple() and update the windows to which each tuple contributes to. Thus, winsets are now accessed and updated concurrently by the  $U_t$  threads. For managing the winsets and synchronizing such access we introduce a second data structure, that we refer to as Window-Hive (W-Hive).

As the T-Gate encapsulates the logic to differentiate between tuples that are *ready* or not, the W-Hive is able to differentiate between the winsets to which incoming tuples are still contributing and the ones whose results can be outputted. It provides two methods: updateWindows(tuple, thread) allows for multiple threads to synchronize and safely create and update active winsets while getNext WinSet() returns the earliest winset no longer being updated by any thread. This method is invoked by the output thread  $O_t$ , in charge of forwarding the operator's output tuples. W-Hive uses similar techniques as the T-Gate to quickly find the right location of where to insert a new winset. To preserve the correctness of ordersensitive functions, each update thread is responsible for a distinct subset of the group-by parameter values K. For this implementation, the number of update threads can be chosen by the user.

Window Merged List (WiML) This design (Fig. 2d) further enhances the parallelization of the aggregate's stages for order-insensitive functions. Operations *Add, S-Merge* and *Update* are performed in parallel by the  $I_t$  threads. Since WiML is designed for order-insensitive functions, input tuples do not need to be sorted before being processed to update the windows they contribute to. The required synchronization needed to ensure that output tuples for a given winset are outputted only after all its contributing tuples have been processed, is managed by the getNextWinSet() method provided by the W-Hive.

## 5 Aggregates and data structures implementations

In this section we present in detail the aggregate algorithmic implementations and their supporting data structures. The pseudocode for the baseline aggregate implementation can be found in Algorithm 2, while the enhanced implementations in Algorithms 3 and 6. The supporting data structures are presented in Algorithms 4 and 5. Methods' names have been chosen according to which of the four main aggregate's stages specified in Section 3 they implement. In the following, the group-by value of each tuple is accessed as tuple.key. If no group-by parameter is defined, it is safe to assume all tuples will refer to the same key value (e.g., null).

#### 5.1 Preliminaries

A skip list [22] is a data structure that maintains elements in an ordered list and supports probabilistically logarithmic search, insertion and deletion operations. Essentially, a skip list can be viewed as a traditional linked list where each node, besides the usual pointer connecting to the next element, has a tower of several pointers that shortcut over the next elements and connect to nodes later in the ordered list. The height of the nodes is randomly distributed so that 50% of the nodes have height 1, 25% of them have height 2 and so on. Thus, the higher the level traversed, the sparser the links are and more nodes are skipped. The basic search routine for a key k, is to traverse from the highest level shifting to a lower one every time the current node's key is greater than k.

One of the main benefits of skip lists over standard tree like data structures is that regardless of the data and operation distribution there is no need for rebalancing. This has made it a good candidate for parallel and concurrent implementations, as the one by [29], since rebalancing will typically require expensive synchronization in tree-based implementations.

#### 5.2 Common components

Algorithm 1: Generic window interface and concrete implementation of a window that sums the value of attribute  $A_i$ .

```
interface Window
1
     void processTuple(tuple) // update variables
2
     Tuple produceOutTuple() // produce output tuple
3
4
   class SumWindow : Window
5
     int sum = 0
     void processTuple(tuple)
7
       sum += tuple.A_i
8
     Tuple produceOutTuple()
9
10
       return Tuple(sum)
```

The base component of the aggregate operator is the Window. It represents a time interval and provides functionality to aggregate the tuples that contribute to it. Algorithm 1 shows the window interface and the implementation of a sample sum aggregation.

As discussed in Section 3, the *winset* can be implemented as a hash table to easily support an arbitrary number of windows and to locate them quickly given the tuple's group-by parameter K. In the WiML and  $TuML_{MC}$  implementations, the winset is accessed by multiple threads concurrently. In data streaming applications the domain of keys for the group-by parameters are typically known in advance. Thus, the use of a closed addressing hash table is appropriate; specifically, in our algorithmic implementation we are using the lock-free, linearizable concurrent hash table by Michael [17], mainly due to its implementation simplicity. Alternatively, open addressing schemes like lock-free cuckoo hashing [21] can be used. Furthermore, the hash table (winset) is a building a block in another lock-free and linearizable data structure (W-Hive). Therefore, we avoided designs based on blocking implementations [13] or explicit hardware support [15]. Shun and Blelloch [26] present a high performing phase-concurrent hash table. In this model only operations of the same type proceed concurrently, which is a limitation in the winset use-case since insert and find operations may be concurrent. For the MQ and  $TuML_{SC}$  implementations, since the access to the winset is sequential, a sequential implementation of a hash table is sufficient.

#### 5.3 **Baseline implementations - MQ**

#### Algorithm 2: MQ

```
Add(tuple, input) // One thread per input
11
     queueinput.enqueue(tuple)
12
13
   SMergeUpdateOutput() // One thread
14
15
     if (\exists i : queue_i.isEmpty()) return
16
      input = v: (\forall i:queue_v.peek().ts \leq queue_i.peek().ts)
     tuple = queue<sub>input</sub>.dequeue()
17
     upout (tuple)
18
19
   upout(tuple)
20
21
      windowTSs = getTargetWindowTSs(tuple)
22
     while(windowlist.first().ts<windowsTSs.first())</pre>
23
        winset = windowlist.removeFirst()
        for (window : winset)
24
          forward(window.produceOutTuple()) // See L3
25
      for (wts : windowTSs)
26
        if(!windowlist.contains(wts))
27
          windowlist.insert(wts, new WinSet(wts))
28
        win = windowlist.find(wts).find(tuple.key)
29
        if (win == null)
30
31
          win = new Window()
          windowlist.find(wts).put(tuple.key, win)
32
        win.processTuple(tuple) // See L2
33
```

This baseline implementation is based on the one used in SPEs such as Borealis [2] or StreamCloud [10]. The multi-queue design consists of two main methods (see Algorithm 2). The Add method is used to deliver tuples to the aggregate and placing them in their respective input queue (L12). The queues are protected by a lock to allow concurrent access.

The main work is performed by the second method, SMergeUpdateOutput. It checks all the queues to make sure a tuple has been received from each input (L15). It then reads the tuple with the lowest timestamp among the inputs (L16-17). This guarantees that all tuples will be read in timestamp order.

The currently active winsets are stored in a linked list. The method getTargetWindowTSs creates a list, windowTSs, of the starting timestamps of the windows that the tuple contributes to. If the starting timestamp of a winset in the window list is lower than the earliest timestamp in windowTSs, the aggregated results of the former can be outputted (L22-25). This is safe since all future tuples will have an equal or higher timestamp and will not contribute to the winset. If the new tuple contributes to a time interval that does not have a corresponding winset yet, the winset is created and added to the list (L27). If the window does not exist for the tuple's key, it is also created (L30). Finally, the window processes the tuple (L33).

Figure 3 presents how stages Add, S-Merge, Update and Output (and their respective code lines) are distributed to threads  $I_t$  and  $O_t$  for the MQ implementation (stages assigned to  $I_t$  and  $O_t$  threads are colored in blue and red, respectively).



Figure 3: Visual representation of how stages (and their respective code lines) are distributed to threads for the MQ implementation.

#### 5.4 **TuML**<sub>SC</sub> and **TuML**<sub>MC</sub>

These aggregate designs rely on the T-Gate data structure (API in Table 1 and further description in Section 5.5). The T-Gate is used to pre-sort all arriving tuples and merge them into one stream. In contrast with the MQ implementation, the *S*-*Merge* operation is now executed at the first stage in the pipeline.

TuML<sub>SC</sub> uses a single thread to read the sorted tuples from the T-Gate, update the windows, and output the aggregated results. This is done using the method UpdateOutput, which shares much functionality with the MQ design (L37). TuML<sub>MC</sub> allows multiple threads to read from the T-Gate and update the windows

# Algorithm 3: TuML<sub>SC</sub>, TuML<sub>MC</sub>

```
AddSMerge(tuple, input) // One thread per input
34
35
      tgate.insertTuple(tuple, input) // See L89
36
    \texttt{UpdateOutput()} \quad \textit{// TuML}_{SC} \quad \textit{only}
37
38
      tuple = tgate.getNextReadyTuple() // See L82
      windowTSs = getTargetWindowTSs(tuple)
39
      // Produce results for windows no longer updated
40
      while(windowlist.first().ts<windowsTSs.first())</pre>
41
        winset = windowlist.removeFirst()
42
43
         for (window : winset)
           forward(window.produceOutTuple())
44
      // Update windows
45
46
      for (wts : windowTSs)
        if(!windowlist.contains(wts))
47
48
           windowlist.insert(wts, new WinSet(wts))
        win = windowlist.find(wts).find(tuple.key)
49
        if (win == null)
50
51
           win = new Window()
           windowlist.find(wts).put(tuple.key, win)
52
        win.processTuple(tuple)
53
54
    Update() // Multiple threads - TuML<sub>MC</sub> only
tuple = tgate.getNextReadyTuple() // See L82
55
56
57
      if (tuple == null) return
      if(¬tuple.hashToThread(threadid)) return
58
59
      whive.updateWindows(tuple) // See L125
60
    \textbf{Output}\left(\right) // One thread - \texttt{TuML}_{MC} only
61
      winset = whive.getNextWinSet() // See L117
62
      if(winset == null) return
63
      for(window : winset)
64
65
         // Output the result of the window
        forward(window.produceOutTuple())
66
```

in parallel. This requires support for concurrent handling of the winsets. The W-Hive (API in Table 1 and further description in Section 5.6) is used to provide lock-free winset management. If the aggregate opertor's function is order-sensitive (e.g., forward only the first received tuple), tuples contributing to the same window cannot be processed in parallel by multiple threads. Hence, a hash function based on the group-by attribute is used to assign input tuples to existing threads.

#### 5.5 T-Gate

#### Algorithm 4: T-Gate

```
Node head, update[maxlevels] // Thread local variables; maxlevels is a
 67
         constant parameter
 68
 69
    def Node
     Node next[maxlevels]
 70
 71
      Tuple tuple
      int input
 72
 73
 74 initializeTGate()
      tail = new Node()
 75
      tmp = new Node()
                                // tmp is the temporary head
76
      for (i=0 to maxlevels-1) // all levels point to tail
 77
78
        tmp.next_i = tail
79
      for (i in input ids)
        insertTuple(new Tuple(), i) //insert one dummy tuple per input
 80
81
 82
    getNextReadyTuple()
      next = head.next_0
83
84
      if(next≠tail ∧ written<sub>next.input</sub> ≠next.tuple)
         head = next
 85
        return next.tuple
86
 87
      return null
 88
89 insertTuple(tuple, input)
 90
     nodeheight = getLevelHeight()
 91
      newnode = new Node(tuple, input)
      curnode = update_{maxlevels-1}
92
     for(i=maxlevels-1 downto 0)
 93
        next = curnode.next<sub>i</sub>
 94
        while(next≠tail ∧ next.ts<tuple.ts)</pre>
 95
          curnode = next
 96
97
          next = curnode.next_i
        update_i = curnode
98
       for(i=0 to nodeheight)
99
         levelinsert(update<sub>i</sub>, newnode, tuple.ts, i)
100
101
       written_{input} = newnode
102
    levelinsert(priornode, newnode, ts, level)
103
104
      while(true)
         next = priornode.nextlevel
105
106
         if(next==tail V next.ts>ts)
107
          newnode.next_{level} = next
           if(CAS(priornode.next<sub>level</sub>, next, newnode)) break
108
109
         else fromNode = next
```

The T-Gate data structure (see Algorithm 4) maintains a merged, timestamp ordered list of the tuples coming from the input streams.

The insertTuple method inserts a tuple at its correct position in the list, given its timestamp. First the number of shortcut levels is decided (L90), according to the standard skip list distribution [29]. Each thread keeps in the update array a pointer to the last accessed node in each level (L92). Since all new tuples added by the same thread will have an equal or higher timestamp than the last inserted one, this lowers the number of nodes that a thread has to examine.

Once all the levels are searched, the node is inserted on each level it should be part of with the use of the levelinsert helper method (L100). This method verifies that the conditions for the prior node in each level still apply, otherwise (in case some newer node has been inserted in between) it traverses the current level of the list until the right position is found. The node is then inserted by using the compare and swap (CAS) atomic instruction. In case of failure, i.e. when another thread achieves an insertion at the same place, the loop retries the search. When the node has been inserted, the written array is updated to hold a reference to the new node (L101). The index into the array is the input stream id. This is done to make sure a tuple is not read until we have received a new tuple with a higher or equal timestamp from all the other input streams.

The getNextReadyTuple method traverses the lowest level of the list to return tuples in timestamp order. Each thread keeps its local head pointer having its own handle to the list and advances this pointer in each successful call. A tuple pointed by the head can be returned if it is not the last one added by any input stream (the latter ensures that if a tuple is returned, it is indeed *ready*). It is useful to point out that in the case of just one tuple per input stream being present in the data structure, according to Definition 1, the tuple with the smallest timestamp is *ready*. However, the presented implementation will not return this tuple until another one with higher timestamp arrives from the same input stream. This is done for implementation simplicity, since it does not compromise the correctness according to Def. 1, and does not affect the high input rate scenarios which we focus in this paper.

Finally, during the initialization of the data structure (L74), one dummy tuple per input is inserted to ensure the correct semantics of the getNextReadyTuple are preserved until all input streams start delivering tuples.

Nodes can be freed when they are no longer accessible (directly or indirectly) from the thread local head and update<sub>maxlevels-1</sub>. For this reason, several memory reclamation techniques such as hazard pointers can be applied [18, 29], while also garbage collection can be exploited. In the Java based implementation of our prototype that is evaluated in Section 7, we rely on the default garbage collector.

#### 5.6 W-Hive

The W-Hive (cf. Algorithm 5) data structure provides lock-free management of winsets. The updateWindows method adds a tuple to each window it contributes

#### Algorithm 5: W-Hive

```
110 Node readhead, inserthead, tail
111
    def Node
112
      Node next[maxlevels]
113
114
       Timestamp ts
       WinSet winset
115
116
117
    getNextWinSet()
       \texttt{if(readhead.next_0==tail) return null}
118
119
       \texttt{if(readhead.next}_{0}.\texttt{ts}{\in}\texttt{written)} \texttt{ return null}
       readhead = readhead.next_0
120
       if(readhead.levels==maxlevels)
121
122
         inserthead = readhead
       return readhead.winset
123
124
    updateWindows(tuple, thread)
125
       windowTSs = getTargetWindowTSs(tuple)
126
127
       written<sub>thread</sub> = windowTSs.first()
128
       for(wints : windowTSs)
         curnode = inserthead
129
130
         for(i = maxlevels-1 downto 0)
           next = curnode.next_i
131
132
           while (next!=tail \land next.ts \leq wints)
              curnode = next
133
             next = curnode.next_i
134
135
           update_i = curnode
         if(curnode.ts != wints)
136
           winset = new WinSet(wints)
137
138
           levels = getLevelHeight()
           newnode = new Node(wints, winset)
139
           curnode = levelinsert(update0, newnode, wints, 0)
140
           if(curnode == newnode)
141
              for(i=1 to levels-1)
142
143
                levelinsert(update<sub>i</sub>, newnode, wints, i)
         win = curnode.winset.find(tuple.key)
144
145
         if(win==null)
146
           win = new Window()
           curnode.winset.put(tuple.key, win)
147
148
         win.processTuple(tuple)
149
    levelinsert(priornode, newnode, wints, level)
150
151
       while(true)
         next = priornode.next<sub>level</sub>
if (level == 0 \land next.ts == wints) return next
152
153
154
         if(next == tail V next.ts > wints)
           newnode.next_{level} = next
155
156
            if (CAS (priornode.next_{level}, next, newnode) break
         else fromNode = next
157
       return newnode
158
```

to. A reference to the earliest such window is saved in the written array for each thread (L127). This is used to keep track of when winsets are no longer being updated (L119). For each window the tuple contributes to, the method traverses the list to locate the winset with the same timestamp as the window. This is done in the same manner as when inserting a node into the T-Gate (L130-135). If the winset is found, it is searched to find the correct window for the tuple's key (L144). If there is no window for the key, a new window is inserted into the winset with the correct key (L147). The tuple is then added to the window. If no winset is found for the timestamp, a new winset and corresponding node to hold it are created. They are inserted into the list in a similar manner to the T-Gate. The difference is that another thread might try to create a winset for the same timestamp concurrently. If this happens and the other thread manages to insert it, then the insertion must be canceled and the other winset will be used instead (L153).

The getNextWinSet operation returns the next winset that is no longer being updated by input tuples. It is assumed that it will only be called by a single thread. If no thread updated any of the windows in the first winset of the list the last time it received a tuple, it can be assumed that no more tuples will contribute to the winset in the future, as each thread receives tuples in timestamp order (L119). If the new head node for the getNextWinSet operation is part of all shortcut levels, it is made the new head node for the updateWindows method. Nodes and winsets with a timestamp lower than the ones referenced by inserthead and readhead can be safely freed or automatically garbage collected.

Figure 4 presents how stages Add, S-Merge, Update and Output (and their respective code lines) are distributed to threads  $I_t$  and  $O_t$  for the TuML<sub>SC</sub> implementation (stages assigned to  $I_t$  and  $O_t$  threads are colored in blue and red, respectively). Similarly, Figure 5 shows how such stages are distributed to threads  $I_t$ ,  $U_t$  and  $O_t$  for the TuML<sub>MC</sub> implementation (stages assigned to  $I_t$ ,  $U_t$  and  $O_t$  threads are colored in blue, red, and yellow, respectively).



Figure 4: Visual representation of how stages (and their respective code lines) are distributed to threads for the  $TuML_{SC}$  implementation.

#### 5.7 WiML

The WiML design (see Algorithm 6) is suitable only for aggregate operator's functions F that are order-insensitive, since it does not sort the tuples prior to insert-



Figure 5: Visual representation of how stages (and their respective code lines) are distributed to threads for the  $TuML_{MC}$  implementation.

Algorithm 6: WiML

```
159
    AddSMergeUpdate(tuple, input) // One thread per input
      whive.updateWindows(tuple, input)
                                             // See L125
160
161
    void Output() // One thread
162
163
      winset = whive.getNextWinSet() // See L117
      if(winset == null) return
164
      for(window : winset)
165
         // Output the result of the window
166
        forward(window.produceOutTuple())
167
```

ing them into their windows. When a tuple arrives it is immediately processed to update the windows it contributes to. This is done in the AddSMergeUpdate method using the W-Hive (L 160). The W-Hive returns the winsets that will no longer be contributed to, which can then be forwarded (L167).

Figure 6 presents how stages Add, S-Merge, Update and Output (and their respective code lines) are distributed to threads  $I_t$  and  $O_t$  for the WiML implementation (stages assigned to  $I_t$  and  $O_t$  threads are colored in blue and red, respectively).



Figure 6: Visual representation of how stages (and their respective code lines) are distributed to threads for the WiML implementation.

### 6 Correctness

In this section we outline proofs of liveness and safety properties of the algorithmic constructions of the data structures, namely lock-freedom and linearizability. *Lock-freedom* guarantees that at least one of the concurrent method call invocations of the data structure will return in a finite number of its own steps [12]. *Linearizability* [14] guarantees that every method call appears to take effect at some point (linearization point) between its invocation and response; more formally, for a linearizable implementation of a data structure, given a history of concurrent operations, there exists a sequential ordering of them, consistent with their real-time ordering and with the sequential semantics of the data structure. Furthermore, we show that the aggregate implementations provide deterministic processing of the stream tuples.

# **Theorem 1.** *The T-Gate implementation presented in algorithm 4 is lock-free and linearizable.*

*Proof.* The getNextReadyTuple method does not contain any loops and returns in a bounded number of its own steps. The insertTuple method contains bounded loops except for the levelinsert subroutine. This will fail to terminate only if the CAS instruction on L108 fails, i.e. in the case a concurrent call of insertTuple from another thread makes progress. Therefore, the T-Gate implementation is lock-free.

The linearization point of the insertTuple method during concurrent calls of the same method, is the successful CAS on L108, as this is when the operation appears to take effect among such calls.

getNextReadyTuple is linearized at the check on L84 when the appropriate cell of the written array is read. In the case of concurrent calls of methods getNextReadyTuple and insertTuple, the linearization point of the latter is the update of the written array on L101. Thus there is a linearization point for all the method calls of the T-Gate implementation.

The W-Hive provides management of winsets and is where the actual aggregate computation takes place. The processTuple call on L148 is an application-specific operation. Naturally, the safety and liveness properties of the

processTuple method call affect the ones of the higher level updateWindows method that includes the former. Thus, the following theorems are shown under the condition that the processTuple call on L148 is linearizable and lock-free (e.g. local computation in the simplest case).

# **Theorem 2.** *The W-Hive implementation presented in algorithm 5 is lock-free and linearizable.*

*Proof.* By definition there are no concurrent calls of getNextWinSet, each such call does not modify any shared variables and returns in a bounded number of its own steps. A call to updateWindows will fail to return only if the CAS

instruction on L156 fails (i.e., if a concurrent call from another thread will have made progress). In the case of concurrent updates of tuples with the same key, the winsets used are lock-free and linearizable, thus so are all the calls to their methods. Therefore, the W-Hive implementation is lock-free.

For concurrent calls to the updateWindows method we distinguish two cases: the ones that successfully add a new winset and the respective holding node to the W-Hive and the ones that update an existing winset. For the latter, the linearization point breaks down to the successful find of the window to be updated in the winset (L144), or the insertion of the respective window in case this does not exist (L147). For the former, the linearization point is the successful CAS instruction on L156. A call to updateWindows concurrent with a call to getNextWinSet is linearized on L127, as this could affect the result of a subsequent check on L119 of getNextWinSet. The linearization point of getNextWinSet can be any of L118 or L119 depending on the successful checks, or L120 otherwise. Thus, there is a linearization point for all the method calls of the W-Hive implementation.

**Lemma 1.** A tuple  $t_i^j$  returned by the getNextReadyTuple method, satisfies the ready definition (cf. Def. 1).

*Proof.* Assume towards a contradiction that  $t_i^j > merge_{ts}$ . Then  $t_i^j$  would be the tuple with the latest timestamp received by its respective input thread. But then the check in line 84 would have failed and  $t_i^j$  would not have been returned.

**Lemma 2.** All tuples contributing to a winset returned by the getNextWinSet satisfy the ready definition (cf. Def. 1).

*Proof.* Assume there is a tuple that does not. As above, the tuple would be the one with the latest timestamp received by its respective input thread. In that case the check in line 119 would have caused null to be returned, as the winset's timestamp would belong to the windows in the written array.  $\Box$ 

**Theorem 3.** The  $TuML_{SC}$ ,  $TuML_{MC}$  and WiML aggregate implementations (Alg. 3, 6) are lock-free and provide deterministic processing of tuples.

*Proof.* All method calls are bounded by a constant and include either local computations for each thread, or calls to data structure implementations that are lock-free (T-Gate,W-Hive). Thus, the implementations are lock-free.

Lemmas 1 and 2 show that the output tuples of the aggregate implementations always consist of ready tuples. Therefore, the aggregate implementations are deterministic (cf. Def. 2).  $\Box$ 

# 7 Evaluation

In this section, we study the performance of the different aggregate operators presented in Section 4 in terms of throughput and latency. We also include the evaluation of a lock-free version of the multi-queue implementation, referred to as  $MQ_{LF}$  and relying on the lock-free queue by Michael and Scott [20], with the sole purpose of showing that MQ's poor performance does not depend only on the use of locks. First, we provide evidence of the superiority of the T-Gate with respect to a lock-free skip list by measuring their maximum throughput and latency. Subsequently, we discuss the improvement enabled by  $TuML_{SC}$  and WiML, compared to MQ and  $MQ_{LF}$ , for different queries and varying number of inputs. Our goal is to show how the TuML<sub>SC</sub> and WiML implementations can achieve higher performance than MQ and MQ<sub>LF</sub> ones by increasing the computing time over the synchronization time of their underlying threads. At the same time, we also show the scalability of the WiML implementation for an increasing number of threads running the Update stage (assigned to  $I_t$  threads in this case, as discussed in Section 5.7). In the last part of the section, we evaluate  $TuML_{MC}$ 's scalability for increasing number of threads running the Update stage, complementing the scalability evaluation of the WiML implementation. Our experiments take into account the aggregate's features that affect throughput and latency: the overall number of keys, the number of windows to which each tuple contributes and the cost of the aggregate function. For each feature, we consider 2 stretching points in order to show how traversing its spectrum (e.g., increasing the overall number of keys) affects the overall performance. All the experiments represent queries that can be found in real-world applications. We take into account aggregate functions that are commonly used and also evaluate highly costly variants when studying how their cost affects the aggregate performance. Our data sets have been collected from real-world applications.

#### 7.1 Evaluation setup

The evaluation has been conducted with an Intel-based workstation with two sockets of 6-core Xeon E5645 (Nehalem) processors with Hyper Threading (24 logical cores in total) and 48 GB DDR3 memory at 1366 MHz. The prototype has been implemented in Java and experiments have been run using the OpenJDK Runtime Environment (IcedTea 2.3.9) with the default garbage collection settings.

We use two datasets that we refer to as *SoundCloud* (*SC*) and *Energy Con*sumption (*EC*). *SC* has been collected from the online audio distribution platform SoundCloud from a subset of approximately 40,000 users exchanging comments about 250,000 songs between 2007 and 2013. Tuples contain comments sent by users in relation to songs and are composed by the attributes  $\langle ts, user, song, cmt \rangle$ . *EC* contains energy consumption readings collected from a set of 243 smart meters between May 2012 and June 2013. Tuples' schema is composed by attributes  $\langle ts, meter, cons \rangle$ .

Each experiment starts with a warm-up phase and ends with a cool-down phase. During the measuring phase (lasting five minutes) tuples are injected at a constant rate; throughput is measured as the average number of tuples/second (t/s) processed by the aggregate operator during the measuring phase while latency is measured as the average timestamp difference between each output tuple and the latest input

Table 2: Parameters for queries used in the evaluation. Identifiers *ID* are composed by 3 letters. The first letter (small or capital, representing smaller or larger values/computation-demands) represents the aggregate parameter being studied (k - overall number of keys, w - window size, f - applied function). The last two letters specify whether *F* is order-sensitive (OS) or order-insensitive (OI).

ID	DS	WS	WA	K	F	Description	
Order-sensitive (OS) functions (MQ, $MQ_{LF}$ , $TuML_{SC}$ )							
k-OS	EC	30	3	meter	first()	Forward the first consumption read-	
						ing, group by meter (243 distinct	
						keys)	
K-OS	SC	30	3	song	first()	Forward the first comment, group	
						by song (40,000 distinct keys)	
w-OS	EC	20	2	meter	first()	Forward the first consumption read-	
						ing, group by meter (each tuple	
						contributes to 10 windows)	
W-OS	EC	40	2	meter	first()	Forward the first consumption read-	
						ing, group by meter (each tuple	
						contributes to 20 windows)	
f-OS	SC	20	2	song	first-mail( <i>cmt</i> )	Forward the first comment contain-	
						ing a mail address, group by song	
F-OS	SC	20	2	song	first-mail/IP( <i>cmt</i> )	Forward the first comment contain-	
						ing a mail or an IP address, group	
						by song	
Order-insensitive (OI) functions (MQ, $MQ_{LF}$ , WiML)							
k-OI	EC	30	3	meter	count()	Count the number of consumption	
						readings, group by meter (243 dis-	
						tinct keys)	
K-OI	SC	30	3	song	count()	Count the number of comments,	
						group by song (40,000 distinct	
						keys)	
w-OI	EC	20	2	meter	avg(cons)	Compute the average consumption,	
						group by meter (each tuple con-	
			_			tributes to 10 windows)	
W-OI	EC	40	2	meter	avg(cons)	Compute the average consumption,	
						group by meter (each tuple con-	
6.01	aa	20	2			tributes to 20 windows)	
1-01	SC	20	2	song	count-mail( <i>cmt</i> )	Count the number of comments	
						containing a mail address, group by	
EOL	an	20	2			Song	
F-01		20	2	song	count-mail/iP( <i>cmt</i> )	count the number of comments	
						group by song	
						group by song	

tuple that produced it during the measuring phase. Finally, presented results are averaged over 10 runs. In each experiment, we deploy one instance of the aggregate, together with injectors, running at dedicated threads and maintaining per-second throughput statistics, and a sink running at a dedicated thread, collecting output tuples and maintaining per-second latency statistics. When running experiments with different input rates, we process data from the EC and SC datasets, modifying only the rate at which tuples are injected. In order to find the maximum throughput and the corresponding latency of a given setup, several experiments (for increasing input rates) are run, as long as results do not indicate the setup cannot sustain the injected rate. Table 2 presents the parameters of the queries used in the evaluation: identifier *ID*, dataset *DS*, window size *WS*, window advance *WA*, group-by parameter *K* and aggregate function *F*. For the *T*-*Gate* and *W*-*Hive*, the maximum possible height of a node (maxlevels) is set to 3 in all experiments.

#### 7.2 Skip list and T-Gate comparison

In this experiment, we evaluate the performance of a lock-free skip list and the T-Gate, measuring the maximum throughput with which EC tuples coming from multiple input streams can be sorted. For this comparison we used the

ConcurrentSkipListMap from Java's java.util.concurrent package, an implementation based also on [29].

Results for 5, 10, 15 and 20 input streams are presented in Fig. 7. It can be noted that, when using a skip list, checking for *ready* tuples (done explicitly since the skip list does not differentiates between tuples that are *ready* or not) results in a cost linear to the number of input streams. Thus, the skip list's throughput degrades while T-Gate throughput grows. For 20 input streams, the skip list is able to sort approximately 1.5 million t/s while the T-Gate reaches approximately 2.2 million t/s (1.5 times better). The T-Gate also achieves a lower sorting latency, approximately 1 ms for 20 input streams against the 1.6 ms latency of the skip list.



Figure 7: T-Gate and Skip List comparison.

#### 7.3 Baseline and new designs comparison

In this set of experiments, we measure the maximum throughput and the latency of the MQ,  $MQ_{LF}$ ,  $TuML_{SC}$  and WiML implementations, quantifying the improvement enabled by the use of concurrent data structures while using the same number of input and output threads.

**Parallelization benefit** We first focus on the average duration of the main operations performed by  $I_t$  and  $O_t$  threads for queries K-OS and K-OI and 20 input streams. Results for the query K-OS are presented in Fig. 8 (we use logarithmic scale to better appreciate the different orders of magnitude). Both  $I_t$ 's and  $O_t$ 's operations are faster for MQ<sub>LF</sub> compared to MQ since the former relies on lockfree queues.  $I_t$ 's duration increases while  $O_t$ 's decreases for TuML<sub>SC</sub> since the *S-Merge* operation is performed by  $I_t$ .  $O_t$  threads, which constitute the bottleneck, will reach 100,000, 130,000 and 160,000 t/s for MQ, MQ<sub>LF</sub> and TuML<sub>SC</sub>.



Figure 10: K-OI Figure 8: K-OS  $I_t$ 's Figure 9: K-OS la- $I_t$ 's and  $O_t$ 's durations. tency evolution. The figure 11: K-OI lations. tency evolution.

Figure 12: Throughput and latency improvements enabled by T-Gate and W-Hive in  $TuML_{SC}$  and WiML implementations.

respectively. Figure 9 compares the latency evolution of the different implementations for an increasing input rate. In all experiments, the latency initially decreases with the increasing rate (lower inter-arrival times at the inputs result in shorter queuing times for input tuples) while it explodes upon saturation of the operator (that is, when the injected load exceeds its maximum throughput). The throughput achieved by each implementation is close to the expected one (from Fig. 8). TuML<sub>SC</sub> achieves a throughput 1.6 times higher than MQ. Average durations for query K-OI are shown in Fig. 10. It can be noticed that  $I_t$ 's duration increase is greater for WiML than TuML<sub>SC</sub> since both *S-Merge* and *Update* operations are performed by the  $I_t$  thread. With WiML, each  $I_t$  thread is able to process 22,000 t/s in parallel. The highest throughput that can be achieved by WiML is in this case given by  $O_t$ , since the latter can only produce up to 2 million t/s. As shown in Fig. 11, WiML achieves a higher throughput and a lower latency compared to MQ<sub>LF</sub> and MQ (approximately 400,000 t/s), independently of the input rate (tuples are processed independently by each  $I_t$  thread).

In the remaining of this section, we study the performance of  $\text{TuML}_{SC}$  and WiML with respect to MQ and MQ<sub>LF</sub> for the different queries presented in Table 2. In all experiments, we experience the same performance behavior when comparing MQ, MQ<sub>LF</sub>, TuML<sub>SC</sub> and WiML implementations, as explained in the following. As the number of inputs increases, the latency of multi-queue implementations (MQ and MQ<sub>LF</sub>) increases while their throughput decreases linearly. This pattern is broken by the TuML<sub>SC</sub>, whose throughput is rather stable as the number of inputs increases. The pattern is even reversed by the WiML, whose throughput actually increases as the number of inputs increases (since stage *Update* is run in parallel by each  $I_t$  thread) while achieving the lowest and almost constant latency.

**Varying number of keys** The overall number of keys in the data affects both the operator's throughput and latency since the higher the number of keys, the higher the number of tuples produced for all windows starting at the same timestamp. Results highlight that  $TuML_{SC}$  and WiML perform better than both MQ and  $MQ_{LF}$ ,

whose throughput decreases linearly with the increasing number of inputs.

With respect to order-sensitive functions, we compare queries k-OS (Fig. 13) and K-OS (Fig. 14). The query k-OS uses the EC dataset (243 distinct keys) while the query K-OS uses the SC dataset (40,000 distinct keys). The upper part of each figure presents the throughput while the bottom part presents the latency (in logarithmic scale). For 20 inputs, TuML<sub>SC</sub> provides the highest throughput (2.9 times better than MQ's for query k-OS) and the lowest latency.



Figure 17: Varying number of keys - evaluation.

With respect to order-insensitive functions, we compare queries k-OI (Fig. 15) and K-OI (Fig. 16). As for order-sensitive functions, both MQ's and MQ<sub>LF</sub>'s throughput decreases for increasing number of inputs. On the other hand, WiML throughput increases accordingly to the number of inputs, achieving a maximum throughput of 2.6 million t/s and 430,000 t/s, respectively. MQ's and MQ<sub>LF</sub>'s latencies increase with the number of inputs while WiML's one remains approximately constant.

**Varying number of windows to which tuples contribute** The rationale for this experiment is that the higher the number of windows to which each input tuple contributes, the higher the duration of the *Update* operation. Also in this case, our enhanced implementations outperform both MQ and  $MQ_{LF}$ . An increasing number of windows to which tuples contribute results in an overall throughput breakdown and latency increase.

We first focus on order-sensitive functions with queries w-OS (Fig. 18) and W-OS (Figure 19). For query w-OS, each tuple contributes to 10 windows. For query W-OS, each tuple contributes to 20 windows.

With respect to order-insensitive functions, we compare queries w-OI (Fig. 20) and W-OI (Fig. 21). Also in this case, WiML outperforms MQ and  $MQ_{LF}$ . For queries w-OI and W-OI, WiML's maximum throughput is of approximately 2.4 and 1.4 million t/s, 19 and 16 times better than MQ, respectively.

**Varying function cost** In this set of experiments, we study how throughput and latency evolve with respect to different function costs. We expect the throughput to decrease and the latency to increase accordingly to the increasing cost of the aggregation function. As observed before,  $TuML_{SC}$  performs better than multi-queue



Figure 18: w-OS Figure 19: W-OS Figure 20: w-OI Figure 21: W-OI

Figure 22: Varying number of windows to which tuples contribute - evaluation.

implementations, although this improvement becomes smaller when running very expensive aggregate functions. The reason for this smaller improvement is due to the dominance of the heavy aggregate function computations over other computations performed by the operator (e.g., the sorting ones) given that both  $TuML_{SC}$  and multi-queue implementations define a single thread dedicated to the former. WiML outperforms multi-queue implementations both in terms of throughput and latency independently of the aggregate function cost (in this case, despite relying on the same overall number of threads, the expensive aggregate function is run in parallel by each  $I_t$  thread).

With respect to order-sensitive functions, we compare queries f-OS (Fig. 23) and F-OS (Fig. 24). When increasing the function cost (query F-OS),  $TuML_{SC}$ 's throughput and latency become really close to MQ's and MQ<sub>LF</sub>'s ones. Throughput



Figure 27: Varying function cost - evaluation.

and latency evolution for order-insensitive functions are evaluated for queries f-OI (25) and F-OI (26). For both experiments, WiML achieves a throughput of approximately 615,000 t/s, 9 times better than MQ.

#### 7.4 TuML<sub>MC</sub> scalability evaluation

In this section, we focus on the scalability of the  $\text{TuML}_{MC}$  implementation. We execute all the previous queries for order-sensitive functions using the  $\text{TuML}_{MC}$  implementation for an increasing number of threads (up to 12, the physical number of cores of the machine used in the evaluation). For each query, we present how the throughput and the latency evolve when considering 5 and 20 inputs streams.



Figure 34: TuML<sub>MC</sub> scaling for increasing number of threads.

Figures 28 and 31 present the throughput and latency evolution for queries k-OS and K-OS. It can be noticed that K-OS scales better than k-OS for an increasing number of threads due to the increase of  $O_t$  operations' duration caused by the higher number of keys. Figures 29 and 32 present the throughput and latency evolution for queries w-OS and W-OS. In this experiment, throughput and latency behave similarly despite the increased duration of  $O_t$  operations. This is because the increased  $O_t$  operations' duration is actually caused by an higher number of windows updated by each tuple, resulting in an higher contention in the underlying W-Hive. Finally, Figures 30 and 33 present the throughput and latency evolution for queries f-OS and F-OS.

#### 7.5 Summary of results

Comparing the implementations that rely on one  $I_t$  thread per input and a single output thread  $O_t$ , both TuML<sub>SC</sub> and WiML perform better than MQ and MQ<sub>LF</sub>, enabling coping with streams of higher speed. The improvement enabled by TuML<sub>SC</sub> is more sensitive to the aggregate parameters than WiML, which clearly outperforms MQ and MQ<sub>LF</sub>. When increasing the number of processing threads, TuML<sub>MC</sub>'s performance increases both in terms of throughput and latency. Moreover, its scaling does not degrade when increasing the number of threads above the number for which the highest rate is achieved.

# 8 Related Work

Parallel execution of data streaming operators has been addressed mainly by means of partitioned parallelism [10, 3], where multiple instances of an operator are assigned to distinct partitions of a given stream. The way tuples are routed to in-

stances (round-robin, hash-based [10] or pane-based [3]) depends on the operator's semantics. It should be noticed that partitioned parallelism is orthogonal to our parallelization technique since we focus on the performance improvement of individual instances of an operator. The work presented in [23] discusses a multithreaded elastic streaming protocol that adjusts the number of processing threads depending on the system load. Similarly to our  $TuML_{MC}$  implementation, the protocol defines a single *work queue* from which multiple *worker threads* consume tuples. Nevertheless, that does not take into account sorting of input tuples, which is one of the key challenges addressed in our work. Moreover, the authors do not discuss improvements enabled by concurrent data structures in the multithreaded environment. With respect to parallel data streaming in the context of multi-core CPUs and GPUs, [24] present a parallel implementation of the aggregate operator and study how it performs on distinct parallel architectures. The aggregate model discussed by the authors differs from ours since windows are tuple-based and the overall number of distinct group-by values is known before-hand and does not vary over time. Moreover, no discussion is provided about deterministic processing in the context of multiple input sources. Parallel processing in multi-core CPUs and GPUs is also discussed by [5] but, differently from us, the authors focus on pattern detection rather than data aggregation and rely on automata-based incremental processing.

As discussed in section 2, one of the challenges in providing deterministic processing is the merging of multiple timestamp sorted input streams. This has been discussed in the context of parallel-distributed SPEs [27, 10] and replicabased fault tolerance protocols for data streaming [2]. Existing approaches for streaming aggregation rely on separated input queues (similar to the MQ protocol). As shown in our evaluation, this merging is not efficient and implementations such as TuML<sub>SC</sub> can drastically improve the overall operator's performance.

# 9 Conclusions and Future Work

Providing the appropriate data structures that best fit the needs of a concurrent application is a key research issue, as emphasized by [25, 19] and also seen in examples such as [8, 30]. In this paper, we study data structures as articulation points in the context of streaming aggregation, analysing the concurrency needs and proposing methods to meet them. We propose proper data structures for managing tuples and windows (T-Gate and W-Hive). Their operations and their lock-free implementations enable better interleaving and hence improve the balancing and the parallelism of the aggregate operator's processing stages. As shown in the extensive evaluation based on real-world datasets, our enhanced aggregate implementations outperform existing ones both in terms of throughput and latency, and are able to handle heavier streams, increasing the processing capacity up to one order of magnitude.

These results and the analysis of the role of data structures as articulation points

to facilitate concurrency and balancing of the work among streaming operators, open up new venues in the broader context of data streaming, including the enhancement of other operators and the enhancement of SPEs architectures.

# References

- [1] Daniel J. Abadi, Don Carney, Ugur Çetintemel, Mitch Cherniack, Christian Convey, Sangdon Lee, Michael Stonebraker, Nesime Tatbul, and Stan Zdonik. Aurora: A new model and architecture for data stream management. *The VLDB Journal*, 12(2):120–139, August 2003.
- [2] Magdalena Balazinska, Hari Balakrishnan, Samuel R Madden, and Michael Stonebraker. Fault-tolerance in the Borealis distributed stream processing system. ACM Transactions on Database Systems (TODS), 33(1):3, 2008.
- [3] Cagri Balkesen, Nesime Tatbul, and M. Tamer Özsu. Adaptive input admission and management for parallel stream processing. In *Proceedings of the 7th ACM international conference on Distributed event-based systems*, DEBS '13, pages 15–26, New York, NY, USA, 2013. ACM.
- [4] Daniel Cederman, Vincenzo Gulisano, Yiannis Nikolakopoulos, Marina Papatriantafilou, and Philippas Tsigas. Brief announcement: concurrent data structures for efficient streaming aggregation. In *Proceedings of the 26th* ACM Symposium on Parallelism in Algorithms and Architectures, pages 76– 78. ACM, 2014.
- [5] Gianpaolo Cugola and Alessandro Margara. Low latency complex event processing on parallel hardware. *J. of Parallel and Distributed Computing*, 72(2):205–218, 2012.
- [6] Alin Dobra, Minos Garofalakis, Johannes Gehrke, and Rajeev Rastogi. Processing complex aggregate queries over data streams. In *Proceedings of the 2002 ACM SIGMOD International Conference on Management of Data*, pages 61–72, New York, NY, USA, 2002. ACM.
- [7] Buğra Gedik, Rajesh R Bordawekar, and S Yu Philip. Celljoin: a parallel stream join operator for the cell processor. *The VLDB Journal*, 18(2):501– 519, 2009.
- [8] Anders Gidenstam, Marina Papatriantafilou, and Philippas Tsigas. Nbmalloc: Allocating memory in a lock-free manner. *Algorithmica*, 58(2):304–338, 2010.
- [9] Shenoda Guirguis, Mohamed A Sharaf, Panos K Chrysanthis, and Alexandros Labrinidis. Three-level processing of multiple aggregate continuous queries. In *Data Engineering (ICDE), 2012 IEEE 28th International Conference on*, pages 929–940. IEEE, 2012.
- [10] Vincenzo Gulisano, Ricardo Jimenez-Peris, Marta Patino-Martinez, Claudio Soriente, and Patrick Valduriez. Streamcloud: An elastic and scalable data streaming system. *IEEE Transactions on Parallel and Distributed Systems*, 23(12):2351–2365, 2012.

- [11] Vehbi C Gungor, Dilan Sahin, Taskin Kocak, Salih Ergut, Concettina Buccella, Carlo Cecati, and Gerhard P Hancke. Smart grid technologies: communication technologies and standards. *Industrial informatics, IEEE transactions on*, 7(4):529–539, 2011.
- [12] Maurice Herlihy and Nir Shavit. *The Art of Multiprocessor Programming, Revised Reprint*. Elsevier, 2012.
- [13] Maurice Herlihy, Nir Shavit, and Moran Tzafrir. Hopscotch Hashing. In Proceedings of the 22Nd International Symposium on Distributed Computing, DISC '08, pages 350–364, Berlin, Heidelberg, 2008. Springer-Verlag.
- [14] Maurice P Herlihy and Jeannette M Wing. Linearizability: A correctness condition for concurrent objects. ACM Transactions on Programming Languages and Systems (TOPLAS), 12(3):463–492, 1990.
- [15] Xiaozhou Li, David G. Andersen, Michael Kaminsky, and Michael J. Freedman. Algorithmic Improvements for Fast Concurrent Cuckoo Hashing. In *Proceedings of the Ninth European Conference on Computer Systems*, EuroSys '14, pages 27:1–27:14, New York, NY, USA, 2014. ACM.
- [16] Simon Loesing, Martin Hentschel, Tim Kraska, and Donald Kossmann. Stormy: an elastic and highly available streaming service in the cloud. In *Proceedings of the 2012 Joint EDBT/ICDT Workshops*, pages 55–60, New York, NY, USA, 2012. ACM.
- [17] Maged M. Michael. High performance dynamic lock-free hash tables and list-based sets. In *Proceedings of the fourteenth annual ACM symposium on Parallel algorithms and architectures*, pages 73–82, New York, NY, USA, 2002. ACM.
- [18] Maged M. Michael. Hazard pointers: Safe memory reclamation for lock-free objects. *IEEE Trans. Parallel Distrib. Syst.*, 15(6):491504, June 2004.
- [19] Maged M. Michael. The balancing act of choosing nonblocking features. *Commun. ACM*, 56(9):46–53, 2013.
- [20] Maged M. Michael and Michael L. Scott. Simple, fast, and practical nonblocking and blocking concurrent queue algorithms. In *Proceedings of the fifteenth annual ACM symposium on Principles of distributed computing*, pages 267–275, New York, NY, USA, 1996. ACM.
- [21] Nhan Nguyen and Philippas Tsigas. Lock-Free Cuckoo Hashing. In Proceedings of the 2014 IEEE 34th International Conference on Distributed Computing Systems, ICDCS '14, pages 627–636, Washington, DC, USA, 2014. IEEE Computer Society.
- [22] William Pugh. Skip Lists: A Probabilistic Alternative to Balanced Trees. *Commun. ACM*, 33(6):668–676, June 1990.
- [23] Scott Schneider, Henrique Andrade, Bugra Gedik, Alain Biem, and Kun-Lung Wu. Elastic scaling of data parallel operators in stream processing. In *IEEE International Symposium on Parallel & Distributed Processing*, pages 1–12, Washington, DC, USA, 2009. IEEE Computer Society.
- [24] Scott Schneider, Henrique Andrade, Bura Gedik, Kun-Lung Wu, and Dimitrios S Nikolopoulos. Evaluation of streaming aggregation on parallel hard-

ware architectures. In *Proceedings of the Fourth ACM International Conference on Distributed Event-Based Systems*, pages 248–257, New York, NY, USA, 2010. ACM.

- [25] Nir Shavit. Data structures in the multicore age. *Commun. ACM*, 54(3):76–84, March 2011.
- [26] Julian Shun and Guy E. Blelloch. Phase-concurrent Hash Tables for Determinism. In Proceedings of the 26th ACM Symposium on Parallelism in Algorithms and Architectures, SPAA '14, pages 96–107, New York, NY, USA, 2014. ACM.
- [27] Utkarsh Srivastava and Jennifer Widom. Flexible time management in data stream systems. In *Proceedings of the twenty-third ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems*, pages 263–274, New York, NY, USA, 2004. ACM.
- [28] Michael Stonebraker, Uur Çetintemel, and Stan Zdonik. The 8 requirements of real-time stream processing. *ACM SIGMOD Record*, 34(4):42–47, 2005.
- [29] Håkan Sundell and Philippas Tsigas. Fast and lock-free concurrent priority queues for multi-thread systems. *Journal of Parallel and Distributed Computing*, 65(5):609–627, 2005.
- [30] Martin Wimmer, Francesco Versaci, Jesper Larsson Träff, Daniel Cederman, and Philippas Tsigas. Data structures for task-based priority scheduling. In Proceedings of the 19th ACM SIGPLAN symposium on Principles and practice of parallel programming, pages 379–380. ACM, 2014.