

A Distributed Online Learning Approach for Pattern Prediction over Movement Event Streams with Apache Flink

Ehab Qadah
Fraunhofer IAIS
Sankt Augustin, Germany
Ehab.Qadah@iais.fraunhofer.de

Elias Alevizos
NCSR "Demokritos"
Athens, Greece
alevizos.elias@iit.demokritos.gr

Michael Mock
Fraunhofer IAIS
Sankt Augustin, Germany
Michael.Mock@iais.fraunhofer.de

Georg Fuchs
Fraunhofer IAIS
Sankt Augustin, Germany
Georg.Fuchs@iais.fraunhofer.de

ABSTRACT

In this paper, we present a distributed online prediction system for user-defined patterns over multiple massive streams of movement events, built using the general purpose stream processing framework Apache Flink. The proposed approach is based on combining probabilistic event pattern prediction models on multiple predictor nodes with a distributed online learning protocol in order to continuously learn the parameters of a global prediction model and share them among the predictors in a communication-efficient way. Our approach enables the collaborative learning between the predictors (i.e., "learn from each other"), thus the learning rate is accelerated with less data for each predictor. The underlying model provides online predictions about when a pattern (i.e., a regular expression over the event types) is expected to be completed within each event stream. We describe the distributed architecture of the proposed system, its implementation in Flink, and present experimental results over real-world event streams related to trajectories of moving vessels.

1 INTRODUCTION

In recent years, technological advances have led to a growing availability of massive amounts of continuous streaming data (i.e., data streams observing events) in many application domains such as social networks [23], Internet of Things (IoT) [24], and maritime surveillance [31]. The ability to detect and predict the full matches of a pattern of interest (e.g., a certain sequence of events), defined by a domain expert, is typically important for operational decision making tasks in the respective domains.

An event stream is an unbounded collection of time-ordered data observations in the form of a tuple of attributes that is composed of a value from finite event types along with other categorical and numerical attributes. In this work, we deal with movement event streams. For instance, in the context of maritime surveillance the event stream of a moving vessel consists of spatio-temporal and kinematic information along with the vessel's identification and its trajectory related events, based on the automatic identification system (AIS) [28] messages that are continuously sent by the vessel. Therefore, leveraging event patterns prediction over real-time streams of moving vessels is useful to alert maritime operation managers about suspicious activities (e.g., fast sailing vessels near ports, or illegal fishing) before they happen. However, processing real-time streaming

data with low latency is challenging, since data streams are large and distributed in nature and continuously arrive at a high rate.

In this paper, we present the design and implementation of an online, distributed and scalable pattern prediction system over multiple, massive streams of events. More precisely, we consider event streams related to trajectories of moving objects (i.e., vessels). The proposed approach is based on a novel method that combines a distributed online prediction protocol [8, 16] with an event forecasting method based on Markov chains [2]. It is implemented on top of the Big Data framework for stream processing Apache Flink [13]. We evaluate our proposed system over real-world data streams of moving vessels, which are provided in the context of the *dataAcron* project¹.

The rest of the paper is organized as follows. We discuss the related work and used frameworks in Section 2. In Section 3, we describe the problem of pattern prediction, our proposed approach, and the architecture of our system. The implementation details on top of Flink are presented in Section 4 and the experimental results in Section 5. We conclude in Section 6.

2 RELATED WORK AND BACKGROUND

2.1 Related work

Pattern prediction over event streams. The task of forecasting over time-evolving streams of data can be formulated in various ways and with varying assumptions. One common way to formalize this task is to assume that the stream is a time-series of numerical values, and the goal is to forecast at each time point n the values at some future points $n + 1$, $n + 2$, etc., (or even the output of some function of future values). This is the task of time-series forecasting [25]. Another way to formalize this task is to view streams as sequences of events, i.e., tuples with multiple, possibly categorical, attributes, like *event type*, *timestamp*, etc., and the goal is to predict future events or patterns of events. In this paper, we focus on this latter definition of forecasting (event pattern forecasting).

A substantial body of work on event forecasting comes from the field of temporal pattern mining where events are defined as 2-tuples of the form (*EventType*, *Timestamp*). The ultimate goal is to extract patterns of events in the form either of association rules [1] or frequent episode rules [22]. These methods have been extended in order to be able to learn not only rules for detecting event patterns but also rules for predicting events. For example, in [32], a variant of association rule mining is where the goal is to extract sets of event types that frequently lead to a rare, target event within a temporal window.

© 2018 Copyright held by the owner/author(s). Published in the Workshop Proceedings of the EDBT/ICDT 2018 Joint Conference (March 26, 2018, Vienna, Austria) on CEUR-WS.org (ISSN 1613-0073). Distribution of this paper is permitted under the terms of the Creative Commons license CC-by-nc-nd 4.0.

¹<http://www.datacron-project.eu/>

In [19], a probabilistic model is presented for calculating the probability of the immediately next event in the stream. This is achieved by using standard frequent episode discovery algorithms and combining them with Hidden Markov Models and mixture models. The framework of episode rules is employed in [9] as well. The output of the proposed algorithms is a set of predictive rules whose antecedent is minimal (in number of events) and temporally distant from the consequent. In [35] a set of algorithms is proposed that target batch online mining of sequential patterns, without maintaining exact frequency counts. As the stream is consumed, the learned patterns can be used to test whether a prefix matches the last events seen in the stream, indicating a possibility of occurrence for events that belong to the suffix of the rule.

Event forecasting has also attracted some attention from the field of Complex Event Processing (see [7] for a review of Complex Event Processing). One such early approach is presented in [26]. Complex event patterns are converted to automata, and subsequently, Markov chains are used in order to estimate when a pattern is expected to be fully matched. A similar approach is presented in [2], where again automata and Markov chains are employed in order to provide (future) time intervals during which a match is expected with a probability above a confidence threshold.

Distributed Online Learning. In recent years, the problem of distributed online learning has received increased attention and has been studied in [8, 16, 18, 33, 34]. A distributed online mini-batch prediction approach over multiple data streams has been proposed in [8]. This approach is based on a static synchronization method. The learners periodically communicate their local models with a central coordinator unit after consuming a fixed number of input samples/events (i.e., batch size b), in order to create a global model and share it between all learners. This work has been extended in [16] by introducing a dynamic synchronization scheme that reduces the required communication overhead. It can do so by making the local learners communicate their models only if they diverge from a reference model. In this work, we employ this protocol with event patterns prediction models over multiple event streams.

2.2 Technological Background

In the last years, many systems for large-scale and distributed stream processing have been proposed, including Spark Streaming [12], Apache Storm [14] and Apache Flink [13]. These frameworks can ingest and process real-time data streams, published from different distributed message queuing platforms, such as Apache Kafka [11] or Amazon Kinesis [5]. In this work, we implemented our system in Flink and Kafka (see Section 4).

In the datAcron project, the Flink streaming processing engine has been chosen as a primary platform for supporting the streaming operations, based on an internal comparative evaluation of several streaming platforms. Hence, we used it to implement our system. A predecessor distributed online learning framework has already been implemented in the FERARI project [10] based on Apache Storm.

Apache Flink. Apache Flink is an open source project that provides a large-scale, distributed, and stateful stream processing platform [6]. Flink is one of the most recent and pioneering Big Data processing frameworks. It provides processing models for both streaming and batch data, where the batch processing model

is treated as a special case of the streaming one (i.e., finite stream). Flink’s software stack includes the *DataStream* and *DataSet* APIs for processing infinite and finite data, respectively. These two core APIs are built on top of Flink’s core dataflow engine and provide operations on data streams or sets such as mapping, filtering, grouping, etc.

The two main data abstractions of Flink are *DataStream* and *DataSet*, they represent read-only collections of data elements. The list of elements is bounded (i.e., finite) in *DataSet*, while it is unbounded (i.e., infinite) in the case of *DataStream*. Flink’s core is a distributed streaming dataflow engine. Each Flink program is represented by a data-flow graph (i.e., directed acyclic graph - DAG) that gets executed by Flink’s dataflow engine [6]. The data flow graphs are composed of stateful operators and intermediate data stream partitions. The execution of each operator is handled by multiple parallel instances whose number is determined by the *parallelism* level. Each parallel operator instance is executed in an independent task slot on a machine within a cluster of computers [13].

Apache Kafka. Apache Kafka is a scalable, fault-tolerant, and distributed streaming framework/messaging system [11]. It allows to publish and subscribe to arbitrary data streams, which are managed in different categories (i.e., *topics*) and partitioned in the Kafka cluster. The Kafka Producer API provides the ability to publish a stream of messages to a topic. These messages can then be consumed by applications, using the Consumer API that allows them to read the published data stream in the Kafka cluster. In addition, the streams of messages are distributed and load balanced between the multiple receivers within the same consumer group for the sake of scalability.

3 SYSTEM OVERVIEW

3.1 Pattern prediction on a single stream

For our work presented in this paper, we use the approach presented in [2]. For the sake of self-containment, we briefly describe this approach in the following, first assuming that only a single stream is consumed and then adjusting for the case of multiple streams. We follow the terminology of [3, 21, 35] to formalize the problem we tackle.

3.1.1 Problem formulation. We define an input event and a stream of input events as follows:

Definition 3.1. Each event is defined as a tuple of attributes $e_i = (id, type, \tau, a_1, a_2, \dots, a_n)$, where *type* is the event type attribute that takes a value from a set of finite event types/symbols Σ , τ represents the time when the event tuple was created, the a_1, a_2, \dots, a_n are spatial or other contextual features (e.g., speed); these features are varying from one application domain to another. The attribute *id* is a unique identifier that connects the event tuple to an associated domain object.

Definition 3.2. A stream $s = \langle e_1, e_3, \dots, e_t, \dots \rangle$ is a time-ordered sequence of events.

A user-defined pattern \mathcal{P} is given in the form of a regular expression (i.e., using operators for *sequence*, *disjunction*, and *iteration*) over Σ (i.e., event types) [2]. More formally, a pattern is given through the following grammar:

Definition 3.3. $\mathcal{P} := E \mid \mathcal{P}_1; \mathcal{P}_2 \mid \mathcal{P}_1 \vee \mathcal{P}_2 \mid \mathcal{P}_1^*$, where $E \in \Sigma$ is a constant event type. $;$ stands for sequence, \vee for disjunction and $*$ for Kleene – *. The pattern $\mathcal{P} := E$ is matched by reading

an event e_i iff $e_i.type = E$. The other cases are matched as in standard automata theory.

The problem at hand may then be stated as follows: given a stream s of low-level events and a pattern \mathcal{P} , the goal is to estimate at each new event arrival the number of future events that we will need to wait for until the pattern is satisfied (and therefore a full match is detected).

3.1.2 Proposed approach. As a first step, event patterns are converted to deterministic finite automata (DFA) through standard conversion algorithms [15]. As an example, see Figure 1a for the DFA of the simple sequential pattern $\mathcal{P} = a; d; c$ and an alphabet $\Sigma = \{a, b, c, d\}$ (note that the DFA has no dead states since we need to handle streams and not strings). The next step is to derive a Markov chain that will be able to provide a probabilistic description of the DFA's run-time behavior.

Towards this goal, we use Pattern Markov Chains, as was proposed in [27]. Under the assumption that the input stream is generated by an m -order Markov source, and after performing a transformation step on the initial DFA to handle the m^{th} order (see [27] for more details). It can be shown that there is a direct mapping of the states of the DFA to states of a Markov chain and the transitions of the DFA to transitions of the Markov chain. The transition probabilities are then conditional probabilities on the event types.

We call such a derived Markov chain a Pattern Markov Chain (PMC) of order m and denote by $PMC_{\mathcal{P}}^m$, where \mathcal{P} is the initial pattern and m the assumed order. As an example, see Figure 1b, which depicts the PMC of order 1 for the generated DFA of Figure 1a.

After constructing a PMC, we can use it to calculate the so-called *waiting-time* distributions. Given a specific state of the PMC, a *waiting-time* distribution gives us the probability of reaching a set of absorbing states in n transition from now (absorbing states are states with self-loops and probability equal to 1.0). By mapping the final states of the initial DFA to absorbing states of the PMC (see again Figure 1). Therefore, we can calculate the probability of reaching a final state, or, in other words, of detecting a full match of the original regular expression in n events from now.

In order to estimate the final forecasts, another step is required, since our aim is not to provide a single future point with the highest probability but an interval. Predictions are given in the form of intervals, as $I = (start, end)$. The meaning of such an interval is that the DFA is expected to reach a final state sometime in the future between the *start* and *end* with probability at least some constant threshold θ_{fc} (provided by the user). These intervals are estimated by a single-pass algorithm that scans a waiting-time distribution and finds the smallest (in terms of length) interval that has probability exceeds this threshold (θ_{fc}). For example, Figure 2a shows the *waiting-time* distributions for the non-final states of the DFA in Figure 1, and the computed prediction intervals are depicted in Figure 2b.

The method described above assumes that we know the (possibly conditional) occurrence probabilities of the various event types appearing in a stream (as would be the case with synthetically generated streams). However, this is not always the case in real-world situations. Therefore, it is crucial for a system implementing this method to have the capability to learn the values of the PMC's transition matrix. One way to do this is to use some

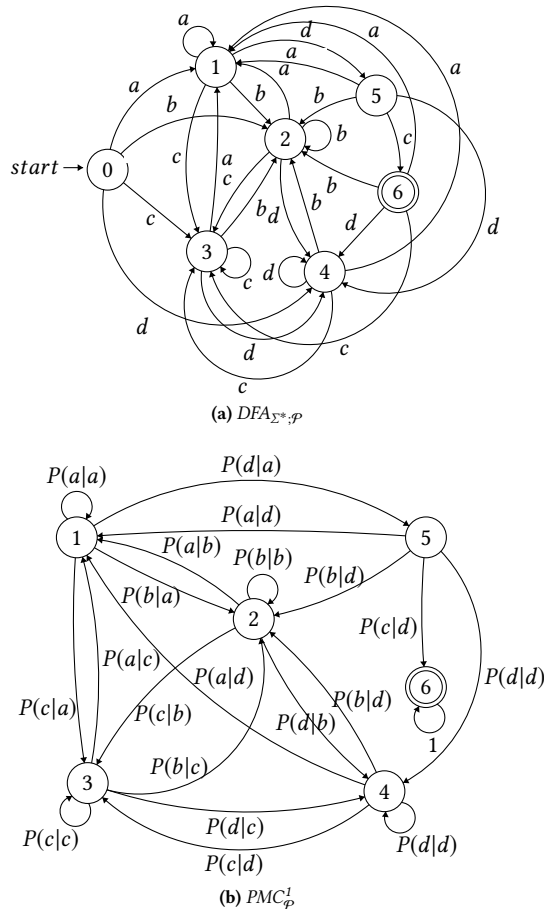


Figure 1: DFA and PMC for $\mathcal{P} = a; d; c$ with $\Sigma = \{a, b, c, d\}$, and order $m = 1$.

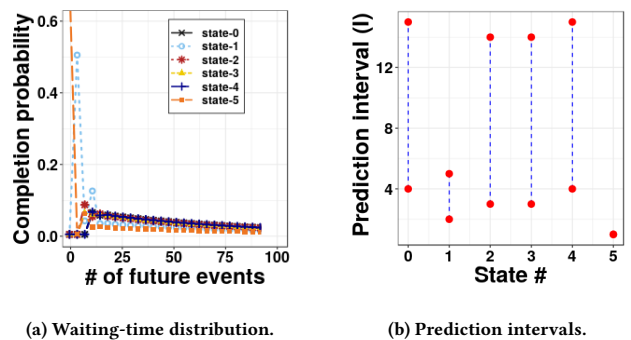


Figure 2: Example of how prediction intervals are produced. $\mathcal{P} = a; d; c$, $\Sigma = \{a, b, c, d\}$, $m = 1$, $\theta_{fc} = 0.5$.

part of the stream to obtain the maximum-likelihood estimators for the transition probabilities [4]. If Π is the transition matrix of a Markov chain with a set of states Q , $\pi_{i,j}$ the transition probability from state i to state j , $n_{i,j}$ the number of observed transitions from state i to state j , then the maximum likelihood estimator

for $\pi_{i,j}$ is given by:

$$\hat{\pi}_{i,j} = \frac{n_{i,j}}{\sum_{k \in Q} n_{i,k}} = \frac{n_{i,j}}{n_i}$$

Executing this learning step on a single node might require a vast amount of time until we arrive at a sufficiently good model. In this paper, we present a distributed method for learning the transition probability matrix.

3.2 Pattern prediction on multiple streams

3.2.1 Problem formulation. Let $O = \{o_1, \dots, o_k\}$ be a set of K objects (i.e., moving objects) and $S = \{s_1, \dots, s_k\}$ a set of real-time streams of events, where s_i is generated by the object o_i . Let \mathcal{P} be a user-defined pattern which we want to apply to every stream s_i , i.e., each object will have its own DFA.

The setting that is considered in this work is then described in the following: we have K input event streams S and a system consisting of K distributed predictor nodes n_1, n_2, \dots, n_k , each of which consumes an input event stream $s_i \in S$. The goal is to provide timely predictions and be able to do this at large-scale. Each node n_i handles a single event stream s_i associated with a moving object $o_i \in O$. In addition, it maintains a local prediction model f_i for the user-defined pattern \mathcal{P} . The f_i model provides the online prediction about the future full match of the pattern \mathcal{P} in s_i for each new arriving event tuple.

In short, we have multiple running instances of an online prediction algorithm on distributed nodes for multiple input event streams. More specifically, the input to our system consists of massive streams of events that describe trajectories of moving vessels in the context of maritime surveillance, where there is one predictor node for each vessel's event stream.

3.2.2 The proposed approach. We designed and developed a scalable and distributed pattern prediction system over a massive input event streams of moving objects. As the base prediction model, we use the PMC forecasting method [2]. Moreover, we propose to enable the information exchange between the distributed predictors/learners of the input event streams, by adapting the distributed online prediction protocol of [16] to synchronize the prediction models, i.e., the transitions probabilities matrix of the PMC predictors.

Algorithm 1 presents the distributed online prediction protocol by dynamic model synchronization on both the predictor nodes and the coordinator. We refer to the PMC's transition matrix Π_i on predictor node n_i by f_i . That is, when a predictor $n_i : i \in [k]$ observes an event e_j it revises its internal model state (i.e., f_i) and provides a prediction report. Then it checks the local conditions (batch size b and local model divergence from a reference model f_r) to decide whether there is a need to synchronize its local model with the coordinator [or not]. f_r is maintained in the predictor node as a copy of the last computed aggregated model \hat{f} from the previous full synchronization step, which is shared between all local predictors/learners. By monitoring the local condition $\|f_i - f_r\|^2 > \Delta$ on all local predictors, we have a guarantee that if none of the local conditions is violated, the divergence (i.e., variance of local models $\delta(f) = \frac{1}{k} \sum_{j=1}^k \|f_i - \hat{f}\|^2$) does not exceed the threshold Δ [16].

On the other hand, the coordinator receives the prediction models from the predictor nodes that requested for model synchronization (violation). Then it tries to keep incrementally querying other nodes for their local prediction models until reaching out all nodes, or the variance of the aggregated model \hat{f} that is

computed from the already received models less or equal than the divergence threshold Δ . Finally, the aggregated model \hat{f} is sent back to the predictor nodes that sent their models after the violation or have been queried by the coordinator.

Algorithm 1: Communication-efficient Distributed Online Learning Protocol

Predictor node n_i : at observing event e_j
 update the prediction model parameters f_i and provide a prediction service ;
if $j \bmod b = 0$ and $\|f_i - f_r\|^2 > \Delta$ **then**
 send f_i to the Coordinator (violation) ;

Coordinator:

receive local models with violation $B = \{f_i\}_{i=1}^m$;
while $|B| \neq k$ and $\frac{1}{|B|} \sum_{f_i \in B} \|f_i - \hat{f}\|^2 > \Delta$ **do**
 add other nodes that have not reported violation for their models $B \leftarrow \{f_l : f_l \notin B \text{ and } l \in [k]\}$;
 receive models from nodes in B ;
 compute a new global model \hat{f} ;
 send \hat{f} to all the predictors in B and set $f_1 \dots f_m = \hat{f}$;
if $|B| = k$ **then**
 set a new reference model $f_r \leftarrow \hat{f}$;

This protocol was introduced for linear models, and has been extended to handle kernelized online learning models [17]. We also employ this protocol for the pattern prediction model, which is internally based on the PMC $PMC_{\mathcal{P}}^m$. This allows the distributed $PMC_{\mathcal{P}}^m$ predictors for multiple event streams to synchronize their models (i.e., the transition probability matrix of each predictor) within the system in a communication-efficient manner.

We propose a *synchronization operation* for the parameters of the models ($f_i = \Pi_i : i \in [k]$) of the k distributed PMC predictors. The operation is based on distributing the maximum-likelihood estimation [4] for the transition probabilities of the underlying $PMC_{\mathcal{P}}^m$ models described by:

$$\hat{\pi}_{i,j} = \frac{\sum_{k \in K} n_{k,i,j}}{\sum_{k \in K} \sum_{l \in L} n_{k,i,l}}$$

Moreover, we measure the divergence of local models from the reference model $\|f_k - f_r\|^2$ by calculating the sum of square difference between the transition probabilities Π_i and Π_r :

$$\|f_k - f_r\|^2 = \sum_{i,j} (\hat{\pi}_{k,i,j} - \hat{\pi}_{r,i,j})^2$$

In general, our approach relies on enabling the collaborative learning among the distributed predictors. Each predictor node receives a stream of events related to a distinct moving object, and the central coordinator is responsible for synchronizing their prediction models using the *synchronization operation*. Moreover, the predictors they only need to share the parameters of their models, not the consumed event streams.

We assume that the underlying event streams belong to the same distribution and share the same behavior (e.g., mobility patterns). We claim this assumption is reasonable in many application domains: for instance, in the context of maritime surveillance, vessels travel through standard routes, defined by the International Maritime Organization (IMO). Additionally, vessels have similar mobility patterns in specific areas such as moving

with low speed and multiple turns near the ports [20, 29]. That allows our system to construct a coherent global prediction model dynamically for all input event streams based on merging their local prediction models.

3.3 Distributed architecture

Our system consumes as an input² an aggregated stream of events coming from a large number of moving objects, which is continuously collected and fed into the system. It allows users to register a pattern \mathcal{P} to be monitored over each event stream of a moving object. The output stream consists of original input events and predictions of full matches of \mathcal{P} , displayed to the end users. Figure 3 presents the overview of our system architecture and its main components.

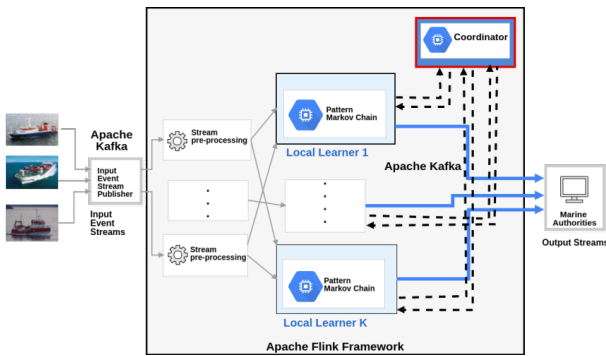


Figure 3: System Architecture.

The system is composed of three processing units: (i) pre-processing operators that receive the input event stream and perform filtering and ordering operations, before partitioning the input event stream to multiple event streams based on the associated moving object (ii) predictor nodes (learners), which are responsible for maintaining a prediction model for the input event streams. Each prediction node is configured to handle an event stream from the same moving object, in order to provide online predictions for a predefined pattern \mathcal{P} (iii) a coordinator node that communicates through Kafka stream channels with the predictors to realize the distributed online learning protocol. It builds a global prediction model, based on the received local models, and then shares it among the predictors.

Our distributed system consists of multiple pre-processing operators, prediction nodes, and a central coordinator node. These units run concurrently and are arranged as a data processing pipeline, depicted in Figure 3. We leverage Apache Kafka as a messaging platform to ingest the input event streams and to publish the resulting streams. Also, it is used as the communication channel between the predictor nodes and the coordinator. Apache Flink is employed to execute the system’s distributed processing units over the input event streams: the pre-processing operators, the prediction units, and the coordinator node. Our system architecture can be modeled as a logical network of processing nodes, organized in the form of a DAG, inspired by the Flink runtime dataflow programs [6].

²In practice, the aggregated input events stream is composed of multiple event streams (partitions) from a set of moving objects.

4 IMPLEMENTATION DETAILS

In this section, we briefly describe in detail the implementation of our system on top of Apache Flink and Apache Kafka frameworks. Each of the three sub-modules, described in Section 3.3, have been implemented as Flink operations over the Kafka events stream.

Pre-processing and Prediction Operators. Listing 1 shows how the main workflow of the system is implemented as Flink data flow program.

The system ingests the input events stream from a Kafka cluster that is mapped to a *DataStream* of events, which is then processed by an *EventTuplesMapper* to create tuples of $(id, event)$, where the *id* is associated to the identifier of the moving object. To handle events coming in out of order in a certain margin, the stream of event tuples is processed by a *TimestampAssigner*, it assigns the timestamps for the input events based on the extracted creation time. Afterwards, an ordered stream of event tuples is generated using a process function *EventSorter*.

```
DataStream<Event> eventsStream =
    env.addSource(kafkaConsumer);
// Create tuples (id,event) and assign time stamps
DataStream<Tuple2<String,Event>> eventTuplesStream =
    inputEventsStream.map(new EventTuplesMapper())
    .assignTimestampsAndWatermarks(new
        EventTimeAssigner());
// Create the ordered keyed stream
orderedEventsStream =
    eventsStream.keyBy(0).process(new
        EventSorter()).keyBy(0);
// Consume the events by the predictors
LocalPredictorNode predictorNode =new
    LocalPredictorNode<Event>(P);
DataStream<Event> processedEventsStream =
    orderedEventsStream.map(predictorNode);
```

Listing 1: Flink pipeline for local predictors workflow

The ordered stream is then transformed to a *keyedEventsStream* by partitioning it, based on the ids values, using a *keyBy* operation. A local *predictor* node in a distributed environment is represented by a *map* function over the *keyedEventsStream*. Each parallel instance of the map operator (predictor) always processes all events of the same moving object (i.e., equivalent id), and maintains a bounded prediction model (i.e., $PMC_{\mathcal{P}}^m$ predictor) using the Flink’s Keyed State³. The output streams of the moving objects from the parallel instances of the predictor map functions are sent to a new Kafka stream (i.e., same topic name). They then can be processed by other components like visualization or users notifier.

Moreover, the implementation of the *predictor* map function includes the communication with *coordinator* using Kafka streams. At the beginning of the execution, it sends a registration request to the coordinator. Also at the run-time, it sends its local prediction model as synchronization request, or as a response for a resolution request from the coordinator. These communication messages are published into different Kafka topics as depicted in Table 1.

Coordinator. It manages the distributed online learning protocol operations, which is also implemented as Flink program. The coordinator receives messages from the local predictors

³Keyed State in Flink: <https://ci.apache.org/projects/flink/flink-docs-release-1.3/dev/stream/state.html#keyed-state>

Table 1: Messages to Kafka topics mapping.

| Message | Kafka Topic |
|--|---------------------------|
| <i>RegisterNode</i> , <i>RequestSync</i> , and <i>ResolutionAnswer</i> | LocalToCoordinatorTopicId |
| <i>CoordinatorSync</i> and <i>RequestResolution</i> | CoordinatorToLocalTopicId |

through a Kafka Stream of a topic named "*LocalToCoordinator-TopicId*". It is implemented as a single *map* function over the messages stream, by setting the *parallelism* level of the Flink program to "1". Increasing the parallelism will scale up the number of parallel coordinator instances, for example, in order to handle different groupings of the input event streams. The map operator of the coordinator handles three message types from the predictors: (i) **RegisterNode** that contains a registration request for a new predictor node, (ii) **RequestSync** to receive a local model after violation, (iii) **ResolutionAnswer** to receive a resolution response from a local predictor node. In addition, it sends **CoordinatorSync** messages for all predictors after creating a new global prediction model, or **RequestResolution** to ask the local predictors for their prediction models.

5 EMPIRICAL EVALUATION

In this section, we evaluate our proposed system by analyzing the predictive performance and communication complexity using real-world event streams provided by the datAcron project in the context of maritime monitoring. The used event streams describe critical points (i.e., synopses) of moving vessels trajectories, which are derived from raw AIS messages as described in [30]. In particular, for our evaluation experiments we used a data set of synopses that contains 4,684,444 critical points of 5055 vessels sailing in the Atlantic Ocean during the period from 1 October 2015 to 31 March 2016.

We used the synopses data set to generate a simulated stream of event tuples i.e., (*id*, *timestamp*, *longitude*, *latitude*, *annotation*, *speed*, *heading*), which are processed by the system to attach an extra attribute *type* that represents the event value, where $type \in \Sigma$, and $\Sigma = \Sigma_1 = \{VerySlow, Slow, Moving, Sailing, Stopping\}$, which is based on a discretization of the speed values. That is, Σ_1 includes a simple derived event types based on the speed value that can be used over streams of raw AIS or critical points. Or $\Sigma = \Sigma_2 = \{stopStart, stopEnd, changeInSpeedStart, changeInSpeedEnd, slowMotionStart, slowMotionEnd, gapStart, gapEnd, changeInHeading\}$, which is derived based on the values of the *annotation* attribute that encodes the extracted trajectory movement events [30]. Σ_2 represents the set of possible mobility changes in the vessel's trajectory [30], each critical point has at least one event. Where in the case of multiple values, we generate duplicate points each of which corresponding to one event in the same order of Σ_2 .

In our experiments, we monitor a pattern $\mathcal{P}_1 = Sailing$ with Σ_1 that detects when the vessel is underway (sailing). Likewise, we test a second pattern $\mathcal{P}_2 = changeInHeading; gapStart; gapEnd; changeInHeading$ with Σ_2 that describes a potential illegal fishing activity [2].

Experimental setup. We ran our experiments on single-node standalone Flink cluster deployed on an Ubuntu Server 17.04 with Intel(R) Core(TM) i7-7700 CPU @ 3.60GHz X 8 processors

and 32GB RAM. We used Apache Flink v1.3.2 and Apache Kafka v0.10.2.1 for our tests.

Evaluation criteria. Our goal is to evaluate our distributed pattern prediction system, which enables the synchronization of prediction models (i.e., PMC models) on the distributed predictor nodes. Our proposed system can operate in three different modes of protocols/schemes of models synchronization: (i) static scheme based on synchronizing the prediction models periodically every b of input events in each stream, (ii) continuous, full synchronization for each incoming event (hypothetical), (iii) dynamic synchronization protocol based on making the predictors communicate their local prediction models periodically but only under condition that the divergence of the local models from a reference model exceeds a variance threshold Δ (recommended).

We compare our proposed system against the isolated prediction mode, in which models are computed on single streams only, and compare the predictive performance in terms of :

- (i) $Precision = \frac{\# \text{ of correct predictions}}{\# \text{ of total predictions}}$ is the fraction of the produced predictions that are correct. For each new event in the stream, the predictor provides a prediction interval where the full match of the pattern might occur. Thus, the predictions are temporarily stored until a full match is detected. At that point, all stored prediction intervals are evaluated by considering those intervals where the full match occurred within as correct.
- (ii) $Spread = end(I) - start(I)$ is the width of the prediction interval I , which represents the number of events between the start and the end of I .

Moreover, we study the communication cost by measuring the *cumulative communication* that captures the number of messages, which are required to perform the distributed online learning modes to synchronize the prediction models. Next, we present the experimental results for the patterns $\mathcal{P}_1 = Sailing$ with an order of $m = 2$, and $\mathcal{P}_2 = changeInHeading; gapStart; gapEnd; changeInHeading$ with first order $m = 1$. All experiments are performed with setting the batch size to 100 ($b = 100$), the variance threshold of 2 ($\Delta = 2$), 80% as PMC prediction threshold ($\theta_{fc} = 80\%$), and 200 for the maximum spread.

Experimental results. Figure 4 depicts the average precision scores of predictions models (one prediction model per vessel) of all synchronization modes for the first pattern $\mathcal{P}_1 = Sailing$, namely, isolated without synchronization, continuous (full-sync), static, and our recommended approach based on the dynamic synchronization scheme. It can be clearly seen that all methods of distributed learning outperform the isolated prediction models. The hypothetical method of full continuous synchronization has the highest precision rates, while the static and dynamic synchronization schemes have close precision scores. Consequently, dynamic synchronization is not much weaker than the static synchronization, but requires much less communication, as explained below.

Figure 5 provides the amount of the accumulated communication that is required by the three modes of the distributed online learning, while the isolated approach does not require any communication between the predictors. These results are shown for \mathcal{P}_1 . As expected, a larger amount of communication is required for the continuous synchronization comparing to the static and dynamic approaches. Also, it can be seen that we can reduce the communication overhead by applying the dynamic synchronization protocol (a reduction by a factor of 100) comparing to the

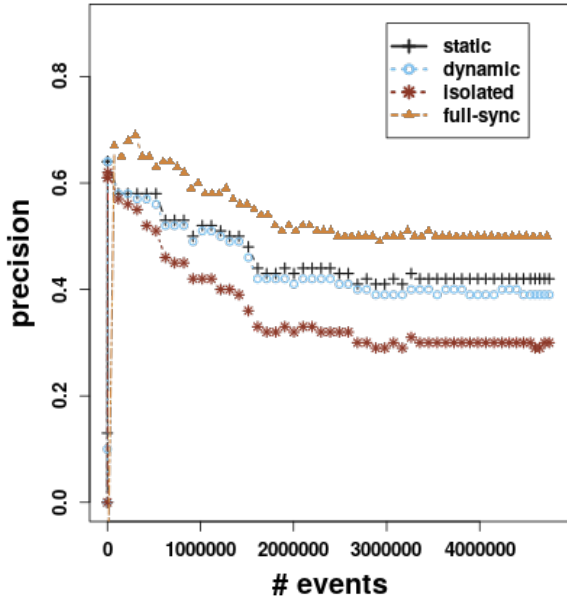


Figure 4: Precision scores with respect to the number of input events over time for \mathcal{P}_1 .

static synchronization scheme, even with a small variance threshold $\Delta = 2$. Furthermore, the dynamic protocol is still preserving a close predictive performance to the static one (see Figure 4). Therefore, we will only consider the dynamic synchronization and the isolated approach in the evaluation of the second pattern.

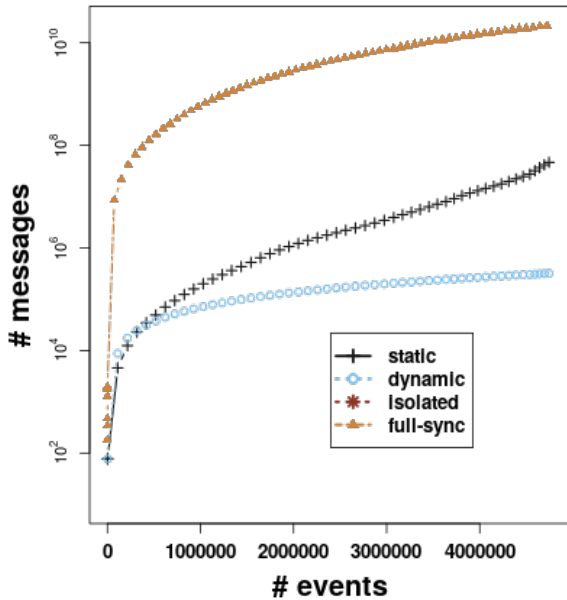


Figure 5: Cumulative communication with respect to the number of input events over time for \mathcal{P}_1 .

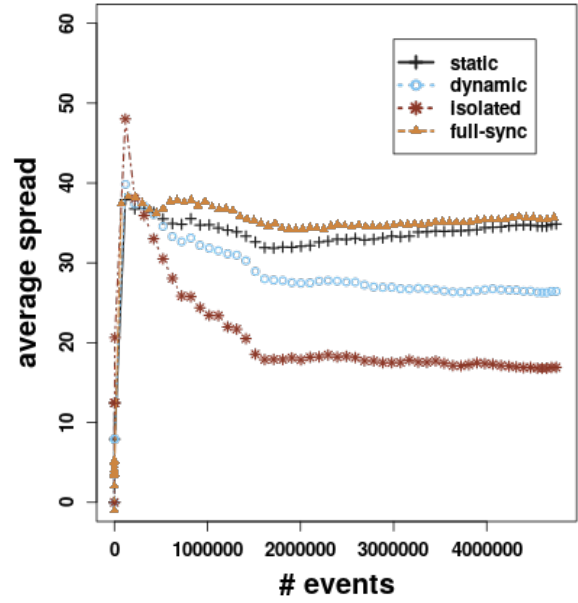


Figure 6: Average spread value for \mathcal{P}_1 .

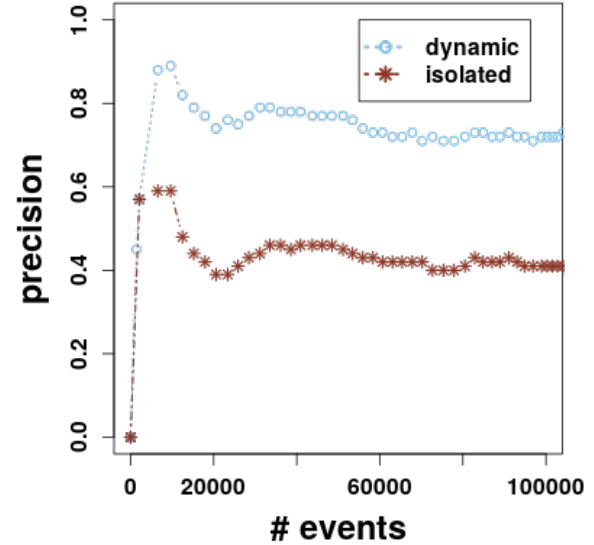


Figure 7: Precision scores of \mathcal{P}_2 for *PLEASURE CRAFT* vessels.

In Figure 4, we also noted that the precision is going down in a first phase and stabilizes then. This seems to be counter-intuitive, as the models should improve when getting more data up to a certain point. For explanation, we have investigated the effect of the distributed synchronization of the prediction models on the average spread value, Figure 6 shows the spread results for all approaches. It can be seen that the spread is higher for the distributed learning based methods comparing to the isolated approach. Furthermore, the average spread is decreasing over time until convergence, as result of confidence increase in the models. This may explain the drop in the precision scores from

the beginning until reaching the convergence. We will investigate further in the interrelation between precision and spread in future work.

For the second, more complex pattern (\mathcal{P}_2), we have found that the precision was worse for a distributed model generated over all vessels than in the model created for each vessel in isolation. This indicates that there is no global model describing the behavior of all models consistently. However, when looking at specific groups of vessels, we achieved an improvement in terms of precision. As initial experiment, we only enable the synchronization of the prediction models associated with vessels that belong to the same vessel class. Currently, this change is technically performed by an extra filter step that passes only one type of vessels, while multiple runs of the system are required for all vessel types. For example, Figure 7 shows the precision scores for vessels of class *PLEASURE CRAFT*. An interesting observation is that the dynamic synchronization approach still has higher precision scores than the isolated approach. This case might seem to contradict of our assumption that the input event streams belong to the same distribution and share the same behavior, but it actually follows the same assumption but between the predictors of vessels within the same type group. We will further investigate the effect of groupings and more patterns in future work.

6 CONCLUSION

In this paper, we have presented a system that provides a distributed pattern prediction over multiple large-scale event streams of moving objects (vessels). The system uses the event forecasting with Pattern Markov Chain (PMC) [2] as the base prediction model on each event stream, and it applies the protocol for distributed online prediction [16] to exchange information between the prediction models over multiple input event streams. Our proposed system has been implemented using Apache Flink and Apache Kafka. In order to show the usefulness and effectiveness of our approach, we empirically tested it against large real-world event streams related to trajectories of moving vessels.

As future work, we will address the open issues emerging from the current findings. Firstly, we will study the interrelation between precision and spread scores by validating the approach over synthetic event streams. Secondly, we will investigate the effect of grouping the input event streams on the predictive performance of our method.

ACKNOWLEDGMENTS

This work was supported by EU Horizon 2020 datAcron project (grant agreement No 687591).

REFERENCES

- [1] Rakesh Agrawal, Tomasz Imieliński, and Arun Swami. 1993. Mining Association Rules Between Sets of Items in Large Databases. In *ACM SIGMOD*.
- [2] Elias Alevizos, Alexander Artikis, and George Paliouras. 2017. Event Forecasting with Pattern Markov Chains. In *Proceedings of the 11th ACM International Conference on Distributed and Event-based Systems*. ACM, 146–157.
- [3] Elias Alevizos, Anastasios Skarlatidis, Alexander Artikis, and Georgios Paliouras. 2015. Complex event recognition under uncertainty: A short survey. *Event Processing, Forecasting and Decision-Making in the Big Data Era (EPForDM)* (2015), 97–103.
- [4] Theodore W Anderson and Leo A Goodman. 1957. Statistical inference about Markov chains. *The Annals of Mathematical Statistics* (1957), 89–110.
- [5] Amazon Web Services (AWS). 2013. Amazon Kinesis. <https://aws.amazon.com/de/kinesis/>. (2013).
- [6] Paris Carbone, Asterios Katsifodimos, Stephan Ewen, Volker Markl, Seif Haridi, and Kostas Tzoumas. 2015. Apache flink: Stream and batch processing in a single engine. *Bulletin of the IEEE Computer Society Technical Committee on Data Engineering* 36, 4 (2015).
- [7] Gianpaolo Cugola and Alessandro Margara. 2012. Processing Flows of Information: From Data Stream to Complex Event Processing. *ACM Comput. Surv.* 44, 3, Article 15 (June 2012), 62 pages. <https://doi.org/10.1145/2187671.2187677>
- [8] Ofer Dekel, Ran Gilad-Bachrach, Ohad Shamir, and Lin Xiao. 2012. Optimal distributed online prediction using mini-batches. *Journal of Machine Learning Research* 13, Jan (2012), 165–202.
- [9] Lina Fahed, Armelle Brun, and Anne Boyer. 2014. Efficient Discovery of Episode Rules with a Minimal Antecedent and a Distant Consequent. In *Knowledge Discovery, Knowledge Engineering and Knowledge Management*. Springer.
- [10] Ioannis Flouris, Vasiliki Manikaki, Nikos Giatrakos, Antonios Deligiannakis, Minos Garofalakis, Michael Mock, Sebastian Bothe, Inna Skarbovska, Fabiana Fournier, Marko Stajcer, et al. 2016. FERARI: A Prototype for Complex Event Processing over Streaming Multi-cloud Platforms. In *Proceedings of the 2016 International Conference on Management of Data*. ACM, 2093–2096.
- [11] The Apache Software Foundation. 2012. Apache Kafka. <https://kafka.apache.org/>. (2012).
- [12] The Apache Software Foundation. 2013. Apache Spark Streaming. <http://spark.apache.org/streaming/>. (2013).
- [13] The Apache Software Foundation. 2014. Apache Flink. <https://flink.apache.org/>. (2014).
- [14] The Apache Software Foundation. 2014. Apache Storm. <http://storm.apache.org/>. (2014).
- [15] John E Hopcroft, Rajeev Motwani, and Jeffrey D Ullman. 2006. Automata theory, languages, and computation. *International Edition* 24 (2006).
- [16] Michael Kamp, Mario Boley, Daniel Keren, Assaf Schuster, and Izchak Sharfman. 2014. Communication-efficient distributed online prediction by dynamic model synchronization. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*. Springer, 623–639.
- [17] Michael Kamp, Sebastian Bothe, Mario Boley, and Michael Mock. 2016. Communication-Efficient Distributed Online Learning with Kernels. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*. Springer, 805–819.
- [18] John Langford, Alex J Smola, and Martin Zinkevich. 2009. Slow learners are fast. *Advances in Neural Information Processing Systems* 22 (2009), 2331–2339.
- [19] Srivatsan Laxman, Vikram Tankasali, and Ryan W. White. 2008. Stream Prediction Using a Generative Model Based on Frequent Episodes in Event Sequences. In *ACM SIGKDD*.
- [20] Bo Liu, Erico N de Souza, Stan Matwin, and Marcin Sydow. 2014. Knowledge-based clustering of ship trajectories using density-based approach. In *Big Data (Big Data), 2014 IEEE International Conference on*. IEEE, 603–608.
- [21] David Luckham. 2008. The power of events: An introduction to complex event processing in distributed enterprise systems. In *International Workshop on Rules and Rule Markup Languages for the Semantic Web*. Springer, 3–3.
- [22] Heikki Mannila, Hannu Toivonen, and A. Inkeri Verkamo. 1997. Discovery of Frequent Episodes in Event Sequences. *Data Mining and Knowledge Discovery* (1997).
- [23] Michael Mathioudakis and Nick Koudas. 2010. Twittermonitor: trend detection over the twitter stream. In *Proceedings of the 2010 ACM SIGMOD International Conference on Management of data*. ACM, 1155–1158.
- [24] Daniele Miorandi, Sabrina Sicari, Francesco De Pellegrini, and Imrich Chlamtac. 2012. Internet of things: Vision, applications and research challenges. *Ad Hoc Networks* 10, 7 (2012), 1497–1516.
- [25] Douglas C Montgomery, Cheryl L Jennings, and Murat Kulahci. 2015. *Introduction to time series analysis and forecasting*. Wiley.
- [26] Vinod Muthusamy, Haifeng Liu, and Hans-Arno Jacobsen. 2010. Predictive Publish/Subscribe Matching. In *DEBS*. ACM.
- [27] Grégory Nuel. 2008. Pattern Markov Chains: Optimal Markov Chain Embedding through Deterministic Finite Automata. *Journal of Applied Probability* (2008).
- [28] International Maritime Organization. 2001. Automatic identification systems. <http://www.imo.org/OurWork/Safety/Navigation/Pages/AIS.aspx>. (2001).
- [29] Giuliana Pallotta, Michele Vespe, and Karna Bryan. 2013. Vessel pattern knowledge discovery from AIS data: A framework for anomaly detection and route prediction. *Entropy* 15, 6 (2013), 2218–2245.
- [30] Kostas Patroumpas, Elias Alevizos, Alexander Artikis, Marios Vodas, Nikos Pelekis, and Yannis Theodoridis. 2017. Online event recognition from moving vessel trajectories. *GeoInformatica* 21, 2 (2017), 389–427.
- [31] Kostas Patroumpas, Alexander Artikis, Nikos Katsouris, Marios Vodas, Yannis Theodoridis, and Nikos Pelekis. 2015. Event Recognition for Maritime Surveillance. In *EDBT*. 629–640.
- [32] R. Vilalta and Sheng Ma. 2002. Predicting rare events in temporal domains. In *ICDM*.
- [33] Lin Xiao. 2010. Dual averaging methods for regularized stochastic learning and online optimization. *Journal of Machine Learning Research* 11, Oct (2010), 2543–2596.
- [34] Feng Yan, Shreyas Sundaram, SVN Vishwanathan, and Yuan Qi. 2013. Distributed autonomous online learning: Regrets and intrinsic privacy-preserving properties. *IEEE Transactions on Knowledge and Data Engineering* 25, 11 (2013), 2483–2493.
- [35] Cheng Zhou, Boris Cule, and Bart Goethals. 2015. A pattern based predictor for event streams. *Expert Systems with Applications* (2015).