(c) Authors|Springer 2022. This is the authors' version of the work. It is posted here for your personal use only. Not for redistribution. The definitive version of the work is published with Springer, DOI: https://doi.org/10.1007/978-3-030-78307-5_1

Processing Big Data in Motion: Core Components and System Architectures with Applications to the Maritime Domain

Nikos Giatrakos, Antonios Deligiannakis, Konstantina Bereta, Marios Vodas, 5 Dimitris Zissis, Elias Alevizos, Charilaos Akasiadis, and Alexander Artikis 6

Abstract Rapidly extracting business value out of Big Data that stream in corporate 7 data centres requires continuous analysis of massive, high-speed data while they are 8 still in motion. So challenging a goal entails that analytics should be performed 9 in memory with a single pass over these data. In this chapter, we outline the 10 challenges of Big streaming Data analysis for deriving real-time, online answers to 11 application inquiries. We review approaches, architectures and systems designed to 12 address these challenges and report on our own progress within the scope of the EU 13 H2020 project INFORE. We showcase INFORE into a real-world use case from the 14 maritime domain and further discuss future research and development directions. 15

KeywordsBig Data · Cross-platform optimisation · Data streams · Data16synopses · Online machine learning · Complex event forecasting · Maritime17situation awareness18

N. Giatrakos (⊠) · A. Deligiannakis Athena Research Center, Athens, Greece e-mail: ngiatrakos@athenarc.gr; adeli@athenarc.gr

K. Bereta · M. Vodas · D. Zissis MarineTraffic, Athens, Greece e-mail: konstantina.bereta@marinetraffic.com; marios.vodas@marinetraffic.com; dzissis@marinetraffic.com

E. Alevizos · C. Akasiadis · A. Artikis NCSR Demokritos, Institute of Informatics and Telecommunications, Athens, Greece e-mail: alevizos.elias@iit.demokritos.gr; cakasiadis@iit.demokritos.gr; a.artikis@iit.demokritos.gr 1

2

3

1 Challenges of Big Streaming Data

Today, organisations and businesses have the ability to collect, store and analyse as ²⁰ much data as they need, exploiting powerful computing machines in corporate data ²¹ centres or the cloud. To extract value out of the raw Big Data that are accumulated, ²² application workflows are designed and executed over these infrastructures engaging simpler (such as grouping and aggregations) or more complex (data mining and ²⁴ machine learning) analytics tasks. These tasks may involve data at rest or data in ²⁵ motion. ²⁶

Data at rest are historic data stored on disks, getting retrieved and loaded for ²⁷ processing by some analytics workflow. Analytics tasks participating in such a ²⁸ workflow perform computations on massive amounts of data, lasting for hours or ²⁹ days. They finally deliver useful outcomes. Using a running example from the ³⁰ maritime domain, historic vessel position data are used to extract Patterns-of-Life ³¹ (PoL) information. These are essentially collections of geometries representing ³² normal navigational routes of vessels in various sea areas [78], used as the basis ³³ for judging anomalies. ³⁴

Data in motion involve Big streaming Data which are unbounded, high-speed 35 streams of data that need to get continuously analysed in an online, real-time 36 fashion. Storing the data in permanent storage is not an option, since the I/O latency 37 would prevent the real-time delivery of the analytics output. Application workflows 38 get a single look on the streaming data tuples, which are kept in memory for a 39 short period of time and are soon stored or discarded to process newly received data 40 tuples. 41

At an increasing rate, numerous industrial and scientific institutions face such ⁴² business requirements for real-time, online analytics so as to derive actionable ⁴³ items and timely support decision-making procedures. For instance, in the maritime ⁴⁴ domain, to pinpoint potentially illegal activities at sea [54] and allow the authorities ⁴⁵ to timely act, position streams of thousands of vessels need to be analysed online. ⁴⁶

To handle the volume and velocity of Big streaming Data, Big Data platforms ⁴⁷ such as Apache Flink [2], Spark [5] or toolkits like Akka [1] have been designed to ⁴⁸ facilitate scaling-out, i.e., parallelising, the computation of streaming analytics tasks ⁴⁹ horizontally to a number of Virtual Machines (VM) available in corporate computer ⁵⁰ clusters or the cloud. Thus, multiple VMs simultaneously execute analytics on ⁵¹ portions of the streaming data undertaking part of the processing load, and therefore ⁵² throughput, i.e., number of tuples being processed per time unit, is increased. This ⁵³ aids in transforming raw data in motion to useful results delivered in real time. ⁵⁴ Big Data platforms also offer APIs with basic stream transformation operators such ⁵⁵ as filter, join, attribute selection, among others, to program and execute streaming ⁵⁶ workflows. However useful these facilities may be, they only focus on a narrow part ⁵⁷ of the challenges that business workflows need to encounter in streaming settings. ⁵⁸

First, Big Data platforms currently provide none or suboptimal support for ⁵⁹ advanced streaming analytics tasks engaging Machine Learning (ML) or Data ⁶⁰ Mining (DM) operators. The major dedicated ML/DM APIs they provide, such as ⁶¹

AO1

MLlib [5] or FlinkML [2], do not focus on parallel implementations of streaming 62 algorithms. 63

Second, Big Data platforms by design focus only on horizontal scalability as 64 described above, while there are two additional types of scalability that are of 65 essence in streaming settings. Vertical scalability, i.e., scaling the computation with 66 the number of processed streams, is also a necessity. Federated scalability, i.e., 67 scaling the computation one step further out, to settings composed of multiple, 68 potentially geo-dispersed computer clusters, is another type of required scalability. 69 For instance, in maritime applications, vessels transmit their positions to satellite or 70 ground-based receivers. These data can be ingested in proximate data centres and 71 communicated only on demand upon executing global workflows, i.e., involving the 72 entire set of monitored vessels, over the fragmented set of streams.

Third, Big Data technologies are significantly fragmented. Delivering advanced 74 analytics requires optimising the execution of workflows over a variety of Big Data 75 platforms and tools located at a number of potentially geo-dispersed clusters or 76 clouds [30, 34, 36]. In such cases, the challenge is to automate the selection of 77 an optimal setup prescribing (a) which network cluster will execute each analytics 78 operator, (b) which Big Data platform available at this cluster, and (c) how to 79 distribute the computing resources of that cluster to the operators that are assigned 80 to it.

Connecting the above challenges to a real-world setting from the maritime ⁸² domain, on a typical day at MarineTraffic,¹ 100GB vessel position data and approx-⁸³ imately 750M messages (volume, velocity—horizontal scalability) are processed ⁸⁴ online. This data is complemented by other data sources such as satellite image data ⁸⁵ of tens of TBs [54]. At any given time, MarineTraffic is tracking over 200K vessels ⁸⁶ in real-time (vertical scalability) over a network of approximately 5K stations ⁸⁷ (federated scalability). Additionally, the analysis engages a variety of Big Data ⁸⁸ platforms including Apache Spark, Flink, Akka and Kafka (details in Sect. 3).

Finally, applications often require an additional level of abstraction on the derived ⁹⁰ analytics results. Consider a vessel that slows down, then makes a U-turn and then ⁹¹ starts speeding up. Such a behaviour may occur in case of an imminent piracy ⁹² event where a vessel attempts to run away from pirates. The application is not ⁹³ interested in knowing the absolute speed, heading or direction information in the ⁹⁴ raw stream. Instead, it wants to receive continuous reports directly on a series of ⁹⁵ detected, *simple* events (slowing down, U-turn, speeding) and the higher ⁹⁶ level, *complex* piracy event or to be able to forecast such events [79]. Complex ⁹⁷ Event Processing (CEP) and Forecasting (CEF) encompass the ability to query for ⁹⁸ business rules (patterns) that match incoming streams on the basis of their content ⁹⁹ and some topological ordering on them (CEP) or to forecast the appearance of ¹⁰⁰ patterns (CEF) [31, 33, 35].

In this chapter, we discuss core system components required to tackle these 102 challenges and the state of the art in their internal architectures. We further describe 103

¹ https://www.marinetraffic.com.

how we advance the state of the art within the scope of the EU H2020 project 104 INFORE. Finally, we showcase the INFORE approach into a real-world use case 105 from the maritime domain. We, however, stress that INFORE applies to any 106 application domain, and we refer the interested reader to [34] for more application 107 scenarios. 108

This chapter relates to the technical priorities (a) Data Management, (b) Data 109 Processing Architectures and (c) Data Analytics of the European Big Data Value 110 Strategic Research & Innovation Agenda [77]. It addresses the horizontal concerns 111 Cloud, HPC and Sensor/Actuator infrastructure of the BDV Technical Reference 112 Model and the vertical concern of Big Data Types and Semantics (Structured data, 113 Time series data, Geospatial data). Moreover, the chapter relates to (a) Knowledge 114 and Learning, (b) Reasoning and Decision Making, (c) Action and Interaction and 115 (d) Systems, Hardware, Methods and Tools, cross-sectorial technology enablers 116 of the AI, Data and Robotics Strategic Research, Innovation and Deployment 117 Agenda [76]. 118

2 Core Components and System Architectures

2.1 The Case for Data Synopses

Motivation There is a wide consensus in the stream processing community [25, 121] 26, 32] that approximate but rapid answers to analytics tasks, more often than not, 122 suffice. For instance, detecting a group of approximately 50 highly similar vessel tra- 123 jectories with sub-second latency is more important than knowing minutes later that 124 the group actually composes 55 such streams with a similarity value accurate to the 125 last decimal. In the latter case, some vessels may have been engaged in a collision. 126 Data synopses techniques such as samples, histograms and sketches constitute a 127 powerful arsenal of data summarisation tools useful across the challenges discussed 128 in the introduction of this chapter. Approximate, with tunable quality guarantees, 129 synopses operators including, but not limited to [25, 26, 32, 46], cardinality (FM 130 Sketches), frequency moment (CountMin, AMS Sketches, Sampling), correlation 131 (Fourier Transforms, Locality Sensitive Hashing [37]), set membership (Bloom 132 Filters) or quantile (GK Quantile) estimation, can replace respective exact operators 133 in application workflows to enable or enhance all three types of required scalability 134 as well as to reduce memory utilisation. More precisely, data summaries leave only 135 a footprint of the stream in memory and they also enhance horizontal scalability 136 since not only is the processing load distributed to a number of available VMs, but 137 also it is shed by letting each VM operate on compact data summaries. Moreover, 138 synopses enable federated scalability since only summaries, instead of the full (set 139 of) streams, can be communicated when needed. Finally, synopses provide vertical 140 scalability by enabling locality-aware hashing [37, 38, 46]. 141

119

Related Work and State of the Art From a research viewpoint, there is a large 142 number of related works on data synopsis techniques. Such prominent techniques 143 are reviewed in [25, 26, 32] and have been implemented into real-world synopses 144 libraries, such as Yahoo!DataSketch [9], Stream-lib [8], SnappyData [57] and 145 Proteus [7]. Yahoo!DataSketch [9] and Stream-lib [8] are libraries of stochastic 146 streaming algorithms and summarisation techniques, correspondingly, but imple- 147 mentations are detached from parallelisation and distributed execution aspects 148 over streaming Big Data platforms. Apache Spark provides utilities for data 149 synopsis via sampling operators, CountMin sketches and Bloom Filters. Moreover, 150 SnappyData's [57] stream processing is based on Spark and its synopses engine can 151 serve approximate, simple sum, count and average queries. Similarly, Proteus [7] 152 extends Flink with data summarisation utilities. Spark utilities, SnappyData and 153 Proteus combine the potential of data summarisation with horizontal scalability, 154 i.e., parallel processing over Big Data platforms, by providing libraries of parallel 155 versions of data synopsis techniques. However, they neither handle all types of 156 required scalability nor cross Big Data platform execution scenarios. 157

INFORE Contribution In the scope of the INFORE project, we have developed 158 a Synopses Data Engine (SDE) [46] that advances the state of the art by tackling 159 all three types of the required scalability and also accounting for sharing synopses 160 common to various running workflows and for cross-platform execution. INFORE 161 SDE goes far beyond the implementation of a library of data summarisation 162 techniques. Instead, it also implements an entire component with its own internal 163 architecture, employing a Synopses-as-a-Service (SDEaaS) paradigm. That is, the 164 SDE is a constantly running service (job) in one or more clusters (federated 165 scalability) that can accept on-the-fly requests for start maintaining, updating and 166 querying a parallel synopsis built on a single high-speed stream (e.g. vessel) of 167 massive data proportions (horizontal scalability) or on a collection of a large number 168 of streams (vertical scalability). The SDEaaS is customisable to specific application 169 needs by allowing dynamic loading of code for new synopses operators at runtime, 170 with zero downtime for the workflows that it serves. 171

The architecture of INFORE SDEaaS [46] is illustrated in Fig. 1a. INFORE's 172 SDEaaS proof-of-concept implementation is based on Apache Kafka and Flink. 173 Nevertheless, the design is generic enough to remain equally applicable to other 174 Big Data platforms. For instance, an equally plausible alternative would be to 175 implement the whole SDE in Kafka leveraging the Kafka Streams API. Nonetheless, 176 Kafka Streams is simply a client library for developing micro-services, lack-177 ing a master node for global cluster management and coordination. Following 178 Fig. 1a, when a request for maintaining a new synopsis is issued, it reaches the 179 RegisterRequest and RegisterSynopsisFlatMaps which produce keys 180 for workers (i.e., VM resources) which will handle this synopsis. Each of this pair of 181 FlatMaps uses these keys for a different purpose. RegisterSynopsis uses the 182 keys to direct queries to responsible workers, while RegisterSynopsis uses the 183 keys to update the synopses on new data arrivals (blue-coloured path). In particular, 184 when a new streaming data tuple is ingested, the HashData FlatMap looks up the 185





keys of RegisterSynopsis to see to which workers the tuple should be directed 186 to update the synopsis. This update is performed by the add FlatMap in the bluecoloured path. The rest of the operators in Fig. 1a are used for merging partial 188 synopses results [11] maintained across workers or even across geo-distributed 189 computer clusters. Please refer to [46] for further details. In Sect. 3.2.3, we analyse 190 the functionality of a domain-specific synopsis building samples of vessel positions. 191

2.2 Distributed Online Machine Learning and Data Mining

Motivation As discussed in Sect. 1, ML/DM APIs such as Spark's MLlib [5] or 193 FlinkML [2] are focused on analysing data at rest. Therefore, advanced analytics 194 tasks on data in motion call for filling the gap of a stream processing-oriented 195 ML/DM module. ML and DM algorithms that can meet the challenges discussed 196 in the introduction of this chapter are those that (1) are online, i.e., restricting 197 themselves on a single pass over the data instead of requiring multiple passes, 198 and (2) can run in a distributed fashion, i.e., they are parallelisable and thus the 199 load can be distributed to parallel learners and parallel predictors across a number 200 of VMs so as to provide the primitives for horizontal scalability over Big Data 201 platforms and computer clusters. There exists a variety of algorithms that satisfy 202 these preliminary requirements in diverse ML/DM categories, including [18, 34, 69] 203 classification (such as (Multiclass) Passive Aggressive Classifiers, Online Support 204 Vector Machines, Hoeffding Trees, Random Forests), clustering (BIRCH, Online k- 205 Means, StreamKM++) and regression (Passive Aggressive Regressor, Online Ridge 206 Regression, Polynomial Regression) tasks. These algorithms are designed or can be 207 adapted to get executed in an online, distributed setting. The primary focus, then, is 208 not on the algorithms themselves, but on the architecture an ML/DM module should 209 be built upon, so that various algorithms can be incorporated and also allow for 210 vertical scalability, federated scalability and cross-platform execution with reduced 211 memory utilisation. 212

Related Work and State of the Art Towards this direction, the two most 213 prominent approaches and modules that exist in the literature are StreamDM [19] 214 and Apache SAMOA [48]. StreamDM is a library of ML/DM algorithms designed 215 to be easily extensible with new algorithms, but dedicated to run on top of the Spark 216 Streaming API [5]. Thus, it does not cover cross-platform execution scenarios, 217 also lacking provisions for vertical and federated scalability. The only framework 218 with a clear commitment to the cross-platform execution goals is Apache SAMOA. 219 SAMOA is portable between Apache Flink, Storm [6] and Samza [4]. When it 220 comes to its model of computation, the architecture of SAMOA follows the Agent-221 based pattern. In other words, an algorithm is a set of distributed processors 222 that communicate with streams of messages. Little more is provided, which is 223 intentional [48], claiming that a more structured model of computation reduces the 224 applicability of the framework.

The state of the art in distributed ML and DM architectures is the Parameter 226 Server (PS) distributed model [51] as illustrated in Fig. 1b, where a set of distributed 227 learners receive portions of the training streams and extract local models in parallel. 228 The local models are from time to time synchronised to extract a global model 229 at the PS side. The global model is then communicated back to learners via a 230 feedback loop (Fig. 1b). Consider for instance a set of learners each handling a 231 subset of vessel streams within the scope of a vessel type classification task. The 232 learners coordinate with the PS sending their locally trained classification models. 233 while the PS responds back with an up-to-date global model. The PS paradigm 234 enhances horizontal and federated scalability via the option of an asynchronous 235 (besides synchronous) synchronisation policy to reduce the effect of stragglers 236 and bandwidth consumption, respectively. In the synchronous policy, learners are 237 communicating with the PS in predefined rounds/batches, while in the asynchronous 238 case each learner decides individually as to when it should send updates to the PS. 239 Performance-wise, the synchronous policy does not encourage enhanced horizontal 240 scalability because when many learners are used, the total utilisation is usually low, 241 should only few stragglers exist. The asynchronous one is the policy of choice in 242 large-scale ML; the processing speed is much higher when many learners are used 243 and the training is more scalable. 244

The PS paradigm has been criticised for limited training speed due to potential 245 network congestion at the PS side and for severely getting affected by low-246 speed links between the learners and the PS. Under these claims, a number of 247 decentralised ML/DM architectures have evolved which employ a more peer-to-248 peer alike structure, where the training rationale is based on gossiping [42, 70]. The 249 drawback of these approaches, though, is that it is unclear how the continuously 250 updated, but decentralised, global model can be directly deployed for real-time 251 inference purposes. This is because knowing the network node holding the updated 252 global model at any given time requires extra communication. Hence, in case we 253 want to train and simultaneously deploy the updated global ML/DM models at 254 runtime, such a decentralised architecture does not seem to mitigate low-speed 255 issues but moves the problem to the prediction, instead of the training, stage.

INFORE Contribution In the scope of the INFORE project, we follow a PS ²⁵⁷ distributed model [51]. As is the case with the SDEaaS described in the previous sec- ²⁵⁸ tion, INFORE's ML/DM module includes provisions for cross-platform execution ²⁵⁹ scenarios by receiving input and output streams in JSON formatted Kafka messages. ²⁶⁰ Moreover, the communication between learners and the PS is performed using a ²⁶¹ lightweight middleware where a generic API for PS and learner (bidirectional) ²⁶² communication is provided. In that, learners can be implemented over any Big ²⁶³ Data platform and run in any cluster, while still being able to participate in the ²⁶⁴ common ML/DM task. Besides learners, INFORE's ML/DM module includes a ²⁶⁵ separate pipeline of parallel predictors that can communicate with the PS in order to ²⁶⁶ receive up-to-date global models continuously extracted during the training process ²⁶⁷ and directly deploy them for inference purposes. ²⁶⁸

AQ3

INFORE's ML/DM module accounts for vertical and boosts horizontal scalability as well. This is achieved by using INFORE's SDEaaS to partition streams 270 to learners or to allow learners to operate on compact stream summaries, correspondingly. Remarkably, to effectively encounter congestions or low-speed links 272 and also allow to easily and effectively deploy/update the developed models, 273 instead of resorting to decentralised approaches [42, 70], we develop our own 274 synchronisation policy termed FGM [67] (Fig. 1b) that improves the employed 275 PS paradigm. The new synchronisation protocol strengthens horizontal (within a 276 cluster) and federated scalability by bridging the gap between synchronous and 277 asynchronous communication. Instead of having learners communicating in predefined rounds/batches (synchronous) or when each one is updated (asynchronous), 279 FGM requires communication only when a concept drift (i.e., the global model has 280 significantly changed based on some criterion) is likely to have occurred. This is 281 determined based on conditions each learner can individually examine.

2.3 Distributed and Online CEF

Motivation Big Data analytics tools mine data views to extract patterns conveying 284 insights into what has happened, and then apply those patterns to make sense of 285 the fresh data that stream in. This only permits to react upon the detection of 286 such patterns, which is often inadequate. In order to allow for proactive decision-287 making, predictive analytics tools that allow to forecast future events of interest 288 are required. Consider, for instance, the ability to forecast and proactively respond 289 to hazardous events, such as vessel collisions or groundings, in the maritime 290 domain. The ability to forecast, as early as possible, a good approximation to 291 the outcome of a time-consuming and resource-demanding computational task 292 allows to quickly identify possible outcomes and save valuable reaction time, 293 effort and computational resources. Diverse application domains possess different 294 characteristics. For example, monitoring of moving entities has a strong geospatial 295 component, whereas in stock data analysis this component is minimal. Domain- 296 specific solutions (e.g. trajectory prediction for moving objects) cannot thus be 297 universally applied. We need a more general Complex Event Forecasting (CEF) 298 framework. 299

Related Work and State of the Art Time-series forecasting is an area with some 300 similarities to CEF, with a significant history of contributions [56]. However, it is 301 not possible to directly apply techniques from time-series forecasting to CEF. Time- 302 series forecasting typically focuses on streams of (mostly) real-valued variables and 303 the goal is to forecast relatively simple patterns. On the contrary, in CEF we are 304 also interested in categorical values, related through complex patterns and involving 305 multiple variables. Another related field is that of prediction of discrete sequences 306 over finite alphabets and is closely related to the field of compression, as any com- 307 pression algorithm can be used for prediction and vice versa [17, 20, 24, 63, 64, 73]. 308

The main problem with these approaches is that they focus exclusively on next 309 symbol prediction, i.e., they try to forecast the next symbol(s) in a stream/string 310 of discrete symbols. This is a serious limitation for CEF. An additional limitation 311 is that they work on single-variable discrete sequences of symbols, whereas CEF 312 systems consume streams of events, i.e., streams of tuples with multiple variables, 313 both numerical and categorical. Forecasting methods have also appeared in the field 314 of temporal pattern mining [22, 50, 71, 75]. A common assumption in these methods 315 is that patterns are usually defined either as association rules [13] or as frequent 316 episodes [53]. From the perspective of CEF, the disadvantage of these methods is 317 that they usually target simple patterns, defined either as strictly sequential or as 318 sets of input events. Moreover, the input stream is composed of symbols from a 319 finite alphabet, as is the case with the compression methods mentioned previously. 320

INFORE Contribution In a nutshell, the current, state-of-the-art solutions for 321 forecasting, even when they are domain-independent, are not suitable for the kind 322 of challenges that INFORE attempts to address. In INFORE, the streaming input 323 can be constantly matched against a set of event patterns, i.e. arbitrarily complex 324 combinations of time-stamped pieces of information. An event pattern can either 325 be fully matched against the streaming data, in which case events are detected, 326 or partially matched, in which case events are forecast with various degrees of 327 certainty. The latter usually stems from stochastic models of future behaviour, 328 embedded into the event processing loop, which project into the future the sequence 329 of events that resulted to a partial event pattern match, to estimate the likelihood of 330 a full match, i.e. the actual occurrence of a particular complex event. 331

Given that INFORE's input consists of a multitude of data streams, interesting 332 events may correlate sub-events across a large number of different streams, with 333 different attributes and different time granularities. For instance, in the maritime 334 domain relevant streams may originate from position signals of thousands of vessels 335 which may be fused with satellite image data [54] or even acoustic signals [40]. It 336 is necessary to allow for a highly expressive event pattern specification language, 337 capable of capturing complex relations between events. Moreover, the actual 338 patterns of what constitutes an interesting event are often not known in advance, 339 and even if they are, event patterns need to be frequently updated to cope with the 340 drifting nature of streaming data. Not only do we need an expressive formalism in 341 order to capture complex events in streams of data, but we also need to do so in a 342 distributed and online manner. 343

Towards this direction, the CEF module of INFORE uses a highly expressive, 344 declarative event pattern specification formalism, which combines logic, probability 345 theory and automata theory. This formalism has a number of key advantages: 346

- It is capable of expressing arbitrarily complex relations and constraints between 347 events. We are thus not limited to simple sequential patterns applied to streams 348 with only numerical or symbolic values. 349
- It can be used for event forecasting and offering support for robust temporal 350 reasoning. By converting a pattern into an automaton, we can then use historical 351

data to construct a probabilistic description of the automaton's behaviour and ³⁵² thus to estimate at any point in time its expected future behaviour. ³⁵³

• It offers direct connections to machine learning techniques for refining event ³⁵⁴ patterns, or learning them from scratch, via tools and methods from the field ³⁵⁵ of grammatical inference. In cases where we only have some historical data and ³⁵⁶ some labels, we must find a way to automatically learn the interesting patterns. ³⁵⁷ This is also the case when there is concept drift in the streaming data and the ³⁵⁸ patterns with which we started may eventually become stale. It is therefore ³⁵⁹ important to be able to infer the patterns in the data in an online manner. ³⁶⁰

INFORE's CEF module is built on top of Apache Kafka and Flink and has the ³⁶¹ ability to handle highly complex patterns in an online manner, constantly updating ³⁶² its probabilistic models. Figure 1d shows one possible scheme (pattern-based) for ³⁶³ structuring multiple parallel CEF pipelines. As shown in the figure, each such ³⁶⁴ pipeline processes a different CEF query [33, 35]. It is composed of a training ³⁶⁵ process, which estimates the probabilities of a future event to occur, as well as a ³⁶⁶ CEF process that utilises these probabilities to actually forecast complex events. ³⁶⁷ Finally, one implementation detail is that each pipeline also receives a subset of the ³⁶⁸ patterns (part1 to partX in Fig. 1d). The role of these loops is similar to the feedback ³⁶⁹ loop of Fig. 1b. Remarkably, the CEF module can also act as a CEP one since it can ³⁷⁰ not only predict but also detect occurred events of interest [14].

2.4 Geo-distributed Cross-Platform Optimisation

Motivation All the aforementioned advanced stream processing techniques and 373 technologies will only serve their goal if they are properly used. Consider, for 374 instance, that we perfectly tune the execution of a synopsis, ML/DM or CEF 375 operator in a specific cluster, but we assign the execution of the downstream operator 376 of a broader workflow to a distant cluster. The execution speed up achieved for one 377 operator may be diminished by network latency of long network paths. Therefore, 378 developing algorithms for optimising the execution of streaming workflows (a) 379 over a network of many clusters located in various geographic areas, (b) across 380 a number of Big Data platforms available in each cluster and (c) simultaneously 381 elastically devoting VMs and resources (CPU, memory, etc.) is a prerequisite 382 to efficiently deliver in practice real-time analytics. Within a cluster, common 383 optimisation objectives include throughput maximisation, execution latency and 384 memory usage minimisation, while in multi-cluster settings communication cost, 385 bandwidth consumption and network latency are also accounted for. Quality-of- 386 Service (QoS) and computer cluster (CPU, memory, storage) capacity constraints 387 also apply to these objectives. 388

Related Work and State of the Art There are a number of works that assign 389 the execution of operators targeting at optimising network-related metrics, such 390 as communication cost and network latency, while executing global analytics 391

workflows across a number of networked machines or computer clusters. The 392 seminal work of SBON [59] seeks to optimise a quantity similar to network 393 usage (*dataRate* × *latency*), but with a squared latency, across multi-hop paths 394 followed by communicated data. An important limitation in SBON is that by 395 using such a blended metric, the optimisation process cannot support constrained 396 optimisation per metric (communication cost or latency). Due to that, also other 397 related techniques [49, 59, 62] which employ blended metrics cannot incorporate 398 resource or QoS constraints while determining operators' assignment to clusters. 399 Although some [49, 62] claim to support latency constraints, this comes after having 400 determined where an operator will be executed. Finally, the approach of Geode [72] 401 purely focuses on minimising bandwidth consumption in the presence of regulatory 402 constraints, but it does not account for network latency. 403

A series of works aim at optimising the execution of analytics operators within 404 a single computer cluster. Such works focus on optimal assignment of operators 405 to VMs such that high performance (mainly, in terms of throughput) and load 406 balancing among VMs is achieved; subject to multiple function, resource and QoS 407 constraints. Related works mainly provide optimisations on load assignment and 408 distribution, load shedding, resource provisioning and scheduling policies inside 409 the cluster. In Medusa [16], Borealis [10], Flux [68] and Nexus [23], the focus is to 410 primarily balance the load, choose appropriate ways to partition data streams across 411 a number of machines and minimise the usage of available resources (CPU cycles, 412 bandwidth, memory, etc.) while maintaining high performance.

Another category of techniques examines the optimisation of network-wide 414 analytics, simultaneously scaling-out the computation of an operator to the VMs 415 of the cluster that undertakes its execution. JetStream [61] trades-off network 416 bandwidth minimisation with timely query answer and correctness, but while 417 exploring the cluster at which an operator will be executed, it restricts itself to 418 the MapReduce rationale (i.e. the operator is executed at the cluster where data 419 rests), nearest site of relevant data presence or a central location. Iridium [60], 420 basically targeting optimisation of analytics over data at rest, assumes control over 421 where relevant data are transferred and moves these data around clusters to optimise 422 query response latency. SQPR [45] and [21] propose more generic frameworks for 423 the constraint-aware optimal execution of global workflows across clusters, and 424 they also optimise resources devoted to each operator execution at each cluster. 425 However, [21, 45] do not account for cross-platform optimisation in the presence 426 of different Big Data technologies.

Systems such as Rheem [12], Ires [27], BigDawg [28] and Musketeer [39] are 428 designed towards cross-platform execution of workflows, but they can only optimise 429 the processing of data at rest,² instead of data in motion. Furthermore, only Rheem 430 accounts for network-related optimisation parameters such as communication cost. 431

² BigDawg supports stream processing over S-Store and Rheem supports JavaStreams, but no alternatives are included to allow for optimising across different streaming platforms.

INFORE Contribution The INFORE Optimiser is the first complete solution for 432 streaming operators [30, 34]. INFORE's Optimiser is not simply the only one which 433 can simultaneously instruct the streaming Big Data platform, cluster and computing 434 resources for each analytics operator, but also it does so for a wide variety of 435 diverse operator classes including (1) synopses, (2) ML/DM, (3) CEF and (4) stream 436 transformations. INFORE's Optimiser incorporates the richest set of optimisation 437 criteria related to throughput, network and computational latency, communication 438 cost, memory consumption and accuracy of SDE operators, and it also accounts for 439 constraints per metric, fostering the notion of Pareto optimality [30, 34].

The internals of INFORE Optimiser are illustrated in Fig. 1c. We use a statistics 441 collector to derive performance measurements from each executed workflow. 442 Statistics are collected via JMX or Slurm³ and are ingested in an ELK stack⁴ 443 while monitoring jobs. A Benchmarking submodule automates the acquisition of 444 performance metrics for SDE, OMLDM and CEF/CEP operators run in different 445 Big Data platforms. The Benchmarking submodule utilises statistics and builds 446 performance (cost) models. Cost models are derived via a Bayesian Optimisation 447 approach inspired by CherryPick [15]. The cost models are utilised by the optimi- 448 sation algorithms [30, 34] to prescribe preferable physical execution plans. 449

3 Real-Life Application to a Maritime Use Case 450

3.1 Background on Maritime Situation Awareness (MSA) 451

According to the US National Concept of Operations for Maritime Domain Awareness,⁵ "Global Maritime Intelligence is the product of legacy, as well as changing intelligence capabilities, policies and operational relationships used to *integrate all available data*, information, and intelligence in order to identify, locate, and track potential *maritime threats*. Global MSA results from the *persistent monitoring* of maritime activities in such a way that *trends and anomalies* can be identified".

Maritime reporting systems are distinguished into two broad categories: *cooper-* 458 *ative* and *non-cooperative*. An example of a *cooperative* maritime reporting system 459 is the Automatic Identification System (AIS) [43]. All commercial vessels above 460 300 gross tonnage are obliged to bear AIS transponders. AIS forms the basis of 461 a lot of MSA applications, such as the MarineTraffic vessel tracking platform. 462 Other cooperative, but not public, maritime reporting systems are the Long Range 463 Identification and Tracking system (LRIT) [44], as well as the Vessel Monitoring 464

³ https://docs.oracle.com/javase/tutorial/jmx/overview/, https://slurm.schedmd.com/.

⁴ https://www.elastic.co/what-is/elk-stack.

⁵ https://web.archive.org/web/20111004213300/http://www.gmsa.gov/twiki/bin/view/Main/ MDAConOps.

478

487

System (VMS) [29] for fishing vessels. Radar on-board or ashore installations can 465 be used as maritime surveillance systems, such as the ones installed by default in 466 a vessel's bridge, as well as in ports. Thermal cameras and satellite imagery can 467 also be used as additional monitoring systems for vessels. Due to the time elapsed 468 between the actual image acquisition from a satellite and its availability on the 469 satellite repository that can be several hours, satellite imagery data do not offer realtime snapshots of the maritime domain but can be used combined with other sources 471 such as AIS to "fill in the gaps" of AIS coverage (e.g., identify the whereabouts of 472 a vessel while its transponder was switched off).

Global and *continuous* monitoring of the maritime domain as well as the 474 identification of trends and anomalies require to address the challenges pointed 475 out throughout this chapter as well as the following generic Big Data challenges 476 described in the scope of the maritime domain:

- *Volume*, the number of available surveillance systems and sensors increases.
- Velocity, applications rely on continuous monitoring (e.g., vessel tracking) and 479 need to process high velocity streaming data in real time.
- Variety, data from heterogeneous surveillance systems should be combined. 481
- Veracity, most of the maritime data sources are heavily prone to noise requiring 482 data cleaning and analysis tasks to filter out unnecessary or invalid information. 483
- Value, as the availability of more sources of maritime data as well as the advanced 484 Big Data processing, ML and AI technologies that are now available can help to 485 maximise the derived knowledge that can be inferred from maritime data. 486

3.2 Building Blocks of MSA Workflows in the Big Data Era

Figure 2b shows an example of a generic workflow, implemented in the Maritime 488 Use Case of the INFORE project, for MSA purposes. Different applications may 489 include a subset of operators of Fig. 2b or implement different steps. In the 490



(a)Anomaly Detection at MarineTraffic.



Fig. 2 MSA infrastructure and workflow

following, we describe the functionality of the workflow operators of Fig. 2b which 491 serve as the building blocks of modern MSA applications. 492

3.2.1 Maritime Data Sources

The kinds of data sources that are provided as input in a typical MSA application 494 (Fig. 2b) are the following: 495

- Vessel positions. Data about vessel positions derive from vessel reporting 496 systems, the most popular of which is AIS. AIS forms the ground of a wide 497 variety of MSA applications. AIS relies on VHF communication: Vessels send 498 AIS messages that contain dynamic information (e.g., information about the 499 current voyage, such as vessel position, speed, heading, etc.) as well as static 500 information (e.g., vessel identifier, dimensions, etc.). For real-time applications, 501 positional data arrive in streaming fashion to the data consumers.
- Data from other sensors. Some applications do not rely only on one source of 503 information. For example, AIS data can be combined with acoustic data, thermal 504 camera data and satellite data. Vessel detection algorithms are applied on this 505 data to extract the positions of vessels. For example, AI techniques are applied 506 on satellite imagery to extract the vessel positions which is important in the cases 507 when a vessel is out of AIS coverage [54].
- Other datasets describing assets and activities in the maritime domain. These 509 are datasets that describe ports, harbours, lighthouses, the boundaries of areas 510 of interest, bathymetry datasets (e.g., for shallow waters estimation), datasets 511 containing vessel schedules, weather data, etc. These datasets are often combined 512 with other data (e.g., vessel positions) in order to enrich the information 513 displayed to the end-users (e.g., the different layers of the MarineTraffic Live 514 Map).

Kafka [3] is used at the data ingestion layer, as a fast, scalable and fault-tolerant 516 messaging system for large data (at rest or in motion) portions. 517

3.2.2 Maritime Data Fusion

Data from multiple sources besides AIS, such as radars and cameras, are available 519 in real time though in order to be used in MSA modules they must be fused 520 together with AIS and create a unified map. This essentially translates to a need 521 for tracking algorithms that can monitor moving objects globally and in real time 522 using overlapping detections from multiple sensors. The Fusion operator in Fig. 2b 523 is a custom operator with distributed implementations in order to achieve this goal. 524 Trackers are comprised of three main components [65, 66]: (a) a method for the 525 assignment of detections to tracks, (b) the prediction of a target's movement and (c) 526 the architecture of the tracker that coordinates how the detections are processed. 527

A detection arriving to the tracker can be assigned to a track using three 528 strategies, and each tracker implementation is based on one of them. The first way 529 is to simply choose the track that is closest to the detection, which has the lowest 530 computational complexity but it is not accurate in cases where two objects move 531 very close to each other. The second method focuses on improving the accuracy in 532 cases where a detection is close to multiple tracks by deferring the final assignment 533 until more detections arrive, thus making a more informed decision but at the cost of 534 significantly increasing the complexity and decreasing the responsiveness (i.e., real-535 time challenge). The third approach stands between the two methods and allows that 536 a detection is assigned to multiple tracks as soon as it arrives, thus increasing the 537 accuracy satisfactorily without increasing complexity. 538

Each moving object is characterised by certain physical parameters and constraints according to which several kinematic models can predict its movement under different conditions. A simple option is to choose one model, such as constant velocity that assumes the object maintains the last speed, but this affects the accuracy when an object manoeuvres. A better option is to use multiple models, such as constant turn and acceleration, at the same time so that the tracker is able to successfully detect a manoeuvring target.

3.2.3 SDE Operator For Trajectory Simplification

The plethora of incoming data from multiple overlapping sources poses a challenge 547 for data processing workflows. A data synopsis technique with which this challenge 548 can be tackled is trajectory simplification, i.e., reducing the amount of data 549 (positions) so that the computational effort required is reduced as well. The ideal 550 goal is to keep only those positions that are adequate in order to recreate the 551 trajectory with minimal losses in the accuracy of the data processing workflow. 552

For that, we use INFORE's SDEaaS (Sect. 2.1) which includes an application- 553 specific synopses, namely STSampler. The STSampler scheme resembles the 554 concept of threshold-guided sampling in [58] but executes the sampling process in a 555 more simplistic, yet effective in practice, way. More precisely, the sampling process 556 is executed in a per stream fashion, i.e., for the currently monitored trajectory of 557 each vessel separately. The core concept is that if the velocity and the direction 558 of the movement of the vessel do not change significantly, the corresponding AIS 559 message is not sampled. The last two reported trajectory positions are cached in 560 the add FlatMap of Fig. 1a. When an AIS message holding information about 561 the current status of the vessel streams in via HashData, the add FlatMap 562 computes the change in the velocity between the lastly cached and the new AIS 563 report, i.e., $\Delta_{vel} = |vel(prev) - vel(now)|$, and compares this value to a velocity 564 threshold T_{vel} . Using the previously cached points, the vector describing the lastly 565 reported direction of the vessel dir(prev) is computed, while using the last cached 566 and the newly reported positions we also compute dir(now). Then, we compare 567 $\Delta_{dir} = |dir(prev) - dir(now)|$ against a direction threshold T_{dir} . If at least one 568 of these deltas does not exceed the corresponding threshold, the newly received 569

AIS message is not included in the sample by the add FlatMap. This holds, 570 provided that a couple of additional spatiotemporal constraints are satisfied: (a) the 571 time difference between the newly received AIS message and the last one that was 572 included in the sample does not exceed a given time interval threshold T_{tdiff} and 573 (b) the distance among the most recently sampled and the current position of the 574 vessel does not surpass a distance threshold T_{dist} . SDEaaS is implemented in Flink 575 instead of Kafka, for the reasons explained in Sect. 2.1. 576

3.2.4 Complex Maritime Event Processing

A very important module of the modern MSA applications is the Maritime Event 578 Detection module. This is essentially a CEP module tailored to the maritime domain. 579 For now, our analysis concentrates on distributed and online CEP, i.e., detecting 580 complex events, while future work will also exploit the potential of CEF (Sect. 2.3). 581 A description of some of the most common vessel events that can occur in the 582 maritime domain is provided below: 583

• <i>Turn</i> : A vessel turns to a different direction.	584
Acceleration: A vessel accelerates.	585
• <i>Route Deviation</i> : The course of a vessel deviates from "common" routes.	586
Shallow waters: A vessel navigates in shallow waters.	587
• Proximity: A vessel is in close distance to another vessel.	588
• Out of coverage. A vessel is out of coverage with respect to one or more vessel	589
monitoring systems such as AIS [47].	590
The events described above are simple events, i.e., they can be computed without	591
depending on other events. Complex events, on the other hand, are events composed	592
from other events. Below we provide examples of complex events:	593
• Ship-to-ship: Transfer of cargo between vessels.	594
• Bunkering: One vessel provides fuel to another vessel.	595
• <i>Tugging</i> : A smaller vessel (a tug) is tugging another vessel.	596
• Piloting: A smaller vessel (pilot vessel) approaches a bigger vessel so that the	597
pilot of the vessel boards the bigger vessel in order to help it navigate into a port	598
where special local conditions apply.	599
• Fishing: A vessel is engaged in fishing activities.	600
For distributed processing of streaming data in the CEP context, the Akka	601
framework is used [1]. Akka adopts an Actor-based architecture based on message-	602

framework is used [1]. Akka adopts an Actor-based architecture based on messagepassing communication, and it is preferred due to the fact that it is more customisable than Spark and Flink. Each Actor, run in parallel instances, is responsible for detecting a simple or complex event as those described above (Fig. 2a and b).

3.2.5 ML-Based Anomaly Detection

The ML algorithms that are relevant to the MSA workflow relate to Deep Neural 607 Network techniques for classifying vessels according to their type (such as cargo, 608 fishing vessel) [54]. Moreover, we are investigating ML-based techniques such as 609 Random Forests for classifying vessel trajectories and recognise simple or complex 610 events in them. This effort is also aided by advanced ML-based operators we have 611 developed to extract the common routes followed by the majority of vessels for 612 every voyage, defined as a pair of origin and destination ports [78]. At the moment, 613 these ML tasks are performed in an offline fashion mostly using Spark's MLlib [5], 614 which we also use to estimate sea-port area regions in [55]. The outcomes of this 615 process performed at the batch layer of Fig. 2a can then be used as added value 616 knowledge to the event detection or the Fusion operator of Fig. 2b. Our ongoing 617 work focuses on incorporating INFORE's module (Sect. 2.2) to materialise ML/DM 618 analytics in an online, real-time fashion, where possible (see restrictions on satellite 619 images in Sect. 3.1). 620

3.2.6 MSA Workflow Optimisation

Across the workflow of Fig. 2b, the INFORE Optimiser is responsible for prescribing the parallelisation Degree, and the provisioned resources for the maintained 623 trajectory synopses (Sect. 3.2.3) determine the computer cluster and the number of Akka Actors devoted to MSA-related CEP tasks (Sect. 3.2.4). The Optimiser can also do the same for ML-based anomaly detection tasks (Sect. 3.2.5). An initial workflow execution plan can be re-optimised and adjusted at runtime to adapt (e.g., by increasing/decreasing the number of Akka Actors) to changing data stream distributions or to a load of concurrently executed maritime workflows. Moreover, the ongoing integration of the INFORE CEF module will allow the Optimiser to prescribe the most efficient implementation among Akka (Sect. 3.2.4) and Flink (Sect. 2.3) options for event processing tasks.

4 Future Research and Development Directions

Future research and development directions mainly lie in the synergies of ML/DM, 634 Synopses, CEP/CEF and optimisation technologies discussed in this chapter. 635

Resource-Constrained ML/DM Resource-Constrained ML/DM goes beyond 636 data processing over distributed, but computationally powerful infrastructures such 637 as computer clusters or the cloud. The objective in resource constrained ML/DM is 638 to bridge the gap between the very high computation and communication demands 639 of state-of-the-art ML algorithms, such as Deep Neural Nets and Kernel Support 640 Vector Machines, and the goal of running such algorithms (e.g. various classifiers) 641

606

621

on a large, heavily distributed system of resource-constrained devices. Resource- ⁶⁴² constrained devices, such as sensors, pose limitations to the power supply, memory, ⁶⁴³ computation and communication capacity. Fast and efficient classifiers requiring ⁶⁴⁴ reduced power and memory should be developed, along with novel algorithms to ⁶⁴⁵ train, apply and update the classifiers. Synergies between synopses and distributed, ⁶⁴⁶ online ML/DM utilities are critical for such tasks. ⁶⁴⁷

Optimisation over Internet of Things (IoT) Platforms Optimisation over Internet of Things (IoT) platforms, since existing optimisation frameworks, should be extended to allow for planning the execution of workflows taking into consideration the whole set IoT features including: (a) resource scarcity, (b) hardware heterogeneity, (c) data heterogeneity, (d) dynamic population of devices, (e) mobility of devices, (f) security aspects over massively distributed architectures, and (g) resilience and accuracy of analytics in the presence of device failures.

CEP/CEF-Oriented Synopses CEP/CEF-Oriented Synopses techniques tailored 655 for CEP/CEF are becoming a necessity. The work in [41] was the first to point out 656 that load shedding schemes tailored for CEP are missing and that shedding the load 657 in CEP significantly differentiates itself from doing so in conventional streaming 658 settings. A few more approaches emerged since then [52, 74], but still little attention 659 has been paid on the distributed environments and the mergeability properties of 660 such techniques [11]. 661

Acknowledgments This work has received funding from the EU Horizon 2020 research and 662 innovation program INFORE under grant agreement No. 825070.

References

1. Akka v. 2.5.32. https://akka.io/. Accessed 15 September 2020.	665
2. Apache Flink v. 1.12. https://flink.apache.org/. Accessed 15 September 2020.	666
3. Apache Kafka v. 2.3. https://kafka.apache.org/. Accessed 15 September 2020.	667
4. Apache Samza v. 1.5.1. http://samza.apache.org/. Accessed 15 September 2020.	668
5. Apache Spark v. 2.4.4. https://spark.apache.org/. Accessed 15 September 2020.	669
6. Apache Storm v. 2.1. https://storm.apache.org/. Accessed 15 September 2020.	670
7. Proteus project. https://github.com/proteus-h2020/. Accessed 15 September 2020.	671
8. Stream-lib. https://github.com/addthis/stream-lib/. Accessed 15 September 2020.	672
9. Yahoo datasketch. https://datasketches.github.io/ Accessed 15 September 2020.	673
10. Abadi, D. J., Ahmad, Y., Balazinska, M., Cetintemel, U., Cherniack, M., Hwang, J. H., Lindner,	674
W., Maskey, A. S., Rasin, A., Ryvkina, E., Tatbul, N., Xing, Y., Zdonik, S. (2005). The design	675
of the borealis stream processing engine. In CIDR.	676
11. Agarwal, P. K., Cormode, G., Huang, Z., Phillips, J. M., Wei, Z., & Yi, K. (2013). Mergeable	677
summaries. ACM Transactions on Database Systems, 38(4), 26:1–26:28.	678
12. Agrawal, D., Chawla, S., Rojas, B., et al. (2018). RHEEM: enabling cross-platform data	679
processing - may the big data be with you! Proceedings of the VLDB Endowment, 11(11),	680
1414.	681
13. Agrawal, R., Imielinski, T., & Swami, A. N. (1993). Mining association rules between sets of	682
items in large databases. In SIGMOD.	683

14. Alevizos, E., Artikis, A., Paliouras, G. (2018). Wayeb: a tool for complex event forecasting. In <i>LPAR</i> .	684 685
15. Alipourfard, O., Liu, H., Chen, J., Venkataraman, S., Yu, M., & Zhang, M. (2017). Cherrypick: Adaptively unearthing the best cloud configurations for big data analytics. In <i>NSDI</i>	686 687
16. Balazinska, M., Balakrishnan, H., Stonebraker, M. (2004). Contract-based load management is fourted littliheted systems in NDM.	688
17. Begleiter, R., El-Yaniv, R., & Yona, G. (2004). On prediction using variable order Markov	689 690
models. Journal of Artificial Intelligence Research, 22, 385-421.	691
18. Benczúr, A., Kocsis, L., & Pálovics, R. (2018). Online machine learning in big data streams.	692
arXiv: 1802.05872.	693
19. Bifet, A., Maniu, S., Qian, J., Tian, G., He, C., & Fan, W. (2015). Streamdm: Advanced data	694
mining in spark streaming. In <i>ICDMW</i> .	695
20. Buhlmann, P., Wyner, A. J., et al. (1999). Variable length Markov chains. The Annals of Statistics, 27(2), 480–513.	696 697
21. Cardellini, V., Grassi, V., Lo Presti, F., & Nardelli, M. (2016). Optimal operator placement for	698
distributed stream processing applications. In DEBS.	699
22. Cho, C., Wu, Y., Yen, S., Zheng, Y., & Chen, A. L. P. (2011). On-line rule matching for event	700
prediction. VLDB Journal, 20(3), 303–334.	701
23. Cipriani, N., Eissele, M., Brodt, A., Grossmann, M., & Mitschang, B. (2009). Nexusds: a	702
flexible and extensible middleware for distributed stream processing. In IDEAS.	703
24. Cleary, J. G., & Witten, I. H. (1984). Data compression using adaptive coding and partial string	704
matching. <i>IEEE Transactions on Communications</i> , 32(4), 396–402.	705
25. Cormode, G., Garofalakis, M. N., Haas, P. J., & Jermaine, C. (2012). Synopses for massive	706
data: Samples, histograms, wavelets, sketches. Foundations and Trends Databases, $4(1-3)$, 1 204	707
1-294. 26 Cormode G. & Vi K. (2020). Small summarias for hig data. Combridge University Press.	708
https://doi.org/10.1017/9781108769938	710
27. Doka, K., Papailiou, N., Tsoumakos, D., Mantas, C., & Koziris, N. (2015). Ires: Intelligent.	711
multi-engine resource scheduler for big data analytics workflows. In SIGMOD.	712
28. Elmore, A. J., Duggan, J., Stonebraker, M., Balazinska, M., Çetintemel, U., Gadepally, V.,	713
Heer, J., Howe, B., Kepner, J., Kraska, T., Madden, S., Maier, D., Mattson, T. G., Papadopoulos,	714
S., Parkhurst, J., Tatbul, N., Vartak, M., Zdonik, S. (2015). A demonstration of the bigdawg	715
polystore system. Proceedings of the VLDB Endowment, 8(12), 1908.	716
29. FAO: VMS for fishery vessels. http://www.fao.org/fishery/topic/18103/en. Accessed 15 May	717
2019.	718
30. Flouris, I., Giatrakos, N., Deligiannakis, A., & Garofalakis, M. N. (2020). Network-wide	719
28 101442	720
31 Flouris I Giatrakos N Garofalakis M N & Deligiannakis A (2015) Issues in complex	721
event processing systems. In <i>IEEE TrustCom/BigDataSE/ISPA</i> (Vol. 2)	723
32. Garofalakis, M. N., Gehrke, J., & Rastogi, R. (Eds.). (2016). Data stream management -	724
processing high-speed data streams. Data-centric systems and applications. Springer. https://	725
doi.org/10.1007/978-3-540-28608-0	726
33. Giatrakos, N., Alevizos, E., Artikis, A., Deligiannakis, A., & Garofalakis, M. N. (2020).	727
Complex event recognition in the big data era: a survey. VLDB Journal, 29(1), 313–352.	728
34. Giatrakos, N., Arnu, D., Bitsakis, T., Deligiannakis, A., Garofalakis, M. N., Klinkenberg, R.,	729
Konidaris, A., Kontaxakis, A., Kotidis, Y., Samoladas, V., Simitsis, A., Stamatakis, G., Temme,	730
F., IOROK, M., YAQUD, E., MONTAGUD, A., PONCE, M., Arndt, H., Burkard, S. (2020). Infore: Interactive cross platform analytics for everyone. In CIVM	731
111 112 112 112 112 112 112 112 112 112	132
recognition in the big data era <i>Proceedings of the VIDR Endowment</i> 10(12) 1996	133
36. Giatrakos, N., Artikis, A., Deligiannakis, A., & Garofalakis, M. N. (2019). Uncertainty-aware	735
event analytics over distributed settings. In <i>DEBS</i> .	736

- 37. Giatrakos, N., Deligiannakis, A., Garofalakis, M. N., & Kotidis, Y. (2020). Omnibus outlier 737 detection in sensor networks using windowed locality sensitive hashing. *Future Generation* 738 *Computer Systems*, 110, 587–609. 739
- Giatrakos, N., Kotidis, Y., & Deligiannakis, A. (2010). PAO: power-efficient attribution of 740 outliers in wireless sensor networks. In DMSN. https://doi.org/10.1145/1858158.1858168 741
- Gog, I., Schwarzkopf, M., Crooks, N., Grosvenor, M. P., Clement, A., & Hand, S. (2015).
 Musketeer: all for one, one for all in data processing systems. In *EuroSys*.
- Goldhahn, R., Braca, P., Ferri, G., Munafo, A., & Lepage, K. (2014). Adaptive bayesian 744 behaviors for AUV surveillance networks. In UAC.
- He, Y., Barman, S., & Naughton, J. F. (2014). On load shedding in complex event processing. 746 In *ICDT*. 747
- 42. Hegedüs, I., Danner, G., & Jelasity, M. (2019). Gossip learning as a decentralized alternative 748 to federated learning. In *DAIS*. 749
- 43. IMO. (2017). Technical characteristics for an automatic identification system using time 750 division multiple access in the VHF maritime mobile frequency band. Technical report, ITU. 751 https://www.itu.int/dms_pubrec/itu-r/rec/m/R-REC-M.1371-5-201402-I!!PDF-E.pdf 752
- 44. IMO. (2018). Long-range identification and tracking system. Technical report, IMO. http:// 753
 www.imo.org/en/OurWork/Safety/Navigation/Documents/LRIT/1259-Rev-7.pdf 754
- Kalyvianaki, E., Wiesemann, W., Vu, Q. H., Kuhn, D., & Pietzuch, P. (2011). Sqpr: Stream 755 query planning with reuse. In *ICDE*. 756
- Kontaxakis, A., Giatrakos, N., & Deligiannakis, A. (2020). A synopses data engine for 757 interactive extreme-scale analytics. In *CIKM*.
- Kontopoulos, I., Chatzikokolakis, K., Zissis, D., Tserpes, K., & Spiliopoulos, G. (2020). Realtime maritime anomaly detection: detecting intentional AIS switch-off. *International Journal 760 of Big Data Intelligence*, *7*(2), 85–96.
- Kourtellis, N., Morales, G. D. F., & Bifet, A. (2018). Large-scale learning from data streams r62 with apache SAMOA. CoRR abs/1805.11477. http://arxiv.org/abs/1805.11477
- Kumar, V., Cooper, B. F., & Schwan, K. (2005). Distributed stream management using utilitydriven self-adaptive middleware. In *ICAC*.
- Laxman, S., Tankasali, V., & White, R. W. (2008). Stream prediction using a generative model 766 based on frequent episodes in event sequences. In *KDD*. 767
- 51. Li, M., Andersen, D., Park, J. W., et al. (2014). Scaling distributed machine learning with the 768 parameter server. In *OSDI*. 769
- Li, Z., & Ge, T. (2016). History is a mirror to the future: Best-effort approximate complex event 770 matching with insufficient resources. *Proceedings of the VLDB Endowment*, 10(4), 85–96. 771
- Mannila, H., Toivonen, H., & Verkamo, A. I. (1997). Discovery of frequent episodes in event requences. *Data Mining and Knowledge Discovery*, 1(3), 259–289.
- Milios, A., Bereta, K., Chatzikokolakis, K., Zissis, D., & Matwin, S. (2019). Automatic fusion 774 of satellite imagery and AIS data for vessel detection. In *FUSION*. 775
- Millefiori, L. M., Zissis, D., Cazzanti, L., & Arcieri, G. (2016). A distributed approach to 776 estimating sea port operational regions from lots of AIS data. In *IEEE BigData*. 777
- Montgomery, D. C., Jennings, C. L., & Kulahci, M. (2015). Introduction to time series analysis 778 and forecasting. John Wiley & Sons. 779
- 57. Mozafari, B. (2019). Snappydata. In Encyclopedia of Big Data Technologies. Springer.
- Patroumpas, K., Alevizos, E., Artikis, A., Vodas, M., Pelekis, N., & Theodoridis, Y. (2017).
 Online event recognition from moving vessel trajectories. *GeoInformatica*, 21(2), 389–427.
- Pietzuch, P., Ledlie, J., Shneidman, J., Roussopoulos, M., Welsh, M., & Seltzer, M. (2006).
 Network-aware operator placement for stream-processing systems. In *ICDE*.
- 60. Pu, Q., Ananthanarayanan, G., Bodik, P., Kandula, S., Akella, A., Bahl, P., & Stoica, I. (2015).
 785 Low latency geo-distributed data analytics. In SIGCOMM.
 786
- Rabkin, A., Arye, M., Sen, S., Pai, V. S., & Freedman, M. J. (2014). Aggregation and 787 degradation in jetstream: Streaming analytics in the wide area. In *NSDI*.
- Rizou, S. (2013). Concepts and algorithms for efficient distributed processing of data streams. 789 University of Stuttgart. https://doi.org/10.18419/opus-3209 790

- 522
- 63. Ron, D., Singer, Y., & Tishby, N. (1993). The power of amnesia. In NIPS. 791 64. Ron, D., Singer, Y., & Tishby, N. (1996). The power of amnesia: Learning probabilistic 792 automata with variable memory length. Machine Learning, 25(2-3), 117-149. 793 65. Rong Li, X., & Jilkov, V. P. (2003). Survey of maneuvering target tracking. part i. dynamic 794 models. IEEE Transactions on Aerospace and Electronic Systems, 39(4), 1333–1364. 795 66. Rong Li, X., & Jilkov, V. P. (2005). Survey of maneuvering target tracking. part v. multiple-796 model methods. *IEEE Transactions on Aerospace and Electronic Systems*, 41(4), 1255–1321. 797 67. Samoladas, V., & Garofalakis, M. N. (2019). Functional geometric monitoring for distributed 798 streams. In EDBT. 799 68. Shah, M. A., Hellerstein, J. M., Chandrasekaran, S., & Franklin, M. J. (2003). Flux: An 800 adaptive partitioning operator for continuous query systems. In ICDE. 801 69. Silva, J., Faria, E., Barros, R., Hruschka, E., Carvalho, A., Gama, J. (2013). Data stream 802 clustering: A survey. ACM Computing Surveys, 46(1), 1-31. 803 70. Tang, H., Lian, X., Yan, M., Zhang, C., Liu, J. (2018). D²: Decentralized training over 804 decentralized data. In ICML. 805 71. Vilalta, R., & Ma, S. (2002). Predicting rare events in temporal domains. In ICDM. 806 72. Vulimiri, A., Curino, C., Godfrey, P. B., Jungblut, T., Padhye, J., Varghese, G. (2015). Global 807 analytics in the face of bandwidth and regulatory constraints. In NSDI. 808 73. Willems, F. M. J., Shtarkov, Y. M., & Tjalkens, T. J. (1995). The context-tree weighting method: 809 basic properties. IEEE Transactions on Information Theory, 41(3), 653-664. 810 74. Zhao, B., Hung, N. Q. V., & Weidlich, M. (2020). Load shedding for complex event processing: 811 Input-based and state-based techniques. In: ICDE. 812 75. Zhou, C., Cule, B., & Goethals, B. (2015). A pattern based predictor for event streams. *Expert* 813 Systems with Applications, 42(23), 9294–9306. 814 76. Zillner, S., Bisset, D., Milano, M., Curry, E., Robles, A.G., Hahn, T., Irgens, M., Lafrenz, 815 R., Liepert, B., O'Sullivan, B., & Smeulders, A. (Eds.). (2020). Strategic research, innovation 816 and deployment agenda - AI, data and robotics partnership. Third Release. Brussels. BDVA, 817 euRobotics, ELLIS, EurAI and CLAIRE (September 2020). 818 77. Zillner, S., Curry, E., Metzger, A., Auer, S., & Seidl, R. (Eds.). (2017). European big data 819 value strategic research & innovation agenda. Big Data Value Association. 820
- 78. Zissis, D., Chatzikokolakis, K., Spiliopoulos, G., & Vodas, M. (2020). A distributed spatial 821 method for modeling maritime routes. *IEEE Access*, 8, 47556–47568.
 822
- 79. Zissis, D., Chatzikokolakis, K., Vodas, M., Spiliopoulos, G., & Bereta, K.: A data driven 823 approach to maritime anomaly detection. In *MSAW* (2019).
 824

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 825 International License (http://creativecommons.org/licenses/by/4.0/), which permits use, sharing, 826 adaptation, distribution and reproduction in any medium or format, as long as you give appropriate 827 credit to the original author(s) and the source, provide a link to the Creative Commons licence and indicate if changes were made. 829

The images or other third party material in this chapter are included in the chapter's Creative 830 Commons licence, unless indicated otherwise in a credit line to the material. If material is not 831 included in the chapter's Creative Commons licence and your intended use is not permitted by 832 statutory regulation or exceeds the permitted use, you will need to obtain permission directly from 834 the copyright holder. 834



AUTHOR QUERIES

- AQ1. Please check edit to the sentence "In such cases, the challenge is to..." is fine.
- AQ2. Please check the edit to the sentence "However, they neither handle..." is fine.
- AQ3. Please check the sentence "Performance-wise, the synchronous policy..." for clarity.
- AQ4. Please note that the references in the list have been arranged alphabetically and corresponding citations have been changed accordingly. Please check.

uncorrected