



# Streaming Analytics in Edge-Cloud Environment for Logistics Processes

Moritz Von Stietencron, Marco Lewandowski, Katerina Lepenioti, Alexandros Bousdekis, Karl Hribernik, Dimitris Apostolou, Gregoris Mentzas

## ► To cite this version:

Moritz Von Stietencron, Marco Lewandowski, Katerina Lepenioti, Alexandros Bousdekis, Karl Hribernik, et al.. Streaming Analytics in Edge-Cloud Environment for Logistics Processes. IFIP International Conference on Advances in Production Management Systems (APMS), Aug 2020, Novi Sad, Serbia. pp.245-253, 10.1007/978-3-030-57997-5\_29 . hal-03635614

**HAL Id: hal-03635614**

**<https://inria.hal.science/hal-03635614>**

Submitted on 30 Jun 2023

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License



# Streaming Analytics in Edge-Cloud Environment for Logistics Processes

Moritz von Stietencron<sup>1</sup>, Marco Lewandowski<sup>1</sup>, Katerina Lepenioti<sup>2</sup>, Alexandros Bousdekis<sup>2</sup>, Karl Hribernik<sup>1</sup>, Dimitris Apostolou<sup>2,3</sup> and Gregoris Mentzas<sup>2</sup>

<sup>1</sup> BIBA - Bremer Institut für Produktion und Logistik GmbH  
at the University of Bremen, Germany  
{sti,lew,hri}@biba.uni-bremen.de

<sup>2</sup> Information Management Unit (IMU), Institute of Communication and Computer Systems (ICCS), National Technical University of Athens (NTUA), Greece  
{klepenioti,albous,gmentzas}@mail.ntua.gr

<sup>3</sup>Department of Informatics, University of Piraeus, Greece  
dapost@unipi.gr

**Abstract.** The recent advancements in Internet of Things (IoT) technology and the increasing amount of sensing devices that collect and/or generate massive sensor data streams enhances the use of streaming analytics for providing timely and meaningful insights. The current paper proposes a framework for supporting streaming analytics in edge-cloud computational environment for logistics operations in order to maximize the potential value of IoT technology. The proposed framework is demonstrated in a real-life scenario of a large transportation asset in the aviation sector.

**Keywords:** data analytics, machine learning, predictive maintenance, aviation.

## 1 Introduction

Having an expected economic impact of \$2.7 to \$6.2 trillion in 2025 [1], Internet of Things (IoT) has already started to have a crucial role in every aspect of production, logistics and transportation operations [2,3]. Among others, a promising application deals with logistics operations, which usually consist of low value-added activities combined with a high degree of manual work. High information visibility achieved by real-time data processing has the potential to allow improved decision making and thus, increased efficiency and responsiveness [3]. The use of technologies for real-time condition monitoring and streaming data analytics for fleet maintenance can lead to optimization of maintenance costs, vehicle availability, operational efficiency, and lifetime [4]. Although the use of IoT technologies in public transportation is a well-studied area, the transportation means of manufacturing or product assets to maintenance providers are subject to different processes within the supply chain. Such operations pose significant challenges to the scheduling and planning decisions of the maintenance service providers. The current paper proposes a framework for supporting streaming analytics in edge-cloud computational environment for logistics opera-



tions in order to maximize the potential value of IoT technology. It consists of three main components: SENSE for data acquisition and edge analytics, DETECT & PREDICT for cloud analytics, and DECIDE for actions planning. The rest of the paper is organized as follows: Section 2 briefly presents the key enabling technologies of the proposed framework. Section 3 described the proposed framework, while Section 4 demonstrates a case study from the aviation sector. Section 5 concludes the paper and presents our plans for future work.

## **2 Key Enabling Technologies**

### **2.1 Streaming Analytics**

The fast development of IoT promotes an enormous amount of sensing devices that collect and/or generate massive sensor data streams over time for a wide range of fields and applications [5]. Data streams, i.e. sequences of unbounded tuples generated continuously in time, open new perspectives for extracting meaningful and timely insights from data-in-motion through streaming analytics (e.g. detect anomalies, predict future insights, make decisions) [1,5]. By bringing streaming data analytics closer to the source of data (i.e. IoT devices or edge devices) the need for data parallelism and incremental processing is less sensible as the size of the data in the source allows it to be processed rapidly. However, bringing fast analytics on IoT devices introduces its own challenges such as limitation of computing, storage, and power resources at the source of data [1].

### **2.2 Edge and Cloud Computing**

Cloud computing enables ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources that can be rapidly provisioned and released with minimal management effort or service provider interaction [6]. With the growing volume of data generated at the logical extremes of the network, and the fact that network bandwidth is not scaling at the same speed as with computing power, data mitigation is becoming a bottleneck constraining the cloud computing paradigm for delay-sensitive IoT services [7]. Consequently, data must be processed at the edge of the network for shorter response times, more efficient processing and significantly less network pressure [8]. Edge computing refers to the computation at the logical extremes of the network, enabling the conversion of raw signals to contextually relevant information in proximity of the data source [7].

## **3 The Proposed Framework for Streaming Analytics in Edge-Cloud Environment**


The proposed framework aims at supporting streaming analytics in edge-cloud computational environment for logistics operations in order to maximize the potential



value of IoT technology. It consists of three main components, described in the following sections: SENSE for data acquisition and edge analytics, DETECT & PREDICT for cloud analytics, and DECIDE for actions planning.

### 3.1 SENSE for Data Acquisition and Edge Analytics

SENSE performs data acquisition and edge analytics. It has two main levels denoted “SENSEcore” and “SENSEnode”, as shown in Fig. 1. Table 1 provides a summary of the hardware specifications.



**Fig. 1.** (a) SENSE connections; (b) SENSE subcomponents.

**Table 1.** SENSEcore technical specifications.

Product Specification	Technical Information
Processor/ Architecture	i.MX 6Quad / ARM Cortex-A9
Clock Frequency	4 x 1 GHz
Memory	1 GB DDR3 RAM 64 Bit, 1 GB NAND Flash, 16 MB NOR, 4 kB EEPROM
Hardware Interfaces	<ul style="list-style-type: none"> <li>• Ethernet: 1 x 10/100/1000 Mbps/s</li> <li>• USB</li> <li>• Serial: 1 x RS232</li> <li>• CAN Bus</li> <li>• I2C, 2 x SPI, UART, JTAG, SATA</li> </ul>
Power Supply	12-24 V
Wireless Connectivity	802.11 WLAN (WiFi) / Bluetooth
Operating System	Linux 4.1.36

The *SENSEcore* is one unit per monitored asset and consolidates all data from this asset. It also adds metadata for the integration of the asset data with the context model. It allows for performing edge data analysis and state detection and forwards the information set to the cloud.

The *SENSEnode* firmware acquires data from motion sensors (e.g. accelerometer, gyroscope, magnetometer), environment sensors (e.g. barometric pressure, humidity, light intensity, air temperature), etc. at a fixed sampling rate. The acquired raw data is processed on the sensor controller of the *SENSEnode* to derive measurable physical



quantities - e.g. raw value from accelerometer is processed to get acceleration in G's. These processed values are then converted from floating point to fixed point format in order to reduce the data bytes required to be persisted. The data is then mapped to a timestamp and persisted on a SD card. Each node captures and pre-processes specific sensor data from the asset and forwards this dataset to the SENSEcore unit.

Once data acquisition from the SENSEnodes is complete, the forwarding of the data to the cloud is facilitated by a data parsing script. The data parser reads the persisted data in and converts fixed point sensor values to floating point values. Then the sensor data along with device identification, sensor type and the measured physical quantity are batch uploaded to the cloud InfluxDB (in line format) as batches of 2,000 sensor data points.

### 3.2 DETECT & PREDICT for Cloud Analytics

DETECT & PREDICT performs cloud analytics in order to: (i) detect the current condition of the monitored asset and identify abnormal behaviors by continuously processing sensor data streams; and, (ii) estimate remaining useful life and predict future states. It has been developed using the OpenJDK 10 and it is deployed with an application server like Tomcat as the underlying server structure.

At design time, the user (data scientist) uses the following functionalities: (i) *Connection to sensor data streams and definition of queries*: The component has access to a time-series DB (based on InfluxDB) including discrete measurement points and additional meta-information describing the sensor stream, such as “location”, “installation point”, “reference to a technical equipment”, “type”, etc. Based on this connection, it provides the capability to define queries on the data streams using detection and prediction algorithms from various libraries. The query may provide a list of applicable sensor streams for a later analysis based on the provided meta-data, and a subset of the data within each previously selected sensor stream, e.g. based on a time range; (ii) *Engineering, instantiation and deployment of calculation flows*: The component provides a web-based interface in which the algorithms of the toolbox are orchestrated through the use of calculation flows, which can be considered similar to flowchart-based workflows, in order to be linked to continuous data streams and process them to more aggregated information. The component gives the opportunity to establish rules which define time- or event-based triggers for the later deployment.

At runtime, the component performs *deployment, monitoring and persistent storage of calculation flows* in order to collect the latest required data by executing the queries, execute the algorithms defined by the calculation flow and monitor the calculation flows. The outputs are communicated through an event message broker (Apache Kafka) for further processing.

### 3.3 DECIDE for Actions Planning

DECIDE implements a prescriptive analytics approach [9] in order to provide proactive recommendations and plans about optimal actions by receiving streams of predictions about future states (e.g. failures) through the event message broker.



At design time, the user (domain expert) inserts proactive decision rules, i.e. rules that correlate sets of predicted failure modes with sets of (alternative) proactive actions along with their parameters, in order to create the *knowledge base* (implemented with Drools Rule Engine). At runtime, the real-time predictions are received through subscription to prediction message topics and feed into the *inference engine* which retrieves and fires the related rules from the working memory associated with the knowledge base. On this basis, the *proactive decision making* functionality, implemented as a Kafka Topology and incorporating the Brown-UMBC Reinforcement Learning and Planning (BURLAP) library, embeds reinforcement learning algorithms in order to jointly minimize the maintenance costs and maximize the positive human feedback over the generated recommended actions.

The recommendations are accompanied with explainable graphs that visualize the decision model and the resulting policy. The recommended actions are either approved or rejected by the operator according to their knowledge, experience and preference. This *human feedback* is taken into account in two ways: (a) The approved actions formulate the actions plan to be followed once the vehicle arrives for maintenance; (b) It is taken into account in order to perform policy shaping and adapt the decision models for building optimized human-machine collaborations.

## 4 Case Study in Maintenance Service Logistics of Aviation

The proposed framework was validated in FFT Produktionssysteme GmbH & Co. KG, a company with core business, among others, the design, realization and maintenance of production systems in aerospace industry. Typical maintenance projects include both fixed (stationary production equipment, test rigs) and mobile assets (jigs and tools, transportation jigs). The case under consideration deals with the maintenance of mobile assets, which are subjected to a wide range of external influences affecting their condition, while they are maintained on geographically dispersed locations. Specifically, the type of mobile asset is a large transportation jig for major aircraft wing shells (Wing Upper Cover Transportation Jig - WUCTJ) (Fig. 2). It consists of a steel structure (the main jig) and an associated light-weight roof component (Top Weather Protection) made from aluminum profiles. It is 40m long with an empty weight of around 2,350 kg and is designed to transport wing covers of a commercial aircraft between plants of a transnational manufacturer in Germany and the United Kingdom. During its journey, the asset is moved by different means of transportation on land, water, and air (by Airbus Beluga).

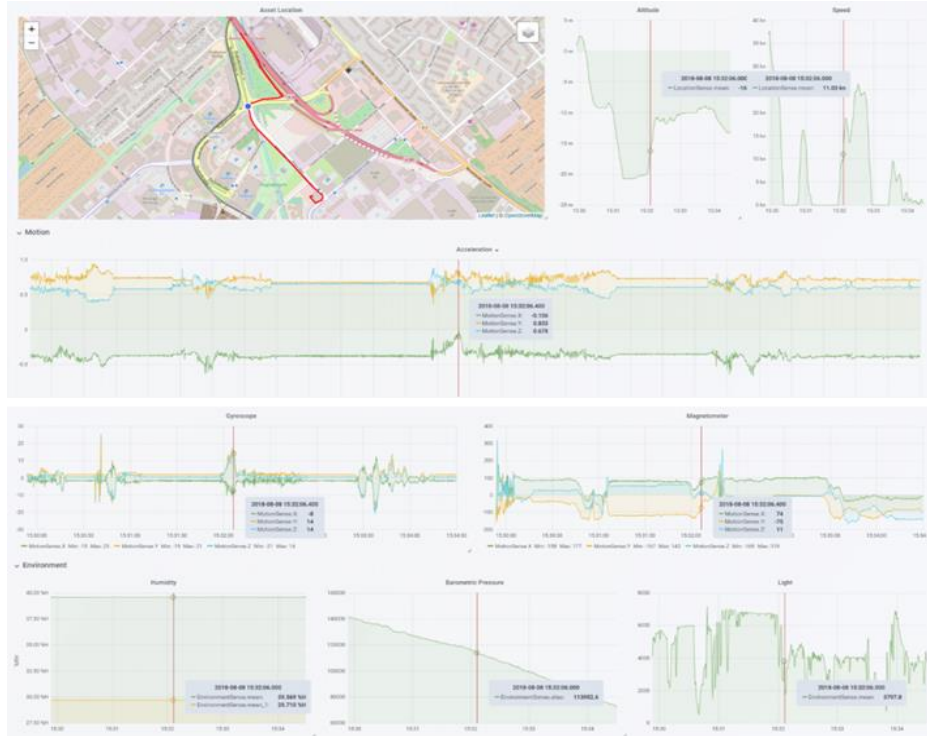
All these environments impose different kinds of stress on the asset. For example, since the cargo hold of the Beluga transport aircraft is not pressurized or air-conditioned, the transportation jig is subjected to extreme temperatures (-60 to +50 °C) additionally to a multitude of loads, shocks and weather influences. There are many situations where damages can occur (e.g. mishandling of the asset during the loading and unloading between different transportation modes). However, FFT, as the manufacturer and maintenance provider of the transportation asset, will see such damages only at the end of a transportation cycle and then, urgent repair activities will



need to be applied. Due to the fact that both production and logistics schedules can be significantly affected by the availability and plannability for transportation jig deployment, there is high benefit of having comprehensive condition information and proactive meaningful insights for maintenance operations of the jig.



**Fig. 2.** The transportation jig during loading at the Airbus Beluga.



**Fig. 3.** Real-time monitoring of the jig.

The SENSE component of the proposed framework acquires and pre-processes sensor data related to GPS location, temperature, humidity, gyroscope, acceleration, magnitude and light. To do this, ten SENSEnodes have been applied to the jig. Fig. 3



shows example graphs of the condition monitoring, as well as the map showing the drive and the current location of the jig. Then, the DETECT & PREDICT component implements an outlier detection algorithm in order to identify abnormal frequencies and amplitudes that lead to predictions about failure modes. Fig. 4 depicts an example of a pattern that, according to the stored predictive analytics models, leads to a prediction of the failure mode “Crack of the light-weight frame”. Based on this trigger, the DECIDE component generates the recommendation “Re-welding of light-weight frame” along with the explainable visualization of the decision model shown in Fig. 5a. Upon user approval, the action is inserted to the maintenance plan (Fig. 5b), and the decision model is updated accordingly.

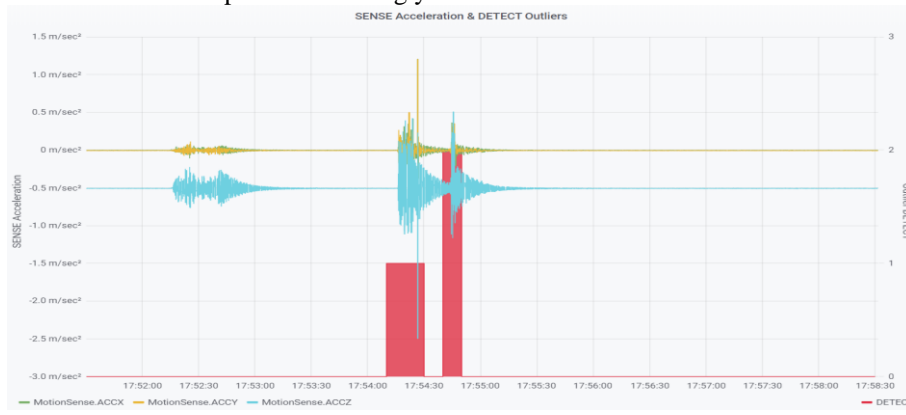


Fig. 4. Outlier detection of the jig health status.

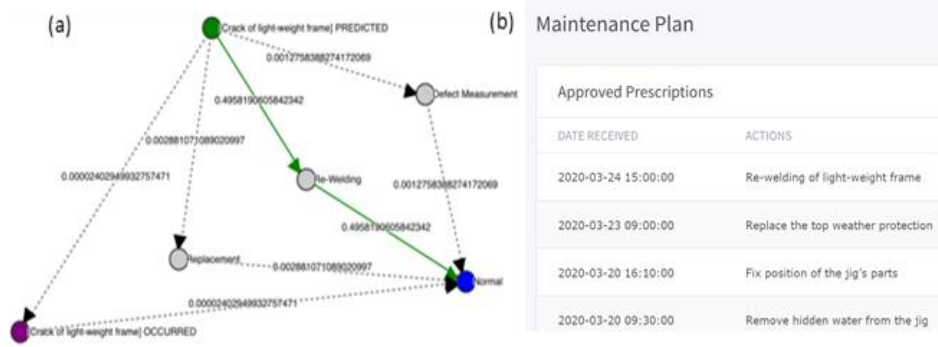


Fig. 5. (a) Explainable visualization of decision making; (b) Maintenance plan.

In order to perform a preliminary evaluation of the proposed framework on the process improvements, we performed simulations involving a typical WUCTJ transport cycle based upon the methodology described in [10]. Specifically, discrete event simulations were performed and compared under an “as-is” and a “to-be” scenario, using ARENA simulation software. The results showed, among others, that the “to-be” scenario has about half the Non-Value Added Cost (49,052 instead of 107,523



Euro), while its “efficiency gain” is more than double comparing to the “as-is”. This translates to less risk of production blocking events and the possibility to increase capacity utilization of the assets. The proposed framework was also evaluated by the domain experts with respect to its initially defined stakeholder requirements.

## 5 Conclusions and Future Work

In this paper, we proposed a framework for supporting streaming analytics in edge-cloud computational environment for logistics operations. The framework consists of three components: SENSE for data acquisition and edge analytics, DETECT & PREDICT for cloud analytics, and DECIDE for actions planning. It was validated in a case study of a transportation jig in the aerospace industry. In our future work, we plan to further validate our approach in additional use cases, as well as to take advantage of legacy data analytics that represent different levels of information.

**Acknowledgements.** This work has been funded by the European Commission project H2020 UPTIME “Unified Predictive Maintenance System” (768634).

## References

1. Mohammadi, M., Al-Fuqaha, A., Sorour, S., Guizani, M.: Deep learning for IoT big data and streaming analytics: A survey. *IEEE Communications Surveys & Tutorials* 20(4), 2923-2960 (2018).
2. Dai, H. N., Wang, H., Xu, G., Wan, J., Imran, M.: Big data analytics for manufacturing internet of things: opportunities, challenges and enabling technologies. *Enterprise Information Systems*, 1-25 (2019).
3. Zafarzadeh, M., Hauge, J. B., Wiktorsson, M., Hedman, I., Bahtijarevic, J.: Real-Time Data Sharing in Production Logistics: Exploring Use Cases by an Industrial Study. In: *IFIP International Conference on Advances in Production Management Systems*, pp. 285-293. Springer, Cham (2019).
4. Killeen, P., Ding, B., Kiringa, I., Yeap, T.: IoT-based predictive maintenance for fleet management. *Procedia Computer Science* 151, 607-613 (2019).
5. Isah, H., Abughofa, T., Mahfuz, S., Ajerla, D., Zulkernine, F., Khan, S.: A Survey of Distributed Data Stream Processing Frameworks. *IEEE Access* 7, 154300-154316 (2019).
6. Miyachi, C.: What is “Cloud”? It is time to update the NIST definition?. *IEEE Cloud Computing* 3, 6-11 (2018).
7. Shi, W., Dustdar, S.: The promise of edge computing. *Computer* 49(5), 78-81 (2016).
8. Cao, H., Wachowicz, M.: An Edge-Fog-Cloud Architecture of Streaming Analytics for Internet of Things Applications. *Sensors* 19(16), 3594 (2019).
9. Lepenioti, K., Bousdekis, A., Apostolou, D., Mentzas, G.: Prescriptive analytics: Literature review and research challenges. *International Journal of Information Management* 50, 57-70 (2020).
10. Nikhil, B. H.: Modelling and Simulation of an Industry 4.0 based Predictive Maintenance system in an Aerospace value chain, Vellore Institute of Technology, School of Mechanical Engineering, Vellore (2019).