

Continuous Outlier Mining of Streaming Data in Flink

Theodoros Toliopoulos^a, Anastasios Gounaris^a, Kostas Tsihclas^a, Apostolos Papadopoulos^a, Sandra Sampaio^b

^a*Department of Informatics, Aristotle University of Thessaloniki, Greece*
{toliop,gounaria,tsichlas,papadopo}@csd.auth.gr

^b*The University of Manchester, UK*
s.sampaio@manchester.ac.uk

Abstract

In this work, we focus on *distance-based outliers* in a metric space, where the status of an entity as to whether it is an outlier is based on the number of other entities in its neighborhood. In recent years, several solutions have tackled the problem of distance-based outliers in data streams, where outliers must be mined continuously as new elements become available. An interesting research problem is to combine the streaming environment with massively parallel systems to provide scalable stream-based algorithms. However, none of the previously proposed techniques refer to a massively parallel setting. Our proposal fills this gap and investigates the challenges in transferring state-of-the-art techniques to Apache Flink, a modern platform for intensive streaming analytics. We thoroughly present the technical challenges encountered and the alternatives that may be applied, of which a micro-clustering-based one is the most efficient. We show speed-ups of up to 2.27 times over advanced non-parallel solutions, by using just an ordinary four-core machine and a real-world dataset. When moving to a three-machine cluster, due to less contention, we manage to achieve both better scalability in terms of the window slide size and the data dimensionality, and even higher speed-ups, e.g., by a factor of more than 11X. Overall, our results demonstrate that outlier mining can be achieved in an efficient and scalable manner. The resulting techniques have been made publicly available as open-source software.

Keywords: distance-based outlier detection, Flink, data streams

1. Introduction

Outlier analysis forms a key mechanism in modern data science and analytics [1], aiming to detect objects that, as defined in [2], appear to be inconsistent with the remainder of the objects in the same dataset. They are used in a variety of applications, such as fraud detection, spam detection and medical diagnosis, just to name a few. One of the most commonly used definitions for an outlier is the distance-based one [3], where an object is considered an outlier if it does not have more than k neighbors within a distance up to R . This has been proven to be a useful and valid definition of outlierness in practice [4]. Continuous outlier detection in data streams deals with the problem of keeping an updated list of all outliers after each new object arrives and/or expires. Apparently, when data grows large, this becomes a challenging task, since applying even an one-pass algorithm to all active data is prohibitively expensive. To improve efficiency and scalability, the main target of this work is to propose massively parallel solutions for continuous outlier detection in data streams.

Briefly, the relevant state-of-the-art falls into two categories. The first category contains efficient non-parallel solutions for streaming outlier detection, e.g., the works described in [5, 6, 7, 8]. The second category contains parallel solutions for outlier detection, where, to date, there is a single proposal that assumes modern distributed computing platforms, such as MapReduce [9]; nevertheless, this solution does not deal with streaming data. The novelty of our proposal, to the best of our knowledge, is that it proposes, for the first time, the combination of both massive parallelism and continuous distance-based outlier detection in a streaming setting. Orthogonally, outlier detection techniques are typically classified as either exact or approximate. We target exact solutions in this work, and more specifically, we aim to transfer the main features of the state-of-the-art non-parallel exact solutions, which are investigated in [10], to a parallel setting in an efficient manner.

Devising efficient parallel solutions for this problem involves addressing a series of important issues. First, outlier detection algorithms in data streams involve windows that cannot be partitioned into non-overlapping partitions, among which no communication is required. Second, low latency is of high significance in order to deliver results in a timely manner. Third, state information needs to be kept between window slides in order to avoid unnecessary recomputations. We provide solutions to the above issues through the application of key ideas obtained from existing non-parallel techniques, such as

the ones described in [8, 5] to the Flink¹ platform.

This work² aspires to become a reference point for all future work on streaming outlier detection in massively parallel settings and its main contributions are summarized as follows:

- (i) We explore a series of implementation alternatives, differing in the algorithmic features they encompass, the data structures they employ and in the way data is partitioned. More specifically, compared to the earlier work described in [11], we present a novel Vantage-Point (VP) tree-based value partitioning technique that is suitable for arbitrary metric spaces compared to grid-based one, as well as the encapsulation of the notion of window slicing, which yields the main alternative to the microcluster-based solution, and thorough investigation of the alternatives with regards to the underlying data structures.
- (ii) We provide thorough experimental evaluation results. Compared to the earlier work described in [11] that performed evaluation on a single multi-core machine solely, we additionally investigate the behavior of continuous outlier detection algorithms running on a cluster consisting of multiple multi-core physical machines, and we assess the impact of the new features mentioned above, i.e., the new data partitioning, the additional technique of window slicing and the data structure choices.
- (iii) We offer the source code as an open-source library.³

In summary, we show that our best performing alternative, when tested on a real-world dataset, due to the micro-clustering approach it follows, is more than 4.3X faster than other advanced alternatives when running on a 3-machine cluster. Compared to the state-of-the-art non-parallel solutions, the speed-up using a small cluster is up to 11.12X, and up to 2.27X when employing a commodity four-core machine. Due to resource contention, we observe that, on a single machine, the scalability is not good regarding the data dimensionality and size of the window slide. Both problems are tackled

¹<https://flink.apache.org/>

²This is an extension to our paper titled “Parallel Continuous Outlier Mining in Streaming Data” presented in 5th IEEE International Conference on Data Science and Advanced Analytics (DSAA), 2018 [11].

³The code repository may be accessed at <https://github.com/tatoliop/parallel-streaming-outlier-detection>

in a cluster environment, where naive parallel solutions are up to 510 times slower than our proposal.

The remainder of this work is structured as follows. Section 2 contains background material on parallel streaming platforms and outlier detection algorithms. Section 3 introduces our first parallel solution, which is extended in Sections 4 and 5. Performance evaluation results are offered in Section 6. We close this work with a discussion of related work and issues relating to extensions to our techniques in Sections 7 and 8, respectively.

2. Fundamental Concepts and Background

The purpose of this section is to make the article as self-contained as possible. We split background material into two parts, referring to the main massively parallel platform alternatives for streaming data and the distance-based outlier algorithms that inspired our solutions, respectively.

2.1. Parallel Frameworks and Streaming Semantics

In this section, the massively parallel platforms for streaming data that have been considered are presented in order to explain the reasons why Flink has been selected. In addition, the main streaming semantics of windows and timestamps are explained.

2.1.1. Parallel Frameworks for Streaming Applications

Three of the main parallel streaming platforms were examined, before choosing the most suitable one for the purposes of this work, namely: (i) Storm⁴, (ii) Spark⁵ and (iii) Flink. In summary, Storm provides low latency, similar to Flink, but it incorporates the *at-least-once* semantics, which allows duplicates to pass through the process in case of failures. Spark, like Storm Trident, provides *exactly-once* semantics with the use of micro-batches. The main drawback of micro-batches is the higher latency compared to the continuous processing model of Storm and Flink. Finally, Flink combines the advantages of Storm and Spark, namely low latency regarding the continuous record processing, coupled with the *exactly-once* semantics. In addition, Flink naturally supports both time- and count-based windows, which is not

⁴<https://storm.apache.org/>

⁵<https://spark.apache.org/>



Figure 1: Tumbling Windows

the case for Spark, since micro-batches essentially correspond to time-based sliding windows (see discussion below).

2.1.2. Streaming Semantics

A continuous stream is an infinite sequence of data points. Each data point o is annotated with its arrival time, $o.t$. A common approach for streaming applications is to adopt the notion of *window*, which refers to the most recent data items. Windows are typically small enough so that they can be stored in main memory, either of a single machine or of a parallel cluster. Windowing essentially splits the data stream into either overlapping (*sliding windows*) or non-overlapping (*tumbling windows*) finite sets of data points. Orthogonally, the splitting can be based either on the time of arrival of the data points (*time-based windows*) or on the number of data points (*count-based windows*). In the former case, the window size W is measured in time units, while in the latter case the size corresponds to the number of the most recent data items held. Figure 1 shows a stream discretized in three windows based on time. The windows are non-overlapping with $W = 2$ time units.

In this work, we focus on sliding windows, which generalize the tumbling ones, and our techniques support both time- and count-based windows. Therefore, without any loss of generality, whenever we use the term window, we will be referring to a sliding time-based window. In sliding windows, the magnitude of each slide is denoted as S . Every time the window moves by S , $W.start$ and $W.end$ are increased by S as well. For example, in Figure 1, $S=2$ time units, and in Figure 2, $S=1$ time unit. In each slide, some points may expire, i.e., they are dropped, while new points are included in the current window. Table 1 summarizes the notation used throughout the paper.




Figure 2: Sliding Windows

2.2. Problem Definition and Non-Parallel Solutions

The problem of continuous distance-based outlier detection is defined formally as follows.

Definition 1. *Given a set of objects \mathbb{O} and the threshold parameters R and k , for each window slide S , report all the objects o_i for which the number of neighbors $o_i.nn < k$, i.e., the number of objects o_j , $j \neq i$ for which $dist(o_i, o_j) \leq R$ is less than k .*

Note that the definition is common for both time- and count-based windows, since the window semantics impact solely on the window contents. The only difference to the static case in [1, 3] is the need to report outliers for each slide instead of once for the complete dataset. The main challenges stem from the fact that all active objects need to be continuously assessed during their lifetime, since an object may change its status as many times as the number of slides during which it remains within the window.

We exclusively focus on exact solutions. There are several exact algorithms for continuous outlier detection in data streams, such as exact-Storm [8], Abstract-C [6], LUE [12, 5], DUE [12, 5], COD [12, 5], MCOD [12, 5] and Thresh_LEAP [7]⁶. All these algorithms assume a centralized environment. In a recent impartial comparison presented in [10], these algorithms were compared in terms of their memory consumption and CPU time. The comparisons were made using four data-sets with varying dimensionality and settings of window length W , window slide S , number of neighbors k

⁶Reference implementations are provided in the MOA tool [13] <http://moa.cs.waikato.ac.nz/> and <https://infolab.usc.edu/Luan/Outlier/>

and neighbor range R . This study showed that MCODE is superior across multiple datasets in most stream settings. In addition, Thresh_LEAP and MCODE displayed the lowest memory consumption and CPU times, while exact-Storm, Abstract-C and DUE are the slowest and most memory consuming algorithms. Based on these findings, MCODE has served as our preferred basis for our solution to the problem of parallelization of continuous distance-based outlier detection algorithms. However, in our investigation, we start with a simpler and easier to parallelize algorithm, namely exact-Storm, which contains several key elements in common with MCODE, such as index structures for range queries and safe inliers, and so it represents a preliminary step towards our proposed solution. We also use the time-slicing notion of the Thresh_LEAP algorithm as one of the extensions. The key details of these three algorithms are discussed in the following.

2.2.1. Exact-Storm

Two of the key operations in distance-based outlier detection are the distance computation between objects and the continuous examination of the neighborhood of each object. To avoid a quadratic number of comparisons in each slide, appropriate indices are required. To this end, exact-Storm uses a data structure called ISB to store the data points in nodes. A node is a record containing the data point o , the arrival time of the point $o.t$, the number of succeeding neighbors $o.count_after$ and a list $o.nn_before$ of size k containing the arrival time of the preceding neighbors of o (i.e., each node contains a different data stream object along with some metadata). This data structure is a pivot-based index that provides support for fast range query search in any metric space. The range query, given a data point o and the range R , returns the nodes in the ISB whose distance to o is less than or equal to R .

A sketch of the algorithm steps in each slide is as follows. For each new data point o , a node is created as described above. Then, a range query is issued on the ISB structure to find the neighbors o' of the new node. The result of the range query is used to initialize the values of the new node's $o.count_after$ and $o.nn_before$. If the size of $o.nn_before$ is more than k then the oldest timestamps are removed. For each o' , the value of $o'.count_after$ is increased by 1. Finally the new node is inserted into the ISB. When a data point o expires, meaning that $o.t$ is lower than the window's starting timestamp $W.start$, it is removed from the ISB. Its timestamp, however, is not removed from other nodes' list of preceding neighbors to mitigate

Symbol	Short description
W	The size of the stream window
S	The slide of the stream
$W.start$	The starting timestamp of the window
$W.end$	The ending timestamp of the window
\mathbb{O}	The set of data objects (or points) in the stream
$o_i \in \mathbb{O}$	The i^{th} data object in the stream
$o.id$	The object identifier o (either i for o_i or any other identifier)
$o.value$	The value of o
$o.t$	The arrival time of o
$o.count_after$	The number of succeeding neighbors of o
$o.nn_before$	A list with the arrival time of the preceding neighbors of o
$o.nn$	The count of neighbors of o
\mathbb{PO}	A list containing data points that are potentially outliers
R	The distance threshold in the outlier definition
k	The neighbor count threshold in the outlier definition
$dist(o_i, o_j)$	The distance function between objects o_i and o_j
P	set of Flink partitions, each handled by a separate Flink node

Table 1: Frequently used symbols and interpretation

overheads.

The above steps are applied in each slide. After they have been completed, ISB is scanned for outliers. If a node’s sum of $count_after$ and the size of nn_before , whose timestamp is within the window borders, is lower than k , then the node is an outlier. An optimization that avoids checking all objects in each slide is through the notion of *safe inliers*: if a node’s $count_after$ is more than k , then this node is a safe inlier and does not need to be checked again in future scans as it is guaranteed to have at least k neighbors for the remainder of its lifetime.

In our parallel solution, we adopt both a structure for fast range queries (using, however, an M-tree instead of an ISB) and the notion of safe inliers.

2.2.2. *Thresh_LEAP*

Thresh_LEAP uses the broader notion of time-slicing [14] in order to decrease the number of range queries needed to establish the status of each data point. Each window slide has a separate index structure where the data points are held. This design allows the use of the *minimal probing principle*. Following this principle, each data point searches the succeeding slides' indices for neighbors and then continues with the preceding slides in reverse chronological order. This search stops when the data point has found k neighbors.

Each data point o has a list $o.evill$, where it stores the number of neighbors in each preceding slide. Each slide has a *trigger list*, where the data points with neighbors within the slide are stored. A sketch of the main algorithm steps is as follows. For each new data point o , a range query is issued on the same slide structure. If the search returns less than k neighbors, range queries are issued for each preceding slide in reverse chronological order until k neighbors are found or all slides are searched. The data point o has its metadata updated with the number of neighbors for each slide and is inserted into the *trigger lists* of all the slides where it has at least one neighbor. When a slide expires, all the data points in its *trigger list* are re-processed. Each data point o in the *trigger list* issues a range query for the succeeding slides that have not been searched so far. There is no need to issue range queries on preceding slides as they have been searched in the previous steps of the algorithm. After the previous steps have been completed, the status of each active data point o is evaluated by summing up the content of the $o.evill$ list with its succeeding neighbors, specified in $o.count_after$.

In our solution, we use the time-slicing notion with the *minimal probing principle* in order to decrease the number of range queries per data point.

2.2.3. *MCOD*

The main motivation behind MCODE relates to the fact that range queries are less expensive than brute-force (all-pairs) distance computations but are, nonetheless, still expensive. MCODE (standing for Micro-cluster-based Continuous Outlier Detection) mitigates the need for range queries by drastically reducing the number of data points that need to be addressed during a range query through the creation of micro-clusters and the assignment of data points to them. A micro-cluster has at least $k + 1$ data points all of which are neighbors to each other. Its center can be a data point or just a point in the metric space and has a radius of $\frac{R}{2}$, implying that the maximum

distance between any two objects in the micro-cluster is at most R . Each data point in any micro-cluster is an inlier and does not need to be checked in outlier detection queries. However, a data point that does not belong to a micro-cluster can be either an inlier or an outlier. Such objects are stored in a list $\mathbb{PO} \subseteq \mathbb{O}$.

On average, MCODE stores less metadata for each object than exact-Storm. More specifically, for each o in a micro-cluster, it stores the identifier of its cluster. For each o in \mathbb{PO} , it stores the $o.count_after$ and the expiration time of the k most recent preceding neighbors. MCODE also uses an event queue to store unsafe inliers that are not in any cluster. This event queue is a specific priority queue that keeps the time point at which a non-safe inlier should be re-checked.

A sketch of the algorithm steps is as follows. For each new data point o , if o is within $\frac{R}{2}$ range of a micro-cluster, it is added to it; if there are multiple such micro-clusters, the closest one is picked. Otherwise, if it has at least k neighbors in \mathbb{PO} within a distance of $\frac{R}{2}$, it becomes the center of a new micro-cluster. If none of the above conditions are met, o is added to \mathbb{PO} and possibly to the event queue, if it is not an outlier. At each slide, all the previous non-expired outliers are checked along with the inliers for which the check time has arrived (with the help of the event queue). When a data point o expires, it is removed from the micro-cluster or \mathbb{PO} and the event queue updates the unsafe inliers. If o is removed from a micro-cluster and the points remaining in that micro-cluster are less than $k + 1$, then the cluster is destroyed and each data point of the cluster is processed as a new data point, without however updating their neighbors.

In our final parallel solution, we also adopt the notion of micro-clusters.

3. Simple Solutions

In this section, first, we introduce a baseline approach, which broadly corresponds to a single-partition (i.e., non-parallel) implementation of exact-Storm in Flink without support for range queries. This allows us to explain the engineering approach chosen and to be able to assess the efficiency of the implementation of the parallel streaming solutions for distance-based outlier detection proposed in this work. Then, we proceed to the parallelization of the baseline, where the windows are physically partitioned across multiple Flink nodes and we are thus forced to employ the notion of a *meta-window* for engineering purposes.

3.1. A baseline non-parallel approach in Flink

In this approach, we use two main components, as follows: (i) a *stream handler*; and (ii) a *window processor*. The stream handler applies a `map` function on each stream object and sends it to the window processor. The window processor runs an outlier detection algorithm in each slide.

In the `map` function of the stream handler, each data point is initialized with a null `o.count_after` count and an empty list `o.nn_before`. These records are sent to the single-partitioned window. The window has its own state in which it stores the records and is persistent across slides. In other words, changes made to a data point’s metadata in a slide are kept throughout the data point’s lifetime. Overall, the contents of a window at any point in time include all the active points along with their metadata, i.e., `o.id`, `o.value`, `o.t`, `o.count_after` and `o.nn_before`.

The outlier detection algorithm contains two steps. In the first step, the update of each data point’s metadata is performed. In the case of this algorithm, this is done only for the new arrivals. For every such data point, o , it finds the neighbors, o' in range R , checking all the points in the window; by default, the range is calculated based on the euclidean distance, but any type of metric distance can be employed. If o' is in the same slide as o , then `o.count_after` is increased; otherwise the timestamp $o'.t$ is added to the list `o.nn_before`. For each o' the value of `o'.count_after` is increased by 1. The remaining metadata associated with the older data points, i.e. the `o'.nn_before` values, have already been computed when these objects were inserted in the window, and so, do not need to be recalculated. The second step of the algorithm is to detect the outliers. For each data point in the current window, the algorithm computes the total number of neighbors. This is done by summing the `o.count_after` and the size of the list `o.nn_before`, taking into account only the values t where $t \geq W.start$, as explained previously.

Both time-based and count-based windows are naturally supported in Flink. In this work, we mainly focus on time-based windows, but it is worth pointing out that, even without explicit support of count-based windows, it is straightforward to emulate them through artificially tweaking the initial timestamps, so that a fixed number of objects arrive and expire in each slide, and thus the number of alive objects remains constant during the stream processing.

To gain insights into the performance of the solution, we evaluate this baseline approach using the Stock dataset from [10]. It is a one-dimensional




Figure 3: Average processing time per slide (left) and input stream consumption rate (right) for the baseline approach.

dataset with 1,048,575 data points⁷. Each data point, o , is assigned a unique identifier, $o.id$, and has a numeric value of type `Double`, $o.value$. We employ a machine with an Intel i7-3770K CPU at 3.5GHz, which possesses 4 cores (8 threads) and 32GB of RAM. Figure 3 (left) shows the average processing time for each slide step considering four different values for slide magnitude, S , with $W = 10K$; S is stated as a percentage of W . Figure 3 (right) shows the corresponding input consumption rate, which is equal to the maximum streaming data throughput that the baseline approach can support; in these settings, this throughput can reach 50 objects/sec. We can see that, in general, the average processing time per new arrival increases for both small slides, where a few new points arrive, and relatively large ones, where most of the window contents are replaced.

3.2. A Naive Replicate-to-all Solution

The parallelization of the baseline approach in Section 3.1 yields a naive parallel solution, where the window is split into a set of multiple partitions, P . It is deemed naive in the sense that it does not benefit from any data structures to speed-up range queries. Notwithstanding its simplicity, the parallelization technique needs to efficiently address the challenges relating to (i) the collaboration between physical window partitions to establish whether a point, at a specific time, is an outlier or not by aggregating local statistics; and (ii) the maintenance of the window state across slides, where the state

⁷available from <https://wrds-web.wharton.upenn.edu/wrds>

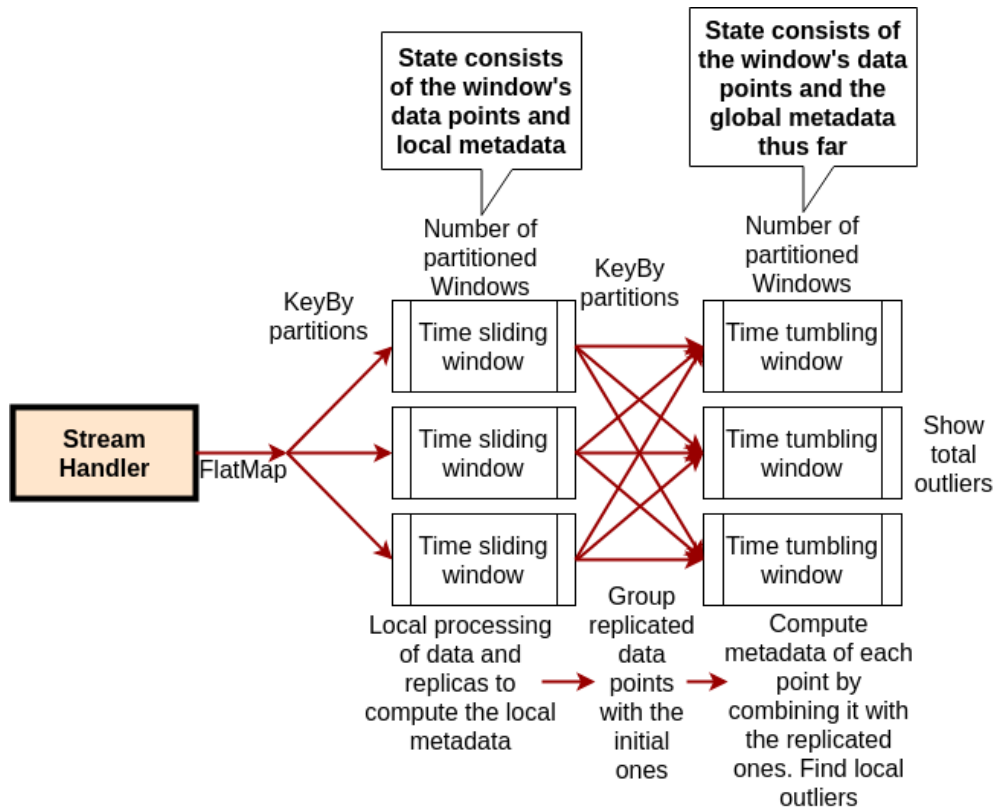


Figure 4: Main implementation rationale with random partitioning, where the parallelized sliding window is followed by a partitioned tumbling meta-window.

includes the object metadata.

The fact that the window is split into multiple partitions and a point may have neighbors belonging to any partition does not allow first to compute the local aggregates for a given point and then to compute the global aggregates into the same window. Therefore, the key idea is to split the window processor into two parts, the sliding window processor and the tumbling window processor, as shown in Figure 4. The former holds the active points in its partition, allowing some temporary replication, as discussed in the following, while the latter keeps the final metadata in each slide. These metadata also include information about detected outliers. Since the metadata evolve in each slide for both new and old points, the window state gets fully updated, and thus the tumbling window semantics apply. Essentially, the second window serves as a *meta-window*.

Algorithm 1 Naive solution

```
procedure STREAMHANDLER
  for each new object  $o$  do
    initialize record
    if there is no timestamp then
      add artificial timestamp
     $o.partition \leftarrow o.id \bmod |P|$ 
     $o.flag \leftarrow 1$ 
    for each partition  $p \in P$  do
      if ( $p == o.partition$ ) then
         $o.flag \leftarrow 0$ 
        send  $o$  to  $p$ 
         $o.flag \leftarrow 1$ 
      else
        send  $o$  to  $p$ 

procedure SLIDINGWINDOWPARTITIONPROCESSOR
  for each slide do
    evict expired objects and old objects with  $o.flag == 1$ 
    compute pairwise distances involving new objects
    update  $o.count\_after$  and  $o.nn\_before$  metadata
    for each object  $o \in \mathbb{PO}$  do
      if ( $o.count\_after \geq k$ ) then
         $\mathbb{PO} \leftarrow \mathbb{PO} \setminus o$ 
    group objects in  $\mathbb{PO}$  by  $o.partition$ 
    and send to a (new) tumbling window

procedure TUMBLINGWINDOWPARTITIONPROCESSOR
  for each slide do
    aggregate  $o.count\_after$  and  $o.nn\_before$  metadata
    for each object  $o$  do
       $o.nn\_prec \leftarrow$  prune  $o.nn\_before$ 
      if ( $o.count\_after + |o.nn\_prec| < k$ ) then
        report  $o$  as an outlier
```

The implementation details are as follows. First, we extend the object record with two new fields, namely $o.flag$ and $o.partition$. The former is a binary variable, where 0 indicates that the object should be kept to the

assigned window during its lifetime, and 1 indicates that the object is redundant and, therefore, should be evicted in the next slide, regardless of the $W.start$ value. The stream handler uses a `flatMap` for not only initializing an object’s extended record, but also to compute the object’s $o.partition$ based on the $o.id$. Then, it dispatches the object to *all* partitions, setting the $o.flag$ to 1 to all partitions different to $o.partition$. In other words, new objects are replicated across all the partitions.

The sliding window processor is responsible for comparing the distances between a) the new objects in the slide and b) the new objects and all the previous window contents, the timestamp of which comes after $W.start$. This leads to the need for updating the $o.count_after$ value for all objects and the $o.nn_before$ value for the new objects only. However, because the metadata for the new objects are local aggregates spread across all partitions, to produce the global aggregates, the updated objects are partitioned again according to $o.partition$ into a new tumbling window, but without replication this time. To save communication cost, not all objects are shuffled, but only those that are not safe inliers. The set of the safe inliers is the complement of the set of potential outliers, i.e., $\mathbb{O} \setminus \mathbb{PO}$. In each slide, the sliding window processor first creates the \mathbb{PO} set by checking whether $o.count_after$ is less than k or not. Then, the tumbling window globally aggregates the local aggregates of non-safe inliers thus producing the exact metadata needed to establish outlieriness. As in [8], the list in $o.nn_before$ may contain preceding neighbors that have expired, and so, a filter is required to identify the alive ones, termed $o.nn_prec$.

Algorithm 1 summarizes the naive approach. In conclusion, the naive solution requires minimal effort from the stream handler at the expense of high communication and computation cost, since new objects are sent to each partition and compared against the full local window contents. In terms of implementation, it requires the notion of *meta-windows* to achieve accurate results.

4. Advanced Solutions: limiting the amount of range queries and replication

The advanced solution extends the naive one in two complementary and orthogonal dimensions, namely through: (i) employing data structures to support fast range queries and processing of window elements, described in Section 4.1; and (ii) through performing value-based partitioning, which

eliminates the need to employ a meta-window and replicate all points on all nodes, described in Section 4.2. The results of (i) are two techniques called *advanced* and *advanced w/ slicing* in the evaluation; these can be combined with two partitioning techniques from (ii), a *grid* and a *tree*-based one.

4.1. Window state management

We can enhance the processing of new and old items in the partitioned window using two techniques. The first technique involves a more advanced approach to holding state in the sliding window, by storing it in a M-tree [15], against which range queries are submitted. M-trees are a part of each partition state, therefore each Flink partition has its own local tree allowing faster execution of local range queries.

An additional enhancement derived from ideas described in [7] and introduces the notion of time-slicing in the window’s state, where a window is divided into a number of slides and each slide has its own index structure, replacing the one structure associated with the whole window. This design allows the application of the *minimal probing principle*, as discussed in Section 2.2.2, having as an index structure in this case the M-tree which allows, as described above, faster range queries. This means that a new data point first issues a range query on the its own slide’s M-tree, and, if the search returns less than k neighbors, it continues issuing range queries on the previous slides’ M-trees. The list of the preceding neighbors on the metadata of the data points is replaced by a list containing the number of neighbors for each preceding slide. Finally, each slide has a *trigger list* containing all of the data points that will be affected when the slide is expired.

Note that, compared to Algorithm 1, the key difference is in the sliding window processor.

4.2. Value-based partitioning

A limitation of the solutions described thus far is that they replicate each new data point to all partitions. This is inevitable, given that the stream handler assigns points to partitions randomly and thus a new object may have neighbors in all partitions. Value-based partitioning addresses this limitation without sacrificing the accuracy of the results. Also, as will be explained below, it eliminates the need to exchange information between partitions during a slide.




Figure 5: Value-based partitioning

4.2.1. Grid-based partitioning

As a starting point, we make an additional assumption that the space is Euclidean and, therefore, it can be partitioned into grid cells. Further, we assume the availability of sample data before execution. Based on the sample data, we can extract min, max and quantile information about the value distribution in each dimension can be extracted, in order to construct the grid cell appropriately.

The rationale behind the approach for a two-dimensional case is illustrated in Figure 5. Each cell is assigned to a single Flink node and we assume that R is much smaller than the size of a grid cell side. However, an object in a cell may have neighbors in other cells as well. The key difference is that these neighbors belong to adjacent cells only; therefore the number of Flink nodes that need to be aware of the arrival of each new object is limited. More specifically, the borders of each cell are extended by a buffer zone of width equal to R . The stream handler sends a new data point (i) to the partition corresponding to its cell with flag $o.flag$ set to 0 and (ii) to all the partitions, the buffer zone of which includes the new data point with flag $o.flag$ set to 1; these partitions form the set AP in Algorithm 2. Due to the assumption that R is much smaller than the size of a grid cell side, each data object is replicated at most 4 times if the data is 2-dimensional; this is because, when it falls near to a cell corner, it may fall into the buffer

Algorithm 2 Advanced solution with value-based partitioning

```
procedure STREAMHANDLER
  for each new object  $o$  do
    initialize record
    if there is no timestamp then
      add artificial timestamp
     $o.partition \leftarrow \text{findGridCell}(o.value)$ 
     $o.flag \leftarrow 0$ 
    send  $o$  to  $o.partition$ 
     $AO \leftarrow \text{findRelevantAdjacentPartitions}(o.value)$ 
     $o.flag \leftarrow 1$ 
    for each partition  $p \in AP$  do
      send  $o$  to  $p$ 

procedure SLIDINGWINDOWPARTITIONPROCESSOR
  for each slide do
    evict expired objects from M-tree
    insert new objects in M-tree
    compute distances
    update  $o.count\_after$  and  $o.nn\_before$  metadata
    for each object  $o \in \mathbb{PO}$  do
      if ( $o.count\_after \geq k$ ) then
         $\mathbb{PO} \leftarrow \mathbb{PO} \setminus o$ 
      else
         $o.nn\_prec \leftarrow \text{prune } o.nn\_before$ 
        if ( $o.count\_after + |o.nn\_prec| < k$ ) then
          report  $o$  as an outlier
```

zone of three other adjacent cells. In the example in Figure 5, the buffer zone borders for the upper-left cell are depicted; points 1,2 and 3 are sent to the Flink node responsible for the upper-left cell along with all the other three points. For d -dimensional data, a data point can be sent to up to 2^d partitions, which grows exponentially in the number of dimensions but, by construction, is at most equal to $|P|$; i.e., the replication is never inferior to that of the naive solution.

According to the partitioning above, each partition has all the necessary information in order to establish object outlierness locally. Therefore, a single sliding window partition processor is required, which incorporates the

responsibility of the tumbling window partition processor in the previous approaches. Algorithm 2 summarizes the advanced solution with value-based partitioning (referred to as *advanced(grid)* in the experiments). Apart from making the tumbling window processing phase obsolete, another difference to the simple advanced algorithm is that objects with $o.flag = 1$ are not dropped until they expire.

In a streaming case, data is inherently volatile. However, dynamic load balancing is left for future work. As shown in Figure 5, the grid cells are not necessarily of the same size. Here, we define cell boundaries statically, taking into consideration only the value distribution per dimension, thus overlooking issues such as actual computation and communication cost per partition.

4.2.2. Tree-based partitioning

The main drawback of grid-based partitioning is that it does not support arbitrary metric spaces. To address this limitation, we can use the M-tree also to perform value-based partitioning. In particular, instead of assuming a grid partitioning, which presupposes an Euclidean space, we can partition the data based on a rather shallow level of the M-tree, as follows. First, fix a particular level ℓ of the M-tree. Then, each Flink node is assigned to a particular node (or multiple nodes) at level ℓ . Level ℓ depends on the choice of R . However, the overlap between nodes at level ℓ depends on the distribution of the data objects. Also, the tree needs to be initially populated with enough elements of the data stream so that it grows at least up to the level ℓ .

The main guideline in implementing a structure such as the M-tree for the partitioning is that the replication of the data points should not significantly exceed the replication from the grid partitioning. The M-tree splits the nodes in such a way that the area covered by each node overlaps with the other nodes. The extent of the overlaps is independent of R and is based only on the quality of the promotion and the partitioning functions used in the tree constructions. In our initial experiments, the partitioning based on a M-tree yielded a replication rate for each data point close to 1.5 to 2 times more than the corresponding rate in the grid partitioning.

Another type of tree that provides fast range queries in a metric space is the Vantage-Point tree (VP-tree) [16]. VP-tree works by selecting a vantage point from the dataset and a corresponding threshold for each node. After this selection, the data points are split into two children nodes, the ones near the vantage point and in radius of the threshold, and the ones outside the

threshold. This means that the tree is binary and there is no overlapping involved between the children of the same node; each tree node comprises just two variables, the vantage point and the threshold. Just as the partitioning with the M-tree, in the VP tree solution, each Flink node is assigned to one or more tree nodes. This implies that we do not need to keep the complete tree after we build it based on a data sample. In other words, we can safely prune the tree nodes beyond a specific level; because the tree is binary, the level can be selected based on the number of Flink partitions, e.g. the level for 16 Flink partitions should be at least 4 (not counting the root level) so that there are at least as many tree nodes as Flink nodes. In our initial experiments, we prune as much as possible and this solution exhibited much less replication than the M-tree one, with a replication rate closer to that of the grid partitioning. This solution is referred to as *advanced(tree)* in the experiments. Similarly to M-trees, VP-trees are sensitive to their construction parameters, too.

Overall, the advanced solution trades i) additional workload on the stream handler and ii) increased memory requirements on the Flink nodes for i) less communication cost between both the stream handler and the Flink nodes, and the nodes themselves, and ii) less computation cost per node. The implementation challenges are mostly related to how the value-based partitioner on the stream handler is efficiently and effectively constructed, while, in the experiments, we assess the impact of the structures for state management.

5. pMCOD: employing micro-clusters

As described in Section 2.2.3, the *MCOD* algorithm is the best non-parallel solution for the distance-based outlier detection problem according to the results in [10]. The *MCOD* algorithm has three distinct features that could be independently implemented in the final parallel solution termed as *pMCOD*:

- the micro-clusters;
- the event-queue; and
- restricted usage of the M-tree for the range queries.

Each of these features is examined in-depth for its usefulness to the final solution hereby, and the accompanying experiments are in Section 6.1.2.

Algorithm 3 pMCOd

procedure STREAMHANDLER

Same as Algorithm 2

procedure SLIDINGWINDOWPARTITIONPROCESSOR

for each slide **do**

evict expired objects from M-tree

if a micro-cluster dissolves **then**

treat all points as new w/o updating their neighbors

insert new objects in M-tree

for each new object o' with $o'.flag = 0$ **do**

if o' belongs to micro-cluster **then**

for each object $o'' \in \mathbb{PO}$ **do**

update o'' metadata due to o'

else

$\mathbb{PO} \leftarrow \mathbb{PO} \cup o'$

compute pairwise distances involving o'

check if a new micro-cluster can be formed

for each object $o \in \mathbb{PO}$ **do**

if $(o.count_after \geq k)$ **then**

$\mathbb{PO} \leftarrow \mathbb{PO} \setminus o$

else

$o.nn_prec \leftarrow \text{prune } o.nn_before$

if $(o.count_after + |o.nn_prec| < k)$ **then**

report o as an outlier

5.1. Micro-clusters

The first version of the *pMCOd* algorithm combines the value-based partitioning and the M-tree with the micro-clusters as shown in Algorithm 3. The motivation behind using micro-clusters is to drastically reduce the number of range queries submitted to the M-tree, as explained in Section 2.2.3. In contrast to the work in [5], this version does not contain an event queue. The sliding window's state consists of the micro-clusters, the potential outliers \mathbb{PO} and the M-tree. The notion of value-partitioning from Section 4 along with the introduction of micro-clusters means that each partition is able to fully report its outliers at a faster rate without the need to communicate with the other partitions.

Each window slide starts with the eviction of the expired data points from

its state and the dissolution of the micro-clusters with $\leq k$ elements. Each data point that belonged to a dissolved micro-cluster is treated as a new data point. For each new data point, the algorithm computes its distance to the micro-clusters and if it belongs to any of them, it proceeds to updating the \mathbb{PO} metadata only. If the data point does not belong to any micro-cluster, it is inserted into the \mathbb{PO} set and a range query is executed to find its neighbors. Based on the number of neighbors of a data point in \mathbb{PO} , a new micro-cluster may be created.

Broadly, the set of potential outliers includes only points that do not belong to a micro-cluster. Also, if a point belongs to a micro-cluster, only the metadata of points in \mathbb{PO} need to be updated. After the update of each data point’s metadata, the algorithm reports the outliers by checking the data points in \mathbb{PO} . The list *o.nn_before* is pruned to contain only the non-expired objects. Each data point for which $o.count_after + |o.nn_before| < k$ holds, is reported as an outlier for the corresponding slide.

5.2. Event queue

The motivation behind the event queue is to quickly re-process data points that are affected by the expiring slide without unnecessary checks of the points in \mathbb{PO} ; it forms the cornerstone of the *COD* algorithm in [12, 5]. Every data point with a preceding neighbor is inserted in the event queue along with the timestamp of the oldest neighbor. The queue is essentially a priority queue, in the sense that the data point that has the oldest neighbor in the active window will be the first data point in the structure that makes sense to check.

Such an event queue is employed in two places. First, during the insertion of a data point, when a data point is being processed by the current window, the timestamp of its oldest neighbor is retrieved and the data point along with that timestamp are saved in the queue. Second, when a slide expires, the first item of the queue is checked to identify as to whether the timestamp belongs to or is older than the expiring slide; if this is the case, the data point needs to be re-processed in order to define its status. If the timestamp is more recent than the expiring slide, then the process stops.

The implementation of the event queue in the *pMCOD* algorithm has been done in two separate ways. The first implementation uses the *Priority Queue* data structure (Sec. 5.2:1 flavor in Figure 11), while the second uses the *TreeSet* (Sec. 5.2:2 flavor in Figure 11) data structure in Java. In both implementations, the event queue is part of the window’s state.

5.3. Usage of M-tree

In the *MCOD* and *pMCOD* algorithms, where the need for range queries has decreased, the use of the M-tree for faster range queries is under question, as to whether its overhead outweighs the benefits it incurs. To investigate this issue, we have implemented and evaluated four different implementation schemes.

In the first flavor (Sec. 5.3:1 flavor in Figure 11), as described in Section 5.1, the M-tree is used to include all the data points of the window in each partition. This means that the tree structure includes both the data points in $\mathbb{P}\mathbb{O}$ and the data points in the micro-clusters.

The second implementation (Sec. 5.3:2 flavor) does not use the M-tree at all. All range queries are done using a custom distance function, such as the euclidean, on the data points of $\mathbb{P}\mathbb{O}$ or on the centers of the micro-clusters.

The third implementation (Sec. 5.3:3 flavor) uses one M-tree to store the data points in $\mathbb{P}\mathbb{O}$. This means that, whenever a range query on the $\mathbb{P}\mathbb{O}$ is issued, it is completed by the M-tree, while a range query on the micro-clusters centers is evaluated with the help of a custom distance function.

The fourth and final flavor (Sec. 5.3:4 flavor) uses two different M-trees: one where the data points in $\mathbb{P}\mathbb{O}$ are stored and one where the centers of the micro-clusters are stored. According to this implementation, every range query issued is processed through one of the two M-trees.

6. Performance Evaluation

The experimental setting is as follows. We focus on presenting the performance as a function of the window size, the slide size, the amount of outliers (controlled through the R and k parameters), the degree of parallelism and the dimensionality of the data. The accuracy is always 100%, since all techniques are exact. We have employed three real and one artificial dataset. These datasets are static and finite, but are adequate to emulate a streaming setting. More specifically, the techniques are implemented so that they can handle infinite streams. But in order to present repeatable results, we employ publicly available datasets to produce finite streams and, unless stated otherwise, the times presented correspond to the average time per slide, aggregated over 200 slides overall. Note that, given that we are in a streaming environment and we report times per slide, the actual full dataset size does not matter.

Technique	Description
<i>naive</i>	see Section 3
<i>advanced</i>	see Section 4; no value-based partitioning and time-slicing
<i>advanced(grid)</i>	see Section 4; grid-based partitioning but no time-slicing
<i>advanced(tree)</i>	see Section 4; tree-based partitioning but no time-slicing
<i>advanced(grid) with slicing</i>	see Section 4; grid-based partitioning and time-slicing
<i>advanced(tree) with slicing</i>	see Section 4; tree-based partitioning and time-slicing
<i>pMCOD</i>	see Section 5

Table 2: Overview of techniques examined in the evaluation

First, we present results using a single multi-core machine, the same machine as the one described in Section 3.1. Then we move to experiments using a small cluster. Initially, we focus on the Stock real-world dataset. Each experiment is repeated 5 times. R and k are set to 0.45 and 50, respectively, yielding approximately 1% of outliers (i.e., the setting is similar to [10]). The default degree of parallelism, i.e., the number of Flink partitions of the window is 16, and each Flink node runs on a single core. Stock is an one-dimensional dataset. To allow a fair comparison, the timestamps are assigned in such a way that all windows are of the same size, and the slide is given as a percentage of W , e.g., a slide of 5% means that the 5% of the window contents are new arrivals.

Table 2 provides an overview of the techniques presented in the previous sections. The first ones, which exhibit the worst performance, are evaluated only on a single-machine setting. In the single machine setting, the goal is to show the benefits of value-based partitioning and micro-clustering. Then, the main flavors, which adopt value-based partitioning and either a time slicing or a micro-clustering approach are investigated in more detail in a cluster setting.

6.1. Algorithm Comparison and Evaluation on a single machine

Initially, we do not activate time-slicing in the advance solutions, pMCOD flavor is as described in Section 5.1, and the grid-based technique is used for value-based partitioning. The main purpose of the experiments using a single machine is to prune the space of the techniques on which it is worth to focus and provide evidence about the higher efficiency of *pMCOD*.




Figure 6: Average (top) and median (bottom) processing time per slide for different slides and window of 10K objects.

In the first experiment, we employ a window of 10K objects, while the slide varies from 5% to 50%. The results are shown in Figure 6, where both the average and the median times are reported. From the figure, we can draw the following observations:

1. *pMCOD* improves upon the advanced solutions with both random and value-based partitioning (labeled as *advanced* and *advanced(grid)*, re-




Figure 7: Throughput comparison for $W = 10K$ and slides of 5%-50%.




Figure 8: The impact of the degree of parallelism P (left), k (middle) and the window size W (right - for slide sizes of 5%-50%) on pMCOD's performance.

spectively) by an order of magnitude for slides up to 20%; e.g., for slide of 10%, it is 10X faster, and for 20%, it is 13X. For slides of 50%, where half of the window points are new in each slide, *pMCOD* is faster than *advanced(grid)* by 2.74 times.

2. *Advanced* dominates *advanced(grid)* for small slides, while the latter is better for large ones, which is mostly attributed to the fact that *advanced(grid)* inherent load imbalance⁸ is outweighed by the benefits of less replication and communication in large slides.

⁸The grid cell used for partitioning is based on an initial sample and thus no guarantees can be provided as to how balanced the workload distribution is throughout stream processing.

3. Despite the non-negligible standard deviation, the trends in the average values are similar to those in the median ones.
4. Relying on parallelism to efficiently evaluate pairwise distances in a brute-force manner is a highly suboptimal choice; *pMCOD* improves upon the *naive* solution by two orders of magnitude; for example, it is 117X faster for slide 10%.

Figure 7 compares *pMCOD* against the throughput results obtained by the baseline technique, described in Section 3.1. The improvements are up to 2076 times, whereas *pMCOD*'s throughput exceeds 33240 new objects per second for $S = 5\%$. This throughput corresponds to the capability of the single machine to process over 11GB of measurement data per day. However, we observe that the throughput degrades as the slide increases.

In the second experiment, we focus on *pMCOD* and we examine the impact of three parameters, namely the degree of parallelism P , k and the window size W . The results are summarized in Figure 8. Regarding the degree of parallelism, the left figure refers to a setting where W is 10K and the slide is 5%. We see that *pMCOD* scales well and the time per slide drops nearly two times when we go from two partitions to four. Since the machine is a 4-core/8-thread one, the gains are small for higher degrees of parallelism. We also see that our default configuration of $P = 16$ yields the highest performance. In the middle figure, we see that, as we increase the k value, the performance degrades. Increasing the k value implies more outliers and higher difficulty in forming micro-clusters; as such, this result is reasonable. Finally, the rightmost figure reveals that *pMCOD* scales well with the size of the window: a ten-fold increase in the window size results in similar increases in the processing time for slides of 5% and 10%, and smaller increases for larger slides. More specifically, for 20% slide magnitude, the increase in the processing time in the 100K window is less than 8 times, and for 50% slide, it is 6.62 times only. We further examine the impact of the input parameters in a cluster setting later.

6.1.1. Using Additional Datasets with More Dimensions

First, we provide results using an artificial dataset, which is generated from a mixture of three Gaussian distributions and taken from [10]. We set $W = 10K$, $R = 0.28$ and $k = 50$. Figure 9 presents the results, where it is shown that the main observations drawn for the Stock real-world dataset hold and *pMCOD* outperforms other solutions by an order of magnitude for slides up to 20%.




Figure 9: Average slide processing time for slide sizes of 5%-50% for the artificial dataset

In the next experiment, we employ two additional real-world datasets, namely Forest Cover (FC)⁹ and TAO¹⁰, considering 2 and 3 dimensions, respectively. We configure their parameters so that always k is 50, while W is kept to 10K. More specifically, for FC, we set $R = 34$ on the 2nd and the 5th dimensions (corresponding to 1.3% outliers - we discard the other 53 dimensions), and for TAO, we set $R = 1.9$ on all three dimensions (corresponding to 0.98% outliers).

The results are shown in Figure 10 (actually, for TAO and $S=50\%$, *advanced* crashed due to memory shortage). The key observations are as follows: (i) *pMCOd* behavior is nearly the same for 2 or 3 dimensions but significantly worse than the behavior for the one-dimensional datasets; the latter is attributed to the known performance degradation of M-tree for higher dimensions and the increased replication; (ii) for the FC dataset, and due to the increased number of outliers compared to the other settings, *advanced(grid)* slightly outperforms *pMCOd* when $S = 50\%$, which means that half of the

⁹ Available from <http://kdd.ics.uci.edu>

¹⁰ Available from <http://www.pmel.noaa.gov>

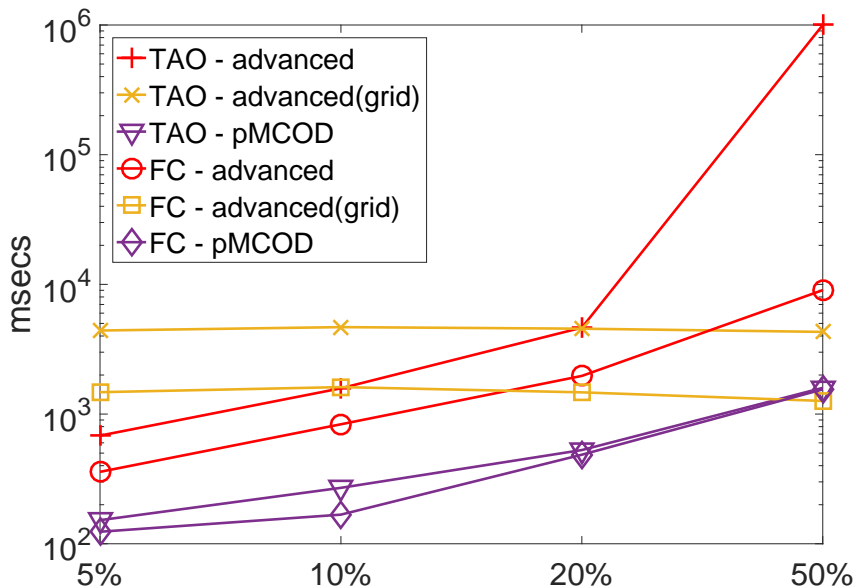


Figure 10: Average slide processing time for slide sizes of 5%-50% for the TAO and FC datasets

points in each window slide are new arrivals; and (iii) *advanced* and *advanced(grid)* are significantly affected by the increase in the number of dimensions considered from 2 to 3.

6.1.2. pMCOd's features

Thus far, *pMCOd* has been the best-performing technique using the rationale in Section 5.1 (first alternative in Section 5.3). To assess the effectiveness of the event queue and the more restricted usage of M-tree, we run experiments over all datasets and window sizes and slides employed previously keeping the portion of outliers approximately at 1%. The normalized results appear in Figure 11. From the figure, the following observations can be extracted:

1. Employing the M-tree for all data (as in our default mode) incurs a significant overhead. Micro-clustering is so effective in our experiments that renders any additional structure that has shown to be advantageous in other settings negligible. If we restrict the usage of the M-tree, as in the third and fourth flavors in Section 5.3, the overhead due to




Figure 11: Normalized times over all datasets, and window and slide sizes for the alternatives of pMCOD in Section 5.

the need of continuous updates observed is 7.7%.

2. The usage of the event queue (see the two rightmost bars) incurs overhead that is not outweighed by the benefits. On average, the algorithm is slowed down by 5-6% compared to the case without any additional structure. Also, the implementation using a *Treaset* outperforms almost consistently the one using *Priority Queue* but by a negligible margin.

The above observations show that the *pMCOD* algorithm runs faster without an event queue and the usage of M-trees for the range queries. The main reason lies in the use of micro-clusters that manage to cover almost all the non-outlier objects. In other words, not only the outliers are a very small fraction of the data, but the \mathbb{PO} set is small as well. As such, the Occam’s razor principle applies, according to which, between models of similar performance we should adopt the simpler one [17]. In the remainder of our experiments, the second and simpler flavor of *pMCOD* that does not employ any additional data structures (see Section 5.3) is used, which means that all the performance benefits of *pMCOD* can be safely attributed to the micro-clustering approach.

6.2. Experiments in a cluster

The experiments in the cluster do not repeat the single-machine ones but build upon them. The main lessons learned thus far is that (i) *pMCOD* outperforms simple *advanced* solutions by up to an order of magnitude and (ii) data structures, such as an event queue and a M-tree, are unnecessary for *pMCOD*. As such, the purpose of the experiments using a cluster is to provide better evidence about the behavior of the partitioning alternatives and the scalability of the approach, especially regarding the slide size and the data dimensionality, where the scalability using a single machine has been problematic. Also, since *pMCOD* dominates the top four techniques in Table 2, we shift our focus on comparing *pMCOD* against the remaining more sophisticated advanced solutions that encapsulate the time-slicing rationale.

The cluster consists of three machines with similar characteristics in terms of CPU speed. Apart from the 4-core/ 8-thread machine with 32 GB of RAM that we used in the previous experiments, we connect through 1 Gbps Ethernet a 6-core/12-thread machine with 64GB of RAM and another one with 8-core/8-threads and 32 GB of RAM. The setup is as a standalone cluster with one *Job Manager* on the 6-core/12-thread machine and one *Task Manager* on each machine, i.e., overall we have three Task Managers.

Before proceeding to the main experiments, we perform some basic tuning, since the results in Figure 8(left) do not hold anymore. We employ the Stock dataset and a window of 10K objects, with the slide varying between 5% and 50% using the *grid partitioning*, in order to gain some insight into the performance of the cluster when the number of the task slots varies. Figure 12 shows that with the number of the task slots being equal to the number of the total cores in the cluster, the running time is on average 25% lower than when the task slots are twice the number of threads and 12% lower than when the task slots match the threads. Therefore, we fix the number of tasks accordingly.

6.2.1. Partitioning evaluation and concluding remarks about the data structures

In the next experiment, we use the *pMCOD* algorithm on the Stock (1-dimensional), FC (2-dimensional) and TAO (3-dimensional) datasets in order to compare the efficiency of the *grid based* and the *VP-tree based* partitioning. Each grid was created from statistics extracted from the whole dataset, while each VP-tree was created from the first 10K data points of each dataset. From Figure 13, the following observations can be extracted:

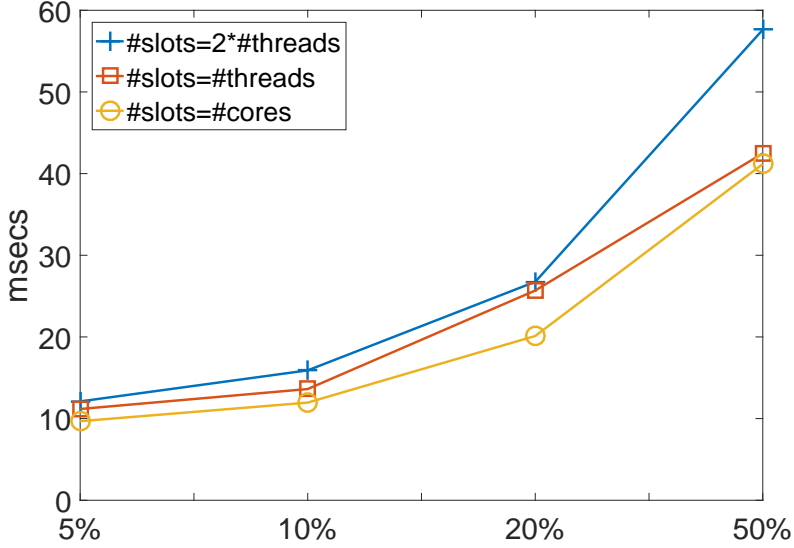


Figure 12: Average processing time per slide for a 10K window with different numbers of Flink task slots.

1. The scalability in the slide size is much better compared to the single-machine cases (we further elaborate on this in the next sub-section).
2. In the 1 and 2-dimensional datasets, the grid-based partitioning is faster than the VP-tree-based one by 1.3 to 3 times. This stems from the fact that the replication rate of the VP-tree partitioning is higher than the replication rate of the grid.
3. In the 3-dimensional dataset, the VP-tree partitioning is faster than the grid partitioning by approximately 1.4 times. This is due to the fact that its replication rate is lower than the rate from the grid partitioning.
4. Using the VP-tree based partitioning, the 3-dimensional dataset times are faster than those for the 2-dimensional dataset in most of the cases, which is due to both the lower replication rate in the 3-dimensional dataset and the difference in the total percentage of the outliers in the two datasets.
5. *pMCO*D scales better with regards to the data dimensionality when the tree-based partitioning is used. For example, the ratios of the 3-dimensional dataset times to those of the single-dimensional for slides 5%, 10%, 20% and 50% are 2.5, 3.1, 2 and 2.1, respectively. If, for

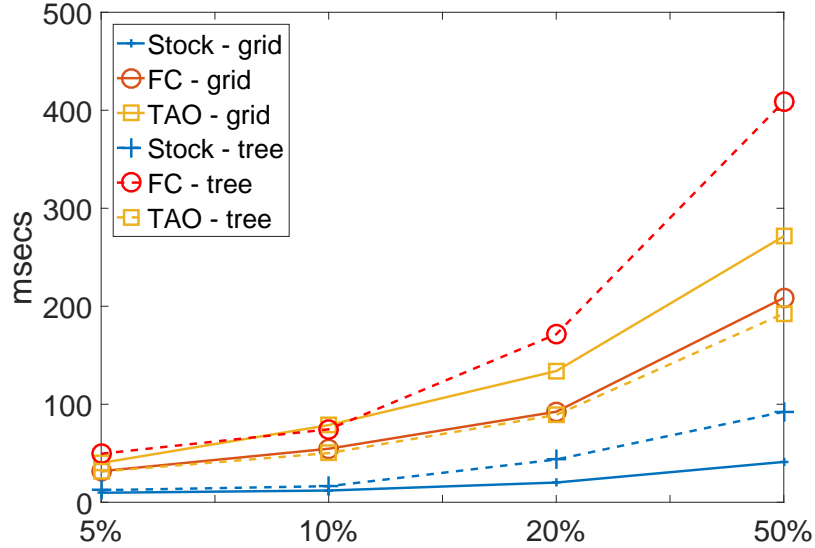


Figure 13: Average processing time per slide for a 10K window with grid based and VP-tree based partitioning.

the single-dimensional dataset, we consider the times of the grid-based partitioning that are lower, the ratios become 3.2, 4.2, 4.4, and 4.7, respectively, which still denote much improved scalability.

6. The observed speed-ups are much higher than previously. For instance, for the Stock dataset, when the slide size is 50%, *pMCOD* is 510X faster than the naive parallel solution.

Based on the above results and the results already presented in Sec. 6.1.2, we can summarize the key insights regarding the data structures employed, as shown in Table 3.

6.2.2. Slicing evaluation - comparison against additional alternatives

In order to extend the evidence that *pMCOD* is the best-performing technique, we further compare it against the flavors of the advanced solutions that follow the time-slicing rationale. Figure 14 shows this comparison for the Stock dataset. The *pMCOD* algorithm remains the clear winner and the speedup over the best performing *advanced w/ slicing* solution reached 4.3X. Also, as previously, the grid partitioning outperforms the tree-based one for the Stock dataset for the advanced algorithm, too.

Type	Summary remarks
M-tree	key data structure for improving the performance of <i>advanced</i> solutions over the naive one; does not play significant role in <i>pMCOD</i> , due to the effectiveness of micro-clusters
event queue	originally proposed in <i>MCOD</i> ; does not need to be incorporated in <i>pMCOD</i> , due to the effectiveness of micro-clusters
tree vs grid	needed for value-based partitioning with no clear winner when both apply; tree-based partitioning is the only choice for non-euclidean spaces

Table 3: Concluding remarks regarding the data structures employed

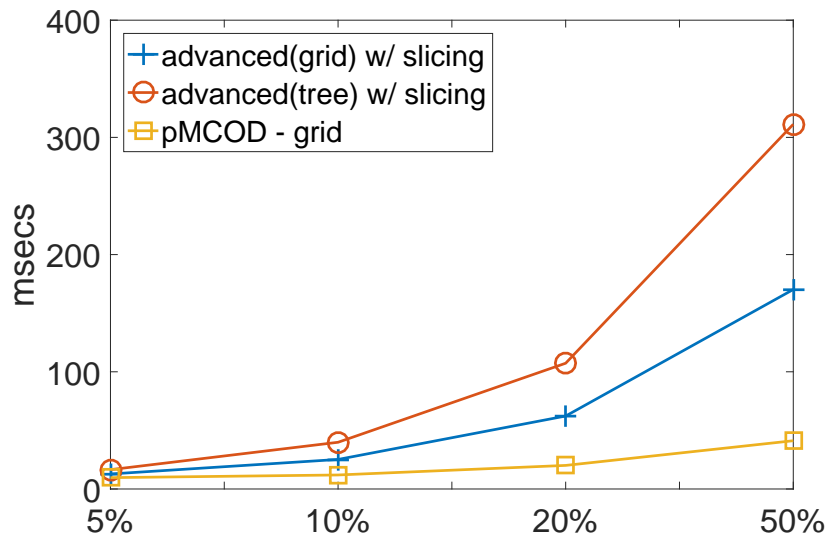


Figure 14: Average slide processing time for slide sizes of 5%-50% of *pMCOD* and slicing-endowed advanced algorithms.

6.2.3. Comparison against results in [10](non-parallel)

Here, we compare our results regarding the Stock dataset against those of non-parallel *MCOD* [5], as evaluated by third parties in [10]. The evaluation in [10] also uses a processor with clock speed at 3.5GHz, but without giving details about the number of cores. Nevertheless, our results can directly compare against those in Figures 6 and 10 in [10]. Due to the log scale used, we can only report approximate values, as shown in Table 4. The speedup in a single machine is between 2.66X and 6.65X, which provides

W	S	MCOD in [10]	MCOD from [10] 1 machine	pMCOD 1 machine	pMCOD 3 machines
10K	5%	100 (approx.)	16.600	15.042	9.681
100K	5%	700 (approx.)	364.66	160.44	41.106
100K	10%	1420 (approx.)	730.36	335.5	74.582
100K	20%	2650 (approx.)	1382.22	771.15	124.277
100K	50%	5270 (approx.)	2731.2	1981.25	291.556

Table 4: Average time to process a slide in different works (in msec).

strong insights into the parallelization efficiency of our solutions on a 4-core machine. However, when we move to a 3-machine cluster, we observe speed-ups up to 21.32X. We rerun the code of [10] on our single-machine setting, as shown in the fourth column of the table. The speed-up is up to 2.27X on a single machine and up to 11.12X on a cluster. These speedups include also any overheads incurred by the Flink framework and thus the evidence regarding the parallelization efficiency is stronger. Overall, comparing the fourth and the fifth column of Table 4 shows the benefits of *pMCOD* over *MCOD* on an even playing field, whereas the leftmost column shows the even higher potential of *pMCOD*.

6.2.4. Further experiments regarding sensitivity analysis and the scalability of pMCOD

With the help of Figure 15, we continue the discussion about the improved scalability regarding the slide size in a cluster setting. First, when the slide size doubles, the running time increases by a smaller margin in the cluster. Further, we can observe that, for large slides (e.g., 50%), the speed-up is super-linear. For example, the times in the cluster for the 1, 2 and 3-dimensional datasets are 7.3, 7.4 and 5.9 times lower than the times on a single machine, respectively. Since the single machine possesses 4 cores and the cluster consists of 18 cores, a linear speed-up corresponds to 4.5 times lower. Superlinear speed-ups are not uncommon in MapReduce-like environments, due to the reduced memory contention of the system as the number of machines increases [18]; i.e., more memory becomes available for processing the same amount of data.

For completeness, we also show the behavior for the Stock dataset, when the slide is only 1%, which corresponds to a setting with much more frequent reporting but fewer modifications across window slides. The results are shown

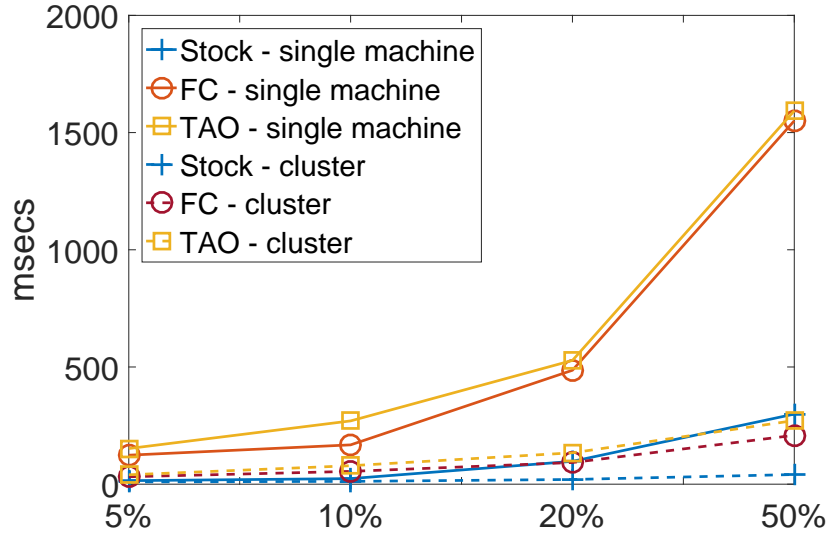


Figure 15: Average slide processing time of *pMCO*D for slide sizes of 5%-50% in the two main experimental settings.

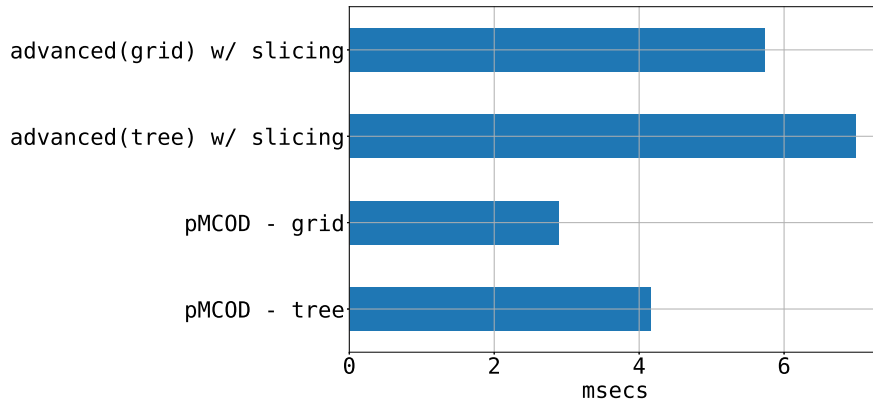


Figure 16: Average slide processing time for the best performing alternatives when the slide is 1%, $W = 10K$ and the Stock dataset is used

in Figure 16.

In the next experiment, we vary the R and k values, as shown in Tables 5 and 6. We draw three main conclusions for the Stock dataset: (i) in general, for the same R , increasing k from 40 to 60 leads to increases in the average processing time by up to 22% for *pMCO*D and 15% for *advanced w/ slicing*;

Algorithm	Slide	R=0.45					k=50			
		$k=50, R=0.45$	K=40	K=45	K=55	K=60	R=0.40	R=0.445	R=0.455	R=0.50
outliers		1.42%	1.10%	1.25%	1.58%	1.73%	1.62%	1.43%	1.40%	1.21%
pMCOd	5%	11.156	10.431	10.493	11.26	11.489	11.477	11.192	10.839	10.795
Adv.(Slicing)		<i>13.917</i>	13.219	14.065	14.188	14.574	14.632	14.712	14.197	14.423
pMCOd	10%	<i>14.254</i>	12.955	13.56	13.851	15.821	14.902	13.402	13.882	14.685
Adv.(Slicing)		29.149	28.68	27.361	25.372	26.343	25.059	28.844	26.372	28.115
pMCOd	20%	22.791	20.695	21.687	22.895	23.69	23.83	23.038	22.221	23.98
Adv.(Slicing)		61.302	61.576	66.542	55.789	70.712	57.85	63.917	55.374	70.339
pMCOd	50%	46.126	41.452	43.227	47.85	50.72	51.091	45.328	45.945	41.048
Adv.(Slicing)		195.944	180.805	198.115	205.499	207.539	169.292	193.763	186.193	213.132

Table 5: Average slide processing times in msec for varying R , k and slide values for the Stock dataset and a 10K window.

Algorithm	Slide	R=1.9					k=50			
		$k=50, R=1.9$	K=40	K=45	K=55	K=60	R=1.80	R=1.85	R=1.95	R=2.00
pMCOd	5%	40.405	37.337	38.638	46.709	46.373	45.482	44.561	41.721	42.128
Adv.(Slicing)		<i>33.434</i>	31.273	30.247	30.846	31.583	31.308	31.12	31.534	33.206
pMCOd	10%	<i>81.697</i>	57.687	66.408	92.979	89.289	82.876	79.68	78.013	79.698
Adv.(Slicing)		56.122	58.148	59.084	57.6	57.451	57.863	57.558	63.19	59.917
pMCOd	20%	139.439	112.156	130.282	154.58	163.037	131.311	134.531	144.411	130.569
Adv.(Slicing)		152.618	145.632	156.731	157.396	151.312	148.428	157.985	157.515	154.742
pMCOd	50%	281.853	243.299	277.185	308.346	323.081	297.086	273.34	283.282	277.06
Adv.(Slicing)		847.95	817.382	814.898	866.342	886.019	895.704	809.55	911.33	947.725

Table 6: Average slide processing times in msec for varying R , k and slide values for the TAO dataset and a 10K window.

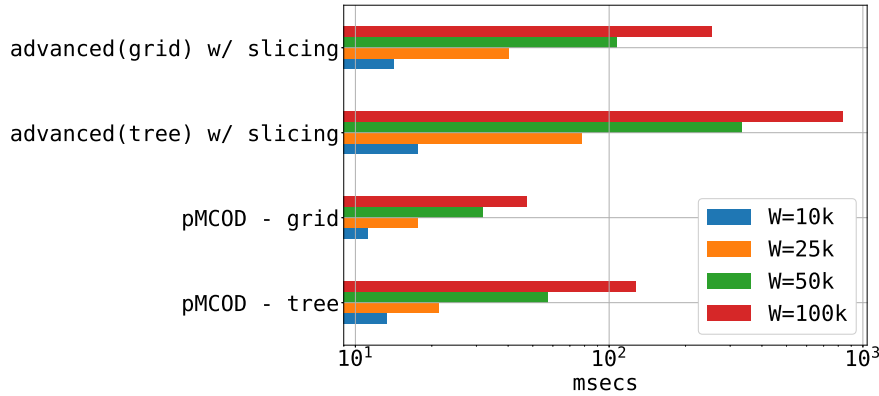


Figure 17: Average slide processing times for 5% slide for the Stock dataset.

and (ii) for the same k , increasing the R results in performance improvements for *pMCOd* but performance degradations for *advanced w/ slicing*, by up to 20% and 25%, respectively; and (iii) *pMCOd* always outperforms *advanced w/ slicing* and the largest observed speedup is 5.2X.

Examining the results for the TAO dataset (Table 6), reveals some in-




Figure 18: Average slide processing times of *pMCOD* for varying window and slide sizes when using a 5d dataset.

interesting features regarding the behavior of the alternative presented: (i) *advanced w/ slicing* may outperform *pMCOD*, especially for slides 5% and 10%; (ii) the highest speedup of *pMCOD* compared to *advanced w/ slicing* is 3.4X and (iii) the performance is more insensitive to R and k modifications apart from *pMCOD* with fixed R and varying k .

We complete our sensitivity analysis through varying the window size for the Stock dataset using 5% slide. As shown in Figure 17, *pMCOD* is the dominant solution by up to 5.3X and the performance gap increases with the window size between the two most efficient techniques.

6.2.5. High-dimensionality experiments and discussion

Finally, we have already argued about the improved scalability of *pMCOD* in data dimensionality up to 3 dimensions. In another experiment, we used a new 5-dimensional dataset by selecting the first 5 dimensions of the FC dataset and normalizing the values. With $R = 0.182$ and $k = 50$ approximately 1.2% data points of each window are outliers while the partitioning is done using the VP-tree based technique. In such a case, the average processing time per slide size is, on average, more than 50 times higher than the corresponding times for Stock; compared to the 2-d FC dataset, the increases are over 12X. This implies that many issues remain to be resolved for

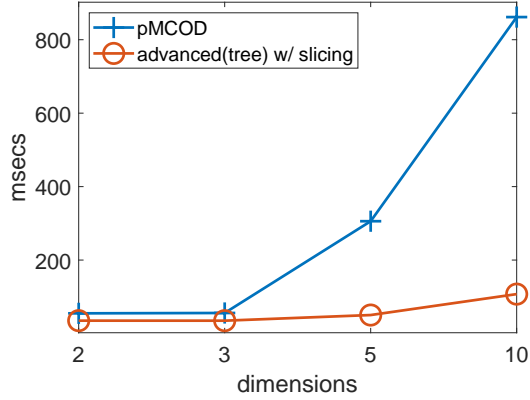


Figure 19: Average slide processing times of *pMCOD* and advanced with slicing for 5% slide and tree-based partitioning for varying number of dimensions using the FC dataset.

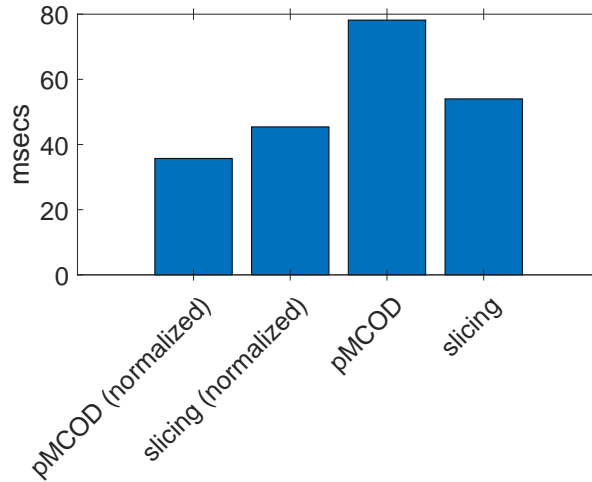


Figure 20: Average slide processing times for 10 random dimensions of the FC dataset when normalizing the data and when there is no normalization.

efficient parallel outlier mining in high-dimensional data. Results are shown in Figure 18.

Next, we also compare *advanced w/ slicing* and *pMCOD* using the FC dataset from which we take the first 2, 3, 5 and 10 dimensions, as shown in Figure 19. From this figure, we verify the observation in the previous experiment that *advanced w/ slicing* may outperform *pMCOD*, and more

specifically, *advanced w/ slicing* scales very well with the number of dimensions. However, this behavior is very dependent on the dataset. For example, we have previously observed that the behavior is different in the TAO dataset (3 dimensions). Also, in Figure 20, we see the performance of the two techniques when taking 10 random dimensions from FC; when normalizing the data (through z-score normalization) the running time of *pMCOD* is an order of magnitude lower and faster than *advanced w/ slicing*.

In summary, developing more efficient solutions for datasets with higher number of dimensions than 3 is left for future work, since distance-based outlier detection is known to be problematic in such cases due to the curse of dimensionality issues [1]. The solutions for high-dimensional settings employ a different approach, such as other metrics for outlier definition, approximations, dimensionality reduction and subspace clustering, which are out of our scope. More specifically, the two most promising approaches to outlier detection in high-dimensional spaces are angle-based solutions, as explained in [19, 20], and distance-based ones after performing judicious feature selection. The former significantly depart from the solutions presented here; the latter, as explained by recent advances in [21], render distance-based outlier detection applicable but report performance differences of orders of magnitude if dimensions are not properly selected.

7. Related Work

Outlier detection has been a topic that has attracted a lot of interest and there are several comprehensive surveys, e.g., [22, 23, 24, 1]. In Section 2, we have already discussed algorithms for outlier detection in streams. The next most related area to our work is parallel algorithms for distance-based outlier detection. A distributed outlier detection algorithm for massive datasets is proposed in [9]. The two key points of this research is the initial partitioning of the data and the different outlier detection algorithm that each partition may run. The partitioning resembles our value-based one and focuses on the workload that each partition receives. Then, in each partition, the exact outlier detection algorithm out of two candidates is chosen. However, none of these candidates is suitable in a streaming setting, since they do not cover or deal with issues related to incremental outlier detection across different slides. There are also some works that assume parallel infrastructures that cannot scale and do not follow the paradigm introduced by MapReduce and its modern extensions, e.g., [25, 26]. Another example is the work in [4],

which delivers an approximate solution for decentralized settings, where the processing nodes conform to a hierarchy and disjoint areas are examined independently. Overall, our work is the first one that combines streaming and massively parallel solutions to the problem of distance-based outlier detection. However, parallel and streaming anomaly detection has been considered for other definitions of outlierness, e.g., as in [27].

A related yet different problem is examined in [28]. In a production distributed environment, a stream of data points may be split across multiple nodes, each holding part of the values of a data point. These parts will eventually need to be aggregated on a core node for outlier detection, but this incurs significant communication cost. The solution proposed is based on compressing local data into a sketch. Another related problem is that of supporting multiple outlier detection queries, i.e., combinations of R and k values. Examples include [29] and [5]. The latter presents multi-query extensions to MCOB and its approach is compatible with the our parallel pMCOB solution; here, we have examined single-query solutions only.

Finally, apart from the platforms discussed, there are additional alternatives. For example, ChronoStream [30] is a prototype system for elastic big stream processing in a distributed environment, providing low latency. However, we have decided to adopt Flink because it combines strong positive features, as discussed in Section 2.1, with mature engineering and a broad user community, while we did not consider elasticity issues in this work.

8. Concluding Remarks

This work targets streaming distance-based outlier detection and provides the first solutions to date to this problem, when examined in a massively parallel setting, such as Flink. We have proposed a series of alternative techniques, with the one termed as *pMCOB* being a clear winner when the number of dimensions is low in the experiments that we have conducted using three real-world and one synthetic dataset. The improvements upon other solutions are significant compared to the second best solution. There are also good speedups, up to 2.27 times, compared to the non-parallel solutions implemented by third parties in [10], when running on a 4-core machine. Even higher performance is achieved in a small cluster with speedups of 11X. Our solutions have been made publicly available. The motivation behind our work is to fill a gap in the currently offered solutions in large-scale streaming big data analytics. Moreover, our solutions aspire to act as a reference point for

future techniques that target both continuous reporting of distance-based outliers and a massively parallel setting; to this end, the alternative techniques are not tailored to Flink but they can be transferred to other similar frameworks. For example, porting these techniques to Spark is an interesting direction for future extensions.

We identify three additional avenues for such future extensions. First, further research is required in order to make grid and VP-tree-based partitioning more practical and adaptively balanced, addressing also the issue of acquiring both initial and evolving metadata to reach efficient partitioning decisions on the fly. Second, the current solutions, similarly to any distance-based outlier detection technique, are sensitive to the input parameters regarding the neighborhood radius and the threshold on the neighbors; evaluating multiple parameters in parallel is a promising approach to tackle this drawback, but, to this end, several issues regarding performance and interpretability need to be resolved using [31], which offers an initial parallel multi-parameter solution, as a starting point. Third, further improvements need to be made in order to address the high-dimensionality problems of data possibly adopting additional definitions of outlierness and thus departing from the distance-based one and leveraging subspace techniques as explained in Section 6.2.5. Finally, another line of future research relates to approximate outlier mining; here, we have provided exact solutions only.

References

- [1] C. C. Aggarwal, *Outlier Analysis*, 2ed, Springer, 2018 (2018).
- [2] R. Johnson, *Applied Multivariate Statistical Analysis*, Prentice Hall, 1992 (1992).
- [3] E. Knorr, R. Ng, V. Tucakov, Distance-based outliers: algorithms and applications, *The VLDB Journal* 8 (3-4) (2000) 237–253 (2000).
- [4] S. Subramaniam, T. Palpanas, D. Papadopoulos, V. Kalogeraki, D. Gunopoulos, Online outlier detection in sensor data using non-parametric models., in: *VLDB*, 2006, pp. 187–198 (2006).
- [5] M. Kontaki, A. Gounaris, A. N. Papadopoulos, K. Tsihlias, Y. Manolopoulos, Efficient and flexible algorithms for monitoring distance-based outliers over data streams, *Inf. Syst.* 55 (2016) 37–53 (2016).

- [6] D. Yang, E. Rundensteiner, M. Ward, Neighbor-based pattern detection for windows over streaming data, in: EDBT, 2009, pp. 529–540 (2009).
- [7] L. Cao, D. Yang, Q. Wang, Y. Yu, J. Wang, E. A. Rundensteiner, Scalable distance-based outlier detection over high-volume data streams, in: ICDE, 2014, pp. 76–87 (2014).
- [8] F. Angiulli, F. Fassetti, Detecting distance-based outliers in streams of data, in: CIKM, 2007, pp. 811–820 (2007).
- [9] L. Cao, Y. Yan, C. Kuhlman, Q. Wang, E. A. Rundensteiner, M. Y. Eltabakh, Multi-tactic distance-based outlier detection, in: ICDE, 2017, pp. 959–970 (2017).
- [10] L. Tran, L. Fan, C. Shahabi, Distance-based outlier detection in data streams, PVLDB 9 (12) (2016) 1089–1100 (2016).
- [11] T. Toliopoulos, A. Gounaris, K. Tsihclas, A. Papadopoulos, S. Sampaio, Parallel continuous outlier mining in streaming data, in: 5th IEEE International Conference on Data Science and Advanced Analytics (DSAA), 2018 (2018).
- [12] M. Kontaki, A. Gounaris, A. N. Papadopoulos, K. Tsihclas, Y. Manolopoulos, Continuous monitoring of distance-based outliers over data streams, in: ICDE, 2011, pp. 135–146 (2011).
- [13] M. Kontaki, A. Gounaris, A. N. Papadopoulos, K. Tsihclas, Y. Manolopoulos, Continuous outlier detection in data streams: an extensible framework and state-of-the-art algorithms, in: SIGMOD, 2013, pp. 1061–1064 (2013).
- [14] S. Wang, E. A. Rundensteiner, Scalable stream join processing with expensive predicates: workload distribution and adaptation by time-slicing, in: EDBT, 2009, pp. 299–310 (2009).
- [15] P. Ciaccia, M. Patella, P. Zezula, M-tree: An efficient access method for similarity search in metric spaces, in: VLDB, 1997, pp. 426–435 (1997).
- [16] P. N. Yianilos, Data structures and algorithms for nearest neighbor search in general metric spaces, in: SODA, Vol. 93, 1993, pp. 311–321 (1993).

- [17] P. Tan, M. Steinbach, V. Kumar, Introduction to Data Mining, Addison-Wesley, 2005 (2005).
- [18] N. J. Gunther, P. Puglia, K. Tomasette, Hadoop superlinear scalability, *Commun. ACM* 58 (4) (2015) 46–55 (2015).
- [19] A. Zimek, E. Schubert, H. Kriegel, A survey on unsupervised outlier detection in high-dimensional numerical data, *Statistical Analysis and Data Mining* 5 (5) (2012) 363–387 (2012).
- [20] G. O. Campos, A. Zimek, J. Sander, R. J. G. B. Campello, B. Micenková, E. Schubert, I. Assent, M. E. Houle, On the evaluation of unsupervised outlier detection: measures, datasets, and an empirical study, *Data Min. Knowl. Discov.* 30 (4) (2016) 891–927 (2016).
- [21] S. Yoon, J. Lee, B. S. Lee, NETS: extremely fast outlier detection from a data stream via set-based processing, *PVLDB* 12 (11) (2019) 1303–1315 (2019).
- [22] M. Gupta, J. Gao, C. C. Aggarwal, J. Han, Outlier detection for temporal data: A survey, *IEEE Trans. Knowl. Data Eng.* 26 (9) (2014) 2250–2267 (2014).
- [23] V. Chandola, A. Banerjee, V. Kumar, Anomaly detection for discrete sequences: A survey, *IEEE Trans. Knowl. Data Eng.* 24 (5) (2012) 823–839 (2012).
- [24] L. Akoglu, H. Tong, D. Koutra, Graph based anomaly detection and description: a survey, *Data Min. Knowl. Discov.* 29 (3) (2015) 626–688 (2015).
- [25] F. Angiulli, S. Basta, S. Lodi, C. Sartori, Distributed strategies for mining outliers in large data sets, *IEEE Trans. Knowl. Data Eng.* 25 (7) (2013) 1520–1532 (2013).
- [26] K. Bhaduri, B. L. Matthews, C. Giannella, Algorithms for speeding up distance-based outlier detection, in: *Proc. of SIGKDD, 2011*, pp. 859–867 (2011).
- [27] L. Rettig, M. Khayati, P. Cudré-Mauroux, M. Piórkowski, Online anomaly detection over big data streams, in: *Big Data, 2015*, pp. 1113–1122 (2015).

- [28] Y. Yan, J. Zhang, B. Huang, X. Sun, J. Mu, Z. Zhang, T. Moscibroda, Distributed outlier detection using compressive sensing, in: Proc. of SIGMOD, ACM, 2015, pp. 3–16 (2015).
- [29] L. Cao, J. Wang, E. A. Rundensteiner, Sharing-aware outlier analytics over high-volume data streams, in: Proc. of SIGMOD, 2016, pp. 527–540 (2016).
- [30] Y. Wu, K.-L. Tan, Chronostream: Elastic stateful stream computation in the cloud, in: ICDE, IEEE, 2015, pp. 723–734 (2015).
- [31] T. Toliopoulos, A. Gounaris, Multi-parameter streaming outlier detection, in: 2019 IEEE/WIC/ACM International Conference on Web Intelligence, WI 2019, Thessaloniki, Greece, October 14-17, 2019, 2019, pp. 208–216 (2019).