BigDataOcean Project: Early anomaly detection from big maritime vessel traffic data

Konstantinos Chatzikokolakis, MarineTraffic, London/UK,
konstantinos.chatzikokolakis@marinetraffic.com
Dimitris Zissis, Department of Product & Systems Design Engineering,
University of the Aegean, Syros/Greece and MarineTraffic, London/UK, dzissis@aegean.gr
Marios Vodas, MarineTraffic, London/UK, marios.vodas@marinetraffic.com
Giannis Tsapelas, Decision and Support Systems Laboratory, National and Technical University of Athens, Athens/Greece, gtsapelas@epu.ntua.gr
Spiros Mouzakitis, Decision and Support Systems Laboratory, National and Technical University of Athens, Athens/Greece, smouzakitis@epu.ntua.gr
Panagiotis Kokkinakos, Decision and Support Systems Laboratory, National and Technical University of Athens, Athens/Greece, pkokkinakos@epu.ntua.gr
Dimitris Askounis, Decision and Support Systems Laboratory, National and Technical University of Athens, Athens/Greece, askous@epu.ntua.gr

Abstract

This paper discusses the concept and results of the BigDataOcean project, and specifically the anomaly detection pilot. While in the past, surveillance had suffered from a lack of data, current tracking technologies have transformed the problem into one of an overabundance of information, with needs which go well beyond the capabilities of traditional processing and algorithmic approaches. The major challenge faced today is developing the capacity to identify patterns emerging within huge amounts of data, fused from various sources and detecting outliers in a timely fashion, to act proactively and minimise the impact of possible threats. Within this context we first define an "anomaly", before proceeding to present the BigDataOcean anomaly detection service; a service for the classification and early detection of anomalous vessel patterns. The service makes use of state-of-the-art big data technologies and novel algorithms which form the basis for a service capable of real time anomaly detection.

1. Introduction

Today numerous maritime systems track vessels during their voyages across the sea. Such is the Automatic Identification System, a collaborative, self-reporting system that allows efficient exchange of navigational data between ships and shore stations, intended to be used primarily for surveillance and safety of navigation purposes in ship to ship use, ship reporting and vessel traffic services (VTS) applications (*ITU, 2014*). Beyond simple collision detection, researchers and scientists are finding out that these data sets provide a new range of possibilities for improving our understanding of what is happening or could happen at sea. As such "anomaly detection" has been identified by operators/analysts of the operational community as an important aspect requiring further research and development (Martineau & Roy, 2011). Anomaly detection can be understood as a method that supports situational assessment to build models of normal data and then attempt to detect deviations from the normal behaviour in observed data that thus, may be of interest for further investigation (*Riveiro et al. (2018), Laxhammar (2011), Brax (2011)*).

To date, the availability of a larger number of sensors does not guarantee a reduction of risk as promised, mostly due to the enormous volumes and the velocity of the data that these systems and operators are faced with. Forecasting complex maritime situations emerging, such as probable collisions or groundings, suspicious activities or vessels' spoofing their identity has become a challenging task for surveillance system operators. They are now exploring complex and heterogeneous data; a situation that may lead to other undesired consequences for the system like uncertainty or time constraint violations in the decision making, or undesired effects for the operator such as fatigue and cognitive overload (*Riveiro et al, 2018*). Time critical computing systems are systems in which the correctness of the system is dependent not only on the accuracy of the result produced, but also on the time in which it was computed; such systems include avionics and marine navigation systems, defence systems, command and control systems, robotics and an ever-increasing number of Internet of Things (IoT) applications. For applications such as navigation, surveillance and others, timeliness is a top priority; making the right decision regarding a collision avoidance manoeuvre is only useful if it is a decision made in due time.

Unfortunately, current state of the art techniques and technologies are incapable of dealing with these growing volumes of high-speed, loosely structured, spatiotemporal data streams that require real-time analysis in order to achieve rapid response times. In this paper we present an adaptation of the lambda architecture as developed in the context of the BigDataOcean, an EU-funded project aiming to improve data sharing and linking between enterprises and entities of the maritime domain and other domains. The project focus is to increase the integration of Big Data and Data Analytics frameworks in blue economy by exploiting the huge potential from cross-sectorial blue data applications and to deliver out-of-the-box value-added Big Data services for maritime applications using advanced queries and analytics. In the following sections we first attempt to accurately define an "anomaly", before proceeding to present the BigDataOcean anomaly detection service; a service for the classification and early detection of anomalous vessel patterns. The service makes use of state-of-the-art big data technologies and novel algorithms which form the basis for a service capable of real time anomaly detection. Preliminary results indicate the efficiency of the proposed methodology when detecting maritime incidents.

2. Problem definition and related literature

The International Maritime Organization (IMO) defines Maritime Domain Awareness (MDA) as "the effective understanding of anything associated with the maritime domain that could impact upon the security, safety, economy, or environment." (*IAMSAR, 2010*). Situational awareness refers to the knowledge of the elements in the maritime space necessary to make well-informed decisions as well as processes involving knowledge and understanding of the environment that are critical to those who need to make decisions within the complex sea area. According to *NATO (2007)* Maritime Situational Awareness (MSA) is defined as "The understanding of military and non-military events, activities and circumstances within and associated with the maritime environment that are relevant for current and future NATO operations and exercises where the Maritime Environment (ME) is the oceans, seas, bays, estuaries, waterways, coastal regions and ports". *Nimmich and Goward (2007)* give a pragmatic view of MDA importance for Maritime security, as well as its economic and social impact.

The concept of anomaly is an important building block for developing situational awareness of the sea environment (*Snidaro et al, 2015*). An anomaly can be considered a critical event to which the system is generally called to react to. Usually, a threshold establishes if input data can be considered unexpected or anomalous, thus raising an exception. The concept of anomaly has a different meaning, depending on the context used as well as the requirements (*Roy, 2008*). *Roy and Davenport (2009)* present a categorisation, based on a taxonomy of the maritime situational facts involved in anomaly detection identified and validated through knowledge acquisition sessions with experts. Over the last century, a large number of diverse maritime situation awareness systems have been developed with different objectives and characteristics depending on what "an anomaly" means for its stakeholder and what are the operational needs. Typical technological systems that are used for Maritime Domain Awareness include Automatic Identification System (AIS), satellites, long range radars and long range Unmanned Aerial Vehicles (UAV). Output data from these systems are being used to provide real-time status of the observed sea environment, early warnings as well as data analytics for decision and policy making. The availability of plethora of sensors, higher data storage capacity, cheaper devices and better database management systems have made it possible to access huge volumes of data related to the Maritime Domain Awareness (*Riveiro and Pallotta, 2018*).

The exploitation of such volumes of data present new opportunities in a large number of applications, including safe vessel traffic management and collision prevention (*Perera et al., 2012*), coastal protection (*Rabasa et al, 2012*), environment protection and safety (*Roarty et al., 2013*), search and rescue missions (*Breivik et al, 2013*), anti-terrorism activities as well as the prevention of illegal activities such as smuggling and piracy (*Axbard, 2016*).

The sheer amount of data that needs to be processed in order to provide real or near-real-time anomaly detection requires the combination of various steps including typical big data computational models with machine learning algorithms (*Abielmona and Rami, 2013*). For instance, *Li et al. (2006)* indicated four steps required to tackle the problem of anomaly detection for moving objects, including microclustering and Support Vector Machine for classification purposes.

Laxhammar (2008) applied a Gaussian Mixture Model as the cluster model and the Expectation-Maximization algorithm as the clustering algorithm, as an extension of the model suggested by *Holst and Ekman (2003)*. The surveillance area was divided into grid cells and for each cell the points' velocities were modelled by a two-dimensional Gaussian distribution. *Korb and Nicholson (2004)* applied the CaMML machine learning tool on the AIS data from the Australian Defence Science and Technology Organization (DSTO). *Lee et al. (2007)* proposed a trajectory clustering algorithm to find similar, or normal patterns. A trajectory was divided into a series of line segments in the partition step which were then clustered by applying the density-based clustering algorithm DBSCAN.

Zhen et al. (2017) presented a similarity measurement method between vessel trajectories based on the spatial and directional characteristics of AIS data. *Zhang et al. (2017)* then applied the method of hierarchical and k-medoids clustering to model and learn the typical vessel sailing pattern within harbour waters. The Naïve Bayes classifier of vessel behaviour was built to classify and detect anomalous vessel behaviour. *Jousselme el al. (2015)* presented the development and implementation of fusion algorithms for maritime anomaly detection, and the definition of associated criteria and measures of performance. *Jousselme et al. (2015)* suggested that adequate uncertainty representation and processing is crucial for this higher-level task where the operator analyses information correlating with his background knowledge.

Those systems are focused on specific scenarios (e.g., collision prevention) and are mostly applied in historical data, in which they attempt to discover vessels' operational pattern. They do not focus on solution's scalability or model's adaptability to "new data", making thus evident the need for more transparent, interpretable and explainable machine learning-based systems (*Riveiro et al. 2018*). In this paper we propose a novel architecture capable to overcome the impediments of huge amount of data arriving at the system with high velocity and detect in near real-time various types of anomalies. Furthermore, the proposed anomaly detection service is modular in the sense that the set of possible anomalies can be further extended without adding computation burden to the system.

3. Approach and preliminary results

In this section we present the design and implementation options selected for the proposed anomaly detection services and we provide the preliminary assessment of the efficiency of those services against specific maritime incidents.

3.1. System architecture

In our approach we have deployed a modified Lambda architecture shown in Fig. 1 below. This scheme allows the decoupling of batch processing (usually performed upon historical data) and real-time analysis, which typically exploits the knowledge extracted from the batch processing. Specifically, in our approach the BigDataOcean batch layer performs the analysis of historical positional data of vessels and extracts port-to-port routes. This is a long-running process which takes several hours to complete. Once completed, the extracted routes are fed into the real-time layer in order to accommodate detection of vessel anomalies in real-time. Upon detection those incidents are displayed to the end user through the service layer. Furthermore, historical data are sent from this layer back to the batch layer at specific time intervals defined from the seasonality of the data, thus replacing previous processed routes with new ones.

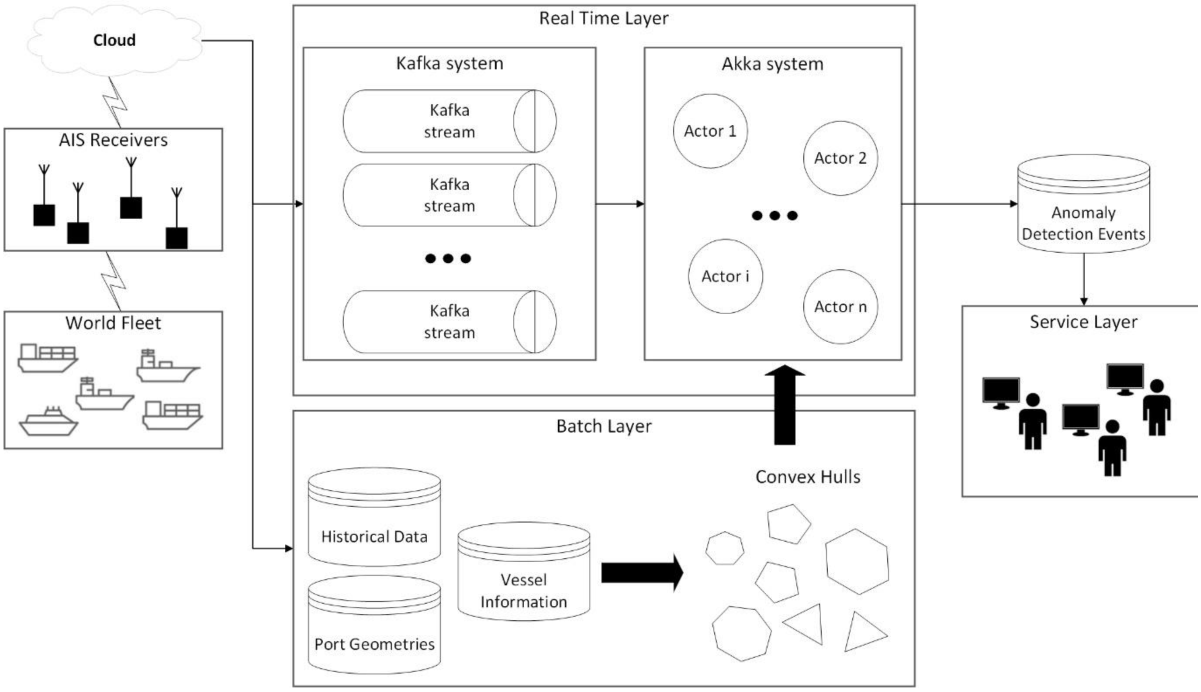

Fig.1: Modified Lambda architecture for the Anomaly Detection service[1]

3.1.1. Batch layer analysis

The BDO batch layer undertakes the role to identify the "safe" route between two ports; a knowledge that can be used to detect vessels travelling out of this "safe" path. This is achieved through data analysis techniques performed upon a huge amount of historical data that have been collected through an AIS network. The batch layer calculates "safe" routes between each pair of ports globally by correlating vessels' positional information (i.e., geographical coordinates broadcasted through AIS) with port geometries so as to link all the positions transmitted from the vessels with a specific port-to-port connection. A "safe" route is a set of convex hulls; each one is indicating an area with dense AIS data transmissions and is being produced through clustering of positional data of vessels travelling from the departure port towards the destination port. Specifically, in our approach we have used K-Means, a clustering algorithm which partitions the data points in groups according to their spatial proximity and density. A normal route is then created as the collection of the convex hulls formed when computing the minimum polygon that encloses all geographical positions in a group. The convex hulls represent the confidence interval, which means that if a vessel that performs the specific port-to-port voyage is travelling out of a convex hull, then an anomaly event is raised. Upon completion of calculation, these routes are fed to the real-time layer, which is then capable to detect possible anomalies in streaming data. The operation of convex hull calculation is executed in sparse time intervals in order to update the model when a significant amount of new data is gathered.

3.1.2. Real Time Layer

In this layer queries are performed on streaming and previously unseen data, thus enabling detection of security incidents in near-real time. More specifically, streaming positional and voyage-related AIS data are consumed and combined with static datasets and data mining models in near real-time to detect anomalous events. This layer includes an Apache Kafka distributed platform to which AIS messages are forwarded from the company's station network. Then, those messages are processed and divided into multiple topics based on the nature of each AIS message (which is identified based on the message type), giving thus, separate topics for the position reports (i.e. messages of type 1, 2 or 3), for the static and voyage related data (i.e., messages of type 5), etc. These topics are consumed by another component

[1] World Fleet icons adapted from https://www.vecteezy.com

responsible to perform distributed computing and identify the maritime security incident types mentioned afore. This component is based on Akka, a free open-source toolkit capable to build powerful, reactive and concurrent applications. It uses actors that evaluate whether the conditions that constitute a situation as "anomaly" exist, once a message is received.

Our platform is currently able to detect in near real-time a collection of four distinct types of possible anomalies, namely route deviations, AIS switch-offs, imminent collisions and groundings. In the following section we define each event type and discuss the algorithms we used in more detail.

- A route deviation happens every time a ship is found outside the normal route it is expected to follow, according to the departure and destination port. The model produced in the batch layer is loaded to the real-time service in order to be used for checking if a position coming from the stream of AIS is within the boundaries the convex hulls dictate. Thus, if a ship deviates at some part of its trajectory our service will detect it and provide it as an indicator that something might be wrong to the service layer.
- An AIS switch off happens when the service stops receiving data from a ship whilst the ship is well within the coverage of AIS base stations. Even though ships (especially large ones) are required by IMO to have their AIS transmitter on, gaps in AIS data are a very frequent phenomenon and they could be a result of many factors, e.g. malfunction of the AIS receiver or of the vessel's AIS transmitter, signal loss due to environmental noise, and intentional switch off. Intentional switch off usually indicates that the ship is about to engage in an illegal activity, except from cases where it passes through areas that have increased piracy, e.g. Somalia, and the switch off is used as a measure of protection.
- We define an imminent collision as an unsafe proximity event between at least two vessels. To handle the complex and demanding computations of determining proximity events, we have followed a distributed approach by partitioning areas of the Earth into a set of identifiable grid cells and using separate Akka actors to monitor each grid cell. Each position message consumed from the Kafka stream is forwarded to the actor responsible to monitor the corresponding area. The actor stores previously received messages and uses the newly arrived message's timestamp to forecast the position of the other vessels in the grid cell based on their past messages. This is achieved by projecting the most recent position of each vessel that the actor is monitoring at the exact time of the newly received message. Then, using the course over ground of those projected positions we determine whether their course is intersecting with the course of the vessel from which the new message was received and, in such case, a new imminent collision event is yielded.
- Groundings occur when a vessel has travelled in swallow sea (i.e. the sea depth was less than the ship's draught). In most cases this would also mean that the vessel has travelled out of the "safe route" and thus a route deviation event would be yielded before the grounding. To identify groundings the real-time layer considers vessel's position and navigation behaviour. The vessel may report through the navigation status field of AIS messages that is not be under command, or its ability to manoeuvre is restricted, or even explicitly report it is aground. Similarly, rapid decrease in vessel's speed, or frequent changes of course over ground in short time also indicate anomalous behaviour of the vessel. This information is correlated with bathymetry data and vessel's draught reported through AIS to increase the accuracy of detected groundings.

In the following subsection we present preliminary results of our approach for anomaly detection. The streams of data are simulated using historical tracks of verified incidents retrieved from the EMSA reported incidents[2]. Our experiments highlight the performance of our approach through examining characteristic examples of such security incidents.

3.2. Preliminary results

In this section we provide the preliminary assessment of the proposed architecture in terms of execution speed and present two indicative cases in which our algorithms achieved early detection of the vessels' anomalous behaviour. The experimental evaluation was performed on a machine with 12 threads (6 cores), an AMD Ryzen™ 5 2600 CPU @ 3.4 GHz and 16 GB of RAM, running Ubuntu 18.04 with

[2] http://www.emsa.europa.eu/

Java OpenJDK 8, Scala 2.12.8 and Akka 2.5.21. Real-time architectures are evaluated in terms of data ingestion delay, CPU and memory usage. The number of actors created in our system is proportional to the number of vessels from which messages are received. This leads to the conclusion that the number of actors created is finite, and thus, increasing the memory and the processing power of the system will alleviate any stress. Thus, the only factor that may affect our system is the volume and the velocity of incoming data that the system should cope with. For our experiments we used a dataset which contains 20M AIS messages received from 4599 vessels sailing in the Greek seas from February 2016 until July 2017 to assess the system's capability to ingest data. In order to test the system's scalability against different volumes of incoming data, we created two more subsets from the dataset containing 1/3 and 2/3 of the initial AIS messages respectively and these were streamed into our real-time layer. We tested multiple configurations of our system with 2 threads, 4 threads, 8 threads and 12 threads for each dataset and the results of our experiments are highlighted in Table I below. Preliminary results demonstrate that our approach achieves under second ingestion time even with 2 threads. In terms of velocity the data are streamed from files, having orders of magnitude higher speed compared to the actual system where the data will be fed into the real-time layer from the receivers through the Internet and thus will have much slower pace.

Table I: Data ingestion results (total execution time in seconds)

#threads	1/3 Subset (seconds)	2/3 Subset (seconds)	Full Dataset (seconds)
2	147	277	841
4	86	159	550
8	60	106	267
12	51	101	161

Besides the system's performance we also evaluate the accuracy of our approach, against two case studies of maritime security incidents. Fig. 2 shows a visualization of a part of the trajectory of the tanker Ovit. This vessel is a Malta registered tanker built in 2011 that ran aground on the Varne Bank in the Dover Strait, England, in the early morning of 18 September 2013, while carrying vegetable oil from Rotterdam to Brindisi. The vessel was using ECDIS for navigation and its passage plan, which was prepared by an inexperienced junior officer, was unsafe as it passed directly over the Varne Bank. The lack of experience of the Officer Of the Watch (OOW) and the false configuration of the ECDIS safety settings resulted in a 19 minute delay till he understood that the vessel was aground and the inability to recover vessel's historical track from the system. Furthermore, Dover coastguard watch officer that was operating the Channel Navigation Information Service was also not qualified and did not issue the warning to Ovit. Under those circumstances Ovit ran aground at 04:34 on 18 September 2013 in the Dover Strait. Fig. 2 (a) below highlights the convex hulls that were produced through the batch-layer for this port-to-port connection (i.e. from Rotterdam to Brindisi) and the past track of the vessel. Black triangles are reported positions where the ship stopped, and white circles are the ones it was moving. It is evident from Fig. 2 (b) and (c) that the ship was either moving outside or on the edge of the normal route a few minutes before running aground. After that the vessel changed its course and headed to the port of Dover.

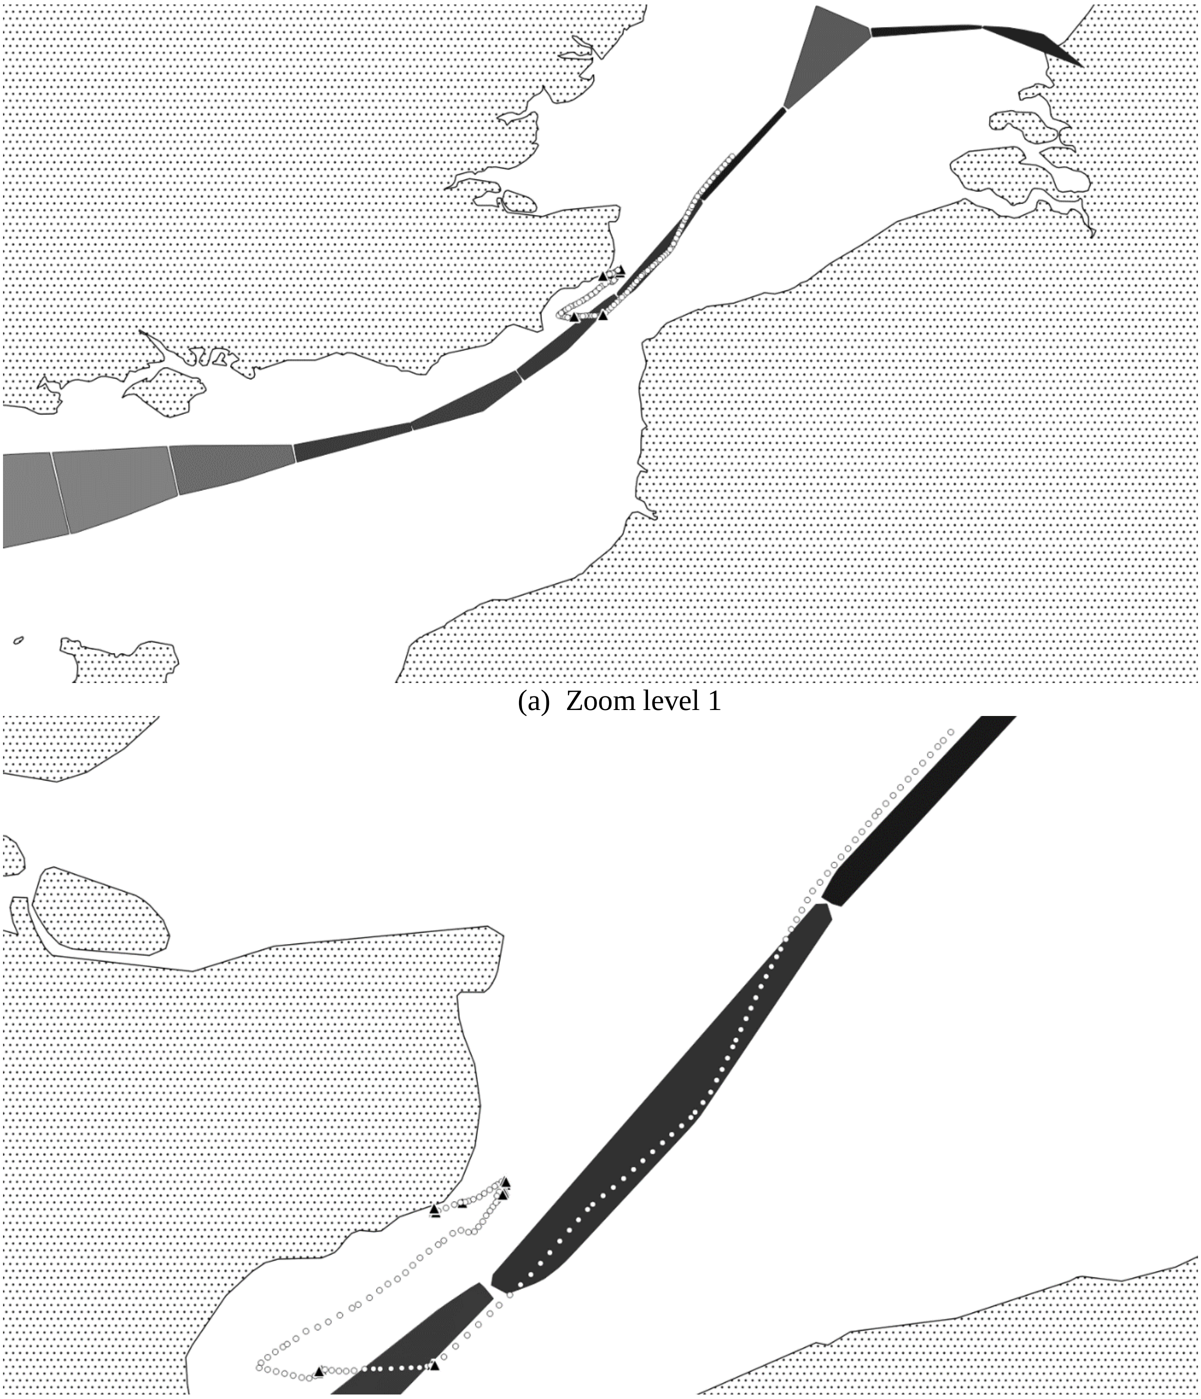

(a) Zoom level 1

(b) Zoom level 2

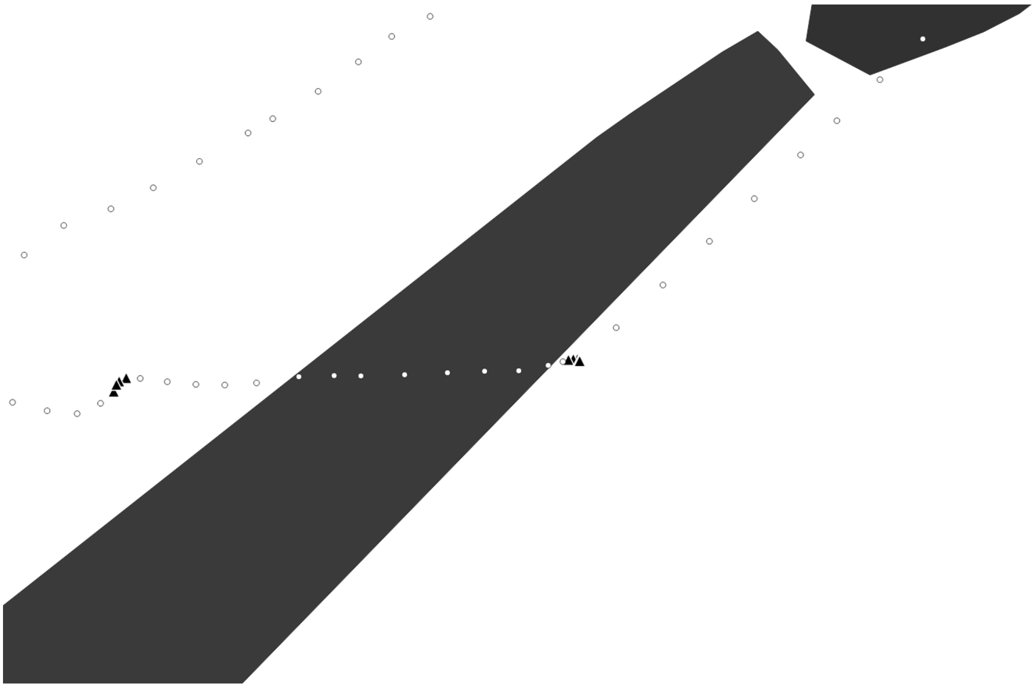

(c) Zoom level 3

Fig. 2: Grounding of oil/chemical tanker Ovit while travelling from Rotterdam to Brindisi[3].

In another example M/V Goodfaith, a bulk carrier sailing under the flag of Cyprus started its last voyage on the 10th of February 2015 from Elefsis, Greece heading to Odessa, Ukraine. The vessel was sailing in ballast condition and was in perfect shape as its special survey maintenance operations were successfully completed on the 9th of February. Its voyage plan included standard passages followed by vessels sailing in the Aegean, Marmaras and Black Sea, but bad weather conditions were prevailing in the area of South Evvoikos and Kafireas Strait. The vessel gradually lowered the engine speed to avoid main engine overspeed, but as the vessel was reaching Kafireas strait the weather conditions had significantly worsened with wind force 9Bfrs and high waves and by the time she was crossing the strait (i.e., at approximately 01:00), Goodfaith was heavily drifting towards the north-west coast of Andros island, where it ran aground at 01:28. Fig. 3 highlights part of the trajectory of the Goodfaith. It is evident from Fig. 3 (a) below that the vessel deviated from the convex hullsof this port-to-port connection much earlier than the grounding incident. Contrary to the Ovit case, this deviation was on purpose, to avoid bad weather conditions. Nevertheless, such deviation yielded early notification of possible anomaly in our system. Finally, Fig. 3 (c) shows the last part of the vessel's route few minutes before running aground.

[3] https://www.gov.uk/maib-reports/grounding-of-oil-chemical-tanker-ovit-on-the-varne-bank-in-the-dover-strait-off-the-south-east-coast-of-england

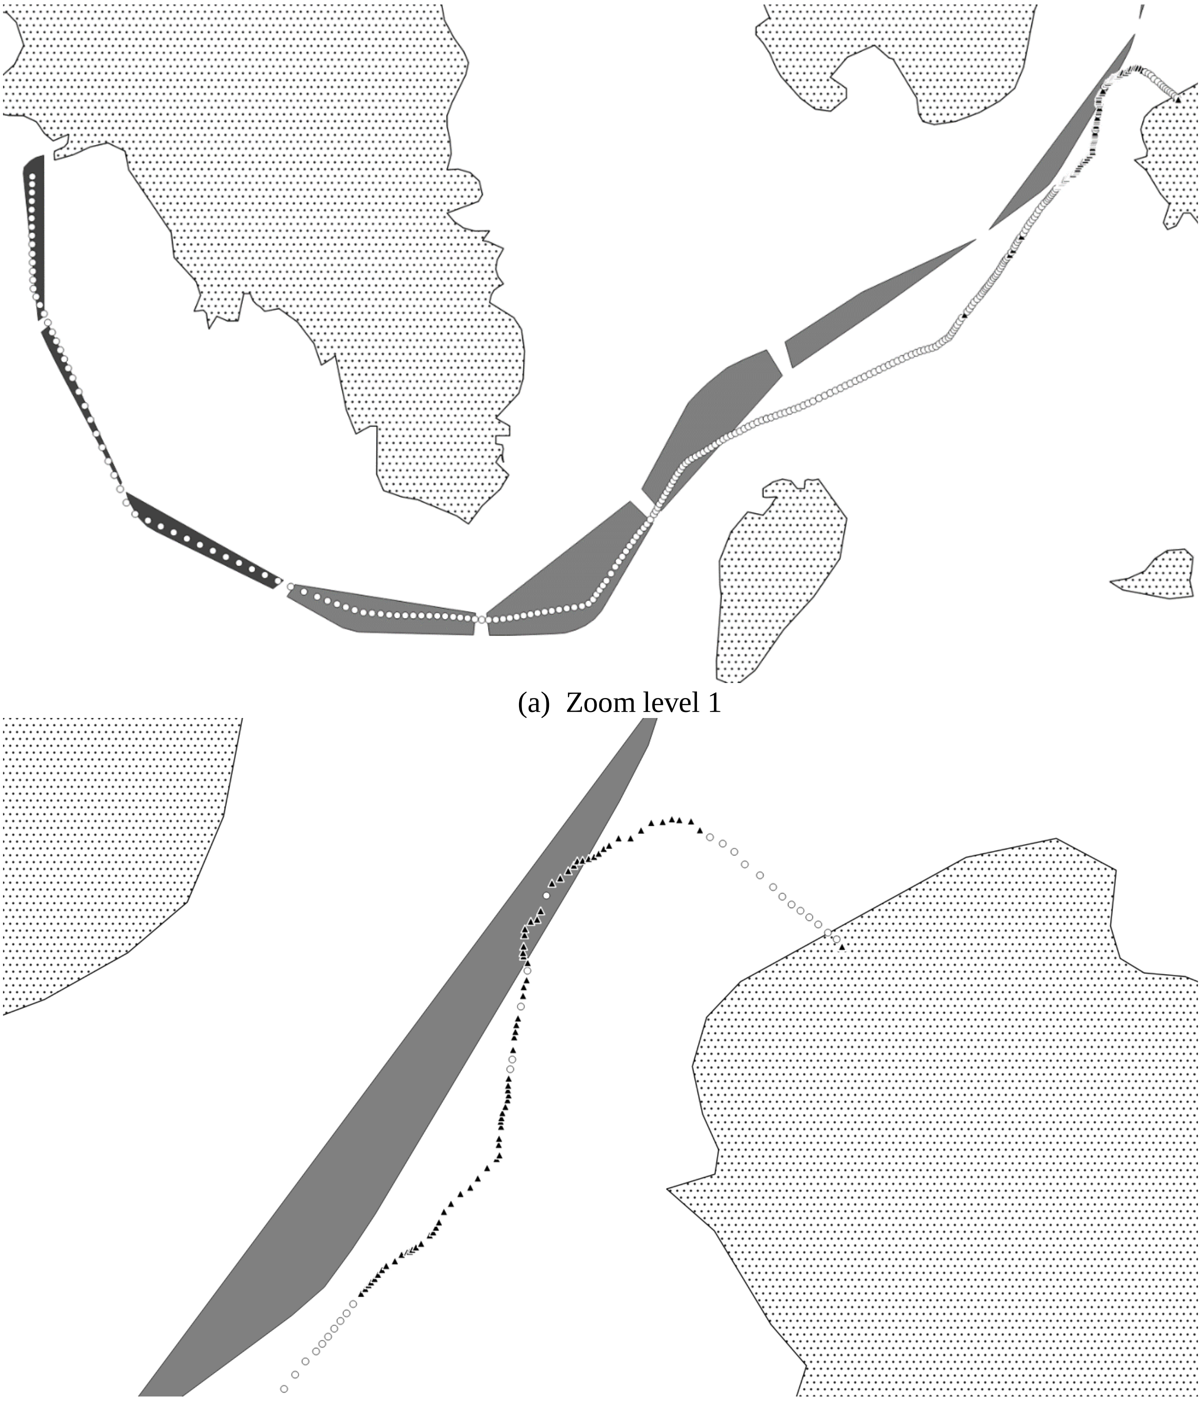

(a) Zoom level 1

(b) Zoom level 2

Fig. 3: Grounding of bulk carrier M/V Goodfaith while travelling from Elefsis, Greece to Odessa, Ukraine[4].

4. Conclusions

In this work we presented the architectural approach to building a real time maritime anomaly detection service, as deployed in the context of the EU funded BigDataOcean project. We show how big data challenges can be overcome to increase our understanding of maritime traffic and improve our monitoring of vessel activities in real time. In addition, we demonstrate the capability of the proposed service to ingest data of high volume and velocity and present some example case studies regarding real

[4] http://www.hbmci.gov.gr/js/investigation%20report/final/07-2015%20GOODFAITH.pdf

world maritime anomalies and evaluate the ability of the presented service to detect these in due time.

References

ABIELMONA, R. (2013), *Tackling big data in maritime domain awareness*, Vanguard, pp. 42-43

AXBARD, S. (2016), *Income Opportunities and Sea Piracy in Indonesia: Evidence from Satellite Data*, American Economic Journal: Applied Economics 8.2, pp. 154-94

BRAX, C. (2011), *Anomaly detection in the surveillance domain*, Diss. Örebro universitet

BREIVIK, Ø. et al. (2013) *Advances in search and rescue at sea*, Ocean Dynamics 63/1, pp. 83-88

HOLST, A.; EKMAN, J. (2003), *Anomaly detection in vessel motion*, Internal Report Saab Systems, Järfälla, Sweden

IAMSAR (2010), *Amendments to the International Aeronautical and Maritime Search and Rescue Manual*

ITU Radiocommunication Sector (ITU-R), (2014), *Technical characteristics for an automatic identification system using time division multiple access in the VHF maritime mobile band* , Recommendation M.1371-5

JOUSSELME, A.L.; PALLOTTA, G. (2015), *Dissecting uncertainty-based fusion techniques for maritime anomaly detection*, Information Fusion (Fusion), 2015 18th International Conference on. IEEE, pp. 34-41

KORB, K.B.; NICHOLSON, A.E. (2004), *Bayesian Artificial Intelligence*, Chapman & Hall/CRC Press

LAXHAMMAR, R. (2008), *Anomaly detection for sea surveillance*, In The 11th International Conference on Information Fusion, pp. 55-62

LAXHAMMAR, R. (2011), *Anomaly detection in trajectory data for surveillance applications*, Diss. Örebro universitet

LEE, J.G.; HAN, J.; WHANG, K.Y. (2007), *Trajectory clustering: a partition-and-group framework*, ACM SIGMOD International Conference on Management of data, pp. 593-604

LI, X.; HAN, J.; KIM, S. (2006), *Motion-Alert: Automatic anomaly detection in massive moving objects*, In Proceedings of the 2006 IEEE Intelligence and Security Informatics Conference (ISI 2006), Berlin, Springer, pp. 166-177

MARTINEAU, E.; ROY , J. (2011), *Maritime anomaly detection: Domain introduction and review of selected literature*, No. DRDC-VALCARTIER-TM-2010-460. DEFENCE RESEARCH AND DEVELOPMENT CANADA VALCARTIER (QUEBEC)

NATO (2007), *NATO Concept for Maritime Situational Awareness*, MCM-0140

NIMMICH, J.L.; GOWARD, D.A. (2007), *Maritime Domain Awareness: The Key to Maritime Security*, Int'l L. Stud. Ser. US Naval War Col., 83: 57

PERERA, L.P.; OLIVEIRA, P.; SOARES, C.G. (2012), *Maritime traffic monitoring based on vessel detection, tracking, state estimation, and trajectory prediction*, IEEE Transactions on Intelligent Transportation Systems 13.3, pp. 1188-1200

RABASA, A.; CHALK, P. (2012), *Non-Traditional Threats and Maritime Domain Awareness in the Tri-Border Area of Southeast Asia: The Coast Watch System of the Philippines*, RAND NATIONAL DEFENSE RESEARCH INST SANTA MONICA CA

RHODES, B.J.; BOMBERGER, N.A.; SEIBERT, M.C.; WAXMAN, A.M. (2005), *Maritime situation monitoring and situation awareness using learning mechanisms*, Military Communications Conference, Atlantic City, NY, USA

RIVEIRO, M.; PALLOTTA, G.; VESPE M. (2018), *Maritime anomaly detection: A review*, Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery 8.5: e1266

ROARTY, H.J. et al. (2013), *Expanding maritime domain awareness capabilities in the Arctic: High frequency radar vessel-tracking*, Radar Conference (RADAR), pp. 1-5

ROY, J. (2008), *Anomaly detection in the maritime domain*, Optics and Photonics in Global Homeland Security IV vol. 6945, International Society for Optics and Photonics

ROY, J.; DAVENPORT, M. (2009), *Categorization of maritime anomalies for notification and alerting purpose*, NATO workshop on data fusion and anomaly detection for maritime situational awareness, La Spezia, Italy, pp. 15–17

SNIDARO, L.; VISENTINI, I.; BRYAN, K. (2015), *Fusing uncertain knowledge and evidence for maritime situational awareness via Markov Logic Networks*, Information Fusion 21, pp. 159-172

ZHANG, Z. et al. (2017), *Review on Machine Learning Approaches in Maritime Anomaly Detection Based on AIS Data*, Electrical Engineering and Automation: Proceedings of the International Conference on Electrical Engineering and Automation (EEA2016), pp. 880-887

ZHEN, R. et al (2017), *Maritime anomaly detection within coastal waters based on vessel trajectory clustering and Naïve Bayes Classifier*, The Journal of Navigation 70.3, pp. 648-670