



## Enabling adaptive analytics at the edge with the Bi-Rex Big Data platform

Riccardo Venanzi<sup>c</sup>, Simon Dahdal<sup>b,\*</sup>, Michele Solimando<sup>c</sup>, Lorenzo Campioni<sup>b</sup>,  
Alberto Cavallucci<sup>c</sup>, Marco Govoni<sup>a,b</sup>, Mauro Tortonesi<sup>b</sup>, Luca Foschini<sup>c</sup>, Loredana Attana<sup>a</sup>,  
Matteo Tellarini<sup>d</sup>, Cesare Stefanelli<sup>b</sup>

<sup>a</sup> Bonfiglioli Riduttori, Via Cav. Clementino Bonfiglioli, 1, Calderara di Reno, 40012, BO, Italy

<sup>b</sup> Distributed Systems Research Group, University of Ferrara, Via Saragat 1, Ferrara, 44122, Italy

<sup>c</sup> Department of Computer Science and Engineering, University of Bologna, Viale Risorgimento, 2, Bologna, 40126, Italy

<sup>d</sup> SACMI, Via Selice Provinciale, 17/A, Imola, 40026, BO, Italy

### ARTICLE INFO

#### Keywords:

Zero defect manufacturing  
Industry 4.0  
Big data  
Edge computing  
MLOps  
Data-driven applications  
Machine learning

### ABSTRACT

Zero Defect Manufacturing (ZDM) is an emergent and disruptive paradigm that aims to optimize industrial process efficiency and sustainability by leveraging innovative and sophisticated data-driven approaches. It is a technology intensive concept that has the ambition of achieving and maintaining “first-time-right” quality goals in spite of varying processes and input material. As a result, developing ZDM applications might become overwhelming for small enterprises due to the multitude of diverse platform, the lack of know-how, and the need to adapt general purpose solutions to meet their needs. The Big Data Innovation and Research Excellence (Bi-Rex) is an Italian consortium that aims to accelerate the industrial innovation process of small enterprises. Within this consortium we developed a Big Data platform that enables adaptive analytics at the IT/OT boundary by leveraging innovative solutions for the safe and automatic deployment of data-driven apps, using MLOps and DevOps techniques and technologies, and evaluated it in real use cases provided by the world leading industrial partners involved in the project.

### 1. Introduction

Industry 4.0 defines the disruptive and revolutionary advancements enabled by IoT in industrial scenario (Boyes et al., 2018; Hofmann and Rüschi, 2017). In fact, IoT represents an extremely powerful, easily adoptable, and cheap solution to interconnect production machines and human operators, paving the way for new advanced production chains and a new era of time-critical, context-aware services for the industries (Corradi et al., 2019, 2021). The sensing capabilities of IoT devices, in particular, enable the acquisition of large amounts of data, which can then be stored and processed, thanks to big data technologies and techniques. This allows for extrapolation of information, the development of data-driven applications, and the optimization of manufacturing processes.

One of the most compelling yet challenging objectives of Industry 4.0 is *Zero Defect Manufacturing (ZDM)* (Powell et al., 2022). ZDM is a new paradigm that advocates the complete elimination of defects through the adoption of smart “predict and prevent” approaches. In this context, ZDM represents a disruptive evolution of quality management

strategies towards a “first-time-right” goal that considers customizable manufacturing processes and production planning as well as quality management and maintenance aspects.

ZDM is a technology intensive concept, that requires capillary sensing, continuous data acquisition, and Big Data platforms implementing sophisticated analytics by leveraging a range of ML and non-ML methods and techniques, including multilayer perceptrons, convolutional neural networks, recurrent neural networks, support vector machines, digital twins, etc. (Caiazzo et al., 2022). On top of this, the analysis of information collected from industrial devices and processes presents an additional set of peculiar challenges, due to imbalanced data sets and data quality issues (imperfectly labeled data, missing data, etc.). The significant efficiency and business opportunities involved addressing these requirements have attracted a multitude of providers that have proposed a wide range of proprietary and open source technologies. However, most (if not all of) these technological solutions are very sophisticated and present a steep learning curve and considerable

\* Corresponding author.

E-mail addresses: [riccardo.venanzi@unibo.it](mailto:riccardo.venanzi@unibo.it) (R. Venanzi), [simon.dahdal@unife.it](mailto:simon.dahdal@unife.it) (S. Dahdal), [michele.solimando@unibo.it](mailto:michele.solimando@unibo.it) (M. Solimando), [lorenzo.campioni@unife.it](mailto:lorenzo.campioni@unife.it) (L. Campioni), [alberto.cavallucci2@unibo.it](mailto:alberto.cavallucci2@unibo.it) (A. Cavallucci), [marco.govoni@bonfiglioli.com](mailto:marco.govoni@bonfiglioli.com) (M. Govoni), [mauro.tortonesi@unife.it](mailto:mauro.tortonesi@unife.it) (M. Tortonesi), [luca.foschini@unibo.it](mailto:luca.foschini@unibo.it) (L. Foschini), [Loredana.Attana@bonfiglioli.com](mailto:Loredana.Attana@bonfiglioli.com) (L. Attana), [matteo.tellarini@sacmi.it](mailto:matteo.tellarini@sacmi.it) (M. Tellarini), [cesare.stefanelli@unife.it](mailto:cesare.stefanelli@unife.it) (C. Stefanelli).

<https://doi.org/10.1016/j.compind.2023.103876>

Received 15 July 2022; Received in revised form 11 January 2023; Accepted 13 February 2023

Available online 2 March 2023

0166-3615/© 2023 The Authors. Published by Elsevier B.V. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

barriers to entry. They thus represent a serious challenge for Small-Medium Enterprises (SMEs) that might lack the sufficient know-how to properly exploit these new technologies. In addition, the significant cost of Research and Development (RnD) departments for business innovation can represent an obstacle for SMEs that dispose of a limited budget (Masood and Sonntag, 2020; Amaral and Peças, 2021).

To address this issue, as part of the Industry 4.0 government plan, the Italian Ministry of Economic Development sponsored the foundation of 8 national competence centers, i.e., public-private consortiums that attract universities and industries, each one dedicated to a specific thematic aspect of Industry 4.0. More specifically, the consortium aims to increase cooperation between public and private sectors and the transfer of skills, technologies, and resources with a strong drive towards open innovation.

Big Data Innovation and Research Excellence (Bi-Rex) is the competence center with a specialized focus on Big Data. Located in Bologna, the capital of Emilia-Romagna, a region of Italy which is well known for its industrial excellences - especially in the automotive and packaging sectors, Bi-Rex includes world known partners such Bonfiglioli Riduttori, SACMI, Aetna Group, IMA, Philip-Morris International, etc. Bi-Rex aims at contributing to the growth of the digital manufacturing enterprises from an Industry 4.0 perspective, fostering technological innovation processes through research and innovation grants co-funded by the Italian Ministry of Economic Development.

This paper presents a novel logical architecture that emerged from the joint work of 2 of the projects funded by Bi-Rex: Big Data for Manufacturing (BD4M) and Dynamic Edge computing for Plant MONitoring (DEEPMON). Those projects aimed at building a Big Data platform of general applicability within Industry 4.0 domain, respectively focusing on the 2 functional layer within a smart factory physical perimeter: the Information Technology (IT) layer and the Operation Technology (OT) layer. In addition, we further explored opportunities to implement adaptive analytics at the IT/OT boundary by realizing an innovative solution that enables the safe and automatic deployment of data-driven apps, more specifically Machine Learning-based applications, through their entire life-cycle using DevOps and MLOps techniques and technologies (Kreuzberger et al., 2022).

Achieving proper convergence between the IT and OT layers and using tools that help to implement the concept of ZDM drives the manufacturing sector toward the goals proposed by the current transition to Industry 5.0 (I5.0). This companion revolution brings three fundamental pillars that push the Industry 4.0 transition forward in the modern industrial landscape (Breque et al., 2021). As we will see in the following sections, our platform goes towards the goals proposed by I5.0, favoring human centrality in the evolutionary process, without neglecting the resilience of the resources and the sustainability of industrial progress.

The Bi-Rex Big Data platform enables a more agile development-to-production process of such applications enabling SMEs to benefit from the extrapolated information from the collected data, permitting the development of various useful applications for ZDM, ranging from the optimization of industrial process to reduce waste and improve sustainability to self-reconfigurable processes that aim to maintain “first-time-right” strategies in spite of varying conditions (input material, degraded machine efficiency due to wear and tear, etc.) through the use of AI and ML techniques (Powell et al., 2022). The comprehensive solution achieved by the Bi-Rex Big Data platform, that emerged as the integration of the outputs from the BD4M and DEEPMON projects, is being deployed and evaluated in real use cases provided by those projects’ world leading industrial partners.

## 2. Related works

Several publications have attempted to offer architectures for the industries to adopt ZDM. Magnanini et al. (2020) present a reference

architecture for industries considering implementing zero-defect manufacturing. The first step toward a data quality management system in ZDM is introduced in Caccamo et al. (2021) that proposes a hybrid design that lowers the limitations and constraints imposed by the needs of the earlier architecture. The four key strategies in the ZDM idea are detection, prediction, prevention, and repair which are presented in detail in Powell et al. (2021) and Caiazzo et al. (2022).

A widely accepted truance is that, in order to be put into practice, ZDM applications of the aforementioned tactics needs data analytics based on a range of ML and non-ML methods and techniques, including multilayer perceptrons, convolutional neural networks, recurrent neural networks, support vector machines, digital twins, etc., in order to extrapolate essential information to apply the ZDM strategies (Caiazzo et al., 2022). In turn, the effective implementation of analytics solutions requires Big Data platforms that are both functional and well structured, since they will be used to collect, store, and process data using a variety of algorithms and techniques.

One of the first work in this direction is presented in Gokalp et al. (2016). Here, Gokalp et al. present a framework that provides conceptual features to ease the adoption of Big Data techniques in future Industry 4.0 enterprises. Another conceptual framework along with its key technologies, applications and some practical scenarios are presented by Zheng et al. in Zheng et al. (2018). These two works are both valid and well done, but they still present a conceptual solution.

More in general, a panoramic and theoretic overview of digital manufacturing for Industry 4.0 is presented by Gerrikagoitia et al. in Gerrikagoitia et al. (2019). There are other relevant works in literature but they address specific use cases or focus in tackling specific aspects of Big Data for the manufacturing. For example, in Bolla et al. (2021), the Matilda platform is presented. This solution aims to enable vertical applications for Industry 4.0 by mainly focusing on the Cloud-5G connectivity aspects. In Villalonga et al. (2020), the authors present a data-driven method for updating edge components in a Industry 4.0 Big Data architecture. This work is very valuable and similar to a feature presented in our solution, but it is strictly focused on the presented method. In Jaskó et al. (2020) is presented a remarkable work strictly focused on MES functionalities and their requirements in Industry 4.0, one of the aspect addressed in our proposed solution. Two theoretical studies and literature reviews for Industry 4.0 paradigm adoption in SMEs are presented in Han and Trimi (2022) and Moeuf et al. (2018). While, on the other hand, two valuable examples of works with a more practical focus are Nikishechkin et al. (2020) and Sahal et al. (2020). In the first one, Nikishechkin et al. address the aspects, requirements and functionalities needed to develop and implement Industry 4.0 platforms for specific parameters monitoring scenarios. The latter, however, addresses Big Data and stream processing platforms for predictive maintenance use cases.

So far, the realization of Big Data solutions specifically designed to enable ZDM, and thus considering continuously changing production contexts, that demand fast react times, has received only limited attention from scientific literature. The development of innovative Big Data methodologies and tools designed to continuously adapt to changes in manufacturing processes, e.g., by re-training, re-deploying, or re-tuning analytics and decision making, represents a key step towards the fulfillment of the challenging ZDM vision and goals.

## 3. Leveraging DevOps and MLOps for adaptive data analytics in the edge

Recently, in Industry 4.0 scenarios the needs for major data flow control, industrial processes management, and IT/OT architecture re-configuration are increasingly assuming a crucial role. These needs are driven by the necessity to dynamically handle a continuously evolving manufacturing context, that demands prompt reactions to changes. For instance, heterogeneous raw input material, unexpected events on the production lines, possible hardware faults, local machine’s settings,

software component life-cycle, and so on, introduce dynamic aspects in the manufacturing process that are very difficult to face for any company, and even more for SMEs that cannot dispose of a large amount of resources and often adopt a very own custom IT solution.

The impressive technological advances, especially at the IT level, pave the way for the auto-reconfiguration of manufacturing processes and the widespread adoption of predict and repair techniques. For instance, if a production line or machine is using a specific sampling rate for a particular feature and its value begins to stray from the range, an ZDM-oriented analytics solution will generate an alert that will cause the dynamic sampling rate to increase because it is essential to better monitor the feature and avoid a serious fault. The system will reconfigure the machinery after confirming the source of the problem, mitigating and reducing any potential waste and product flaws. This machinery's settings reconfiguration aspect comes along with another emerging concept in Industry 4.0 scenarios, the locality. Often, it is necessary to change the settings just of a specific machine/production line within the whole shop floor. This reconfiguration has to be led by IT layer and deployed on the single OT edge component, locally. The IT layer has the global vision of the whole shop floor with edge nodes, and it is responsible to drive the reconfiguration and deployment operations on the OT level.

The application of machine learning techniques represents another remarkable Industry 4.0 scenario that highlights the necessity of continuous integration between IT and OT levels (see Fig. 1). Usually, the OT level is responsible of collecting data from the field, processing those data, storing data locally and finally uploading them to the IT level/Cloud. OT edge nodes do not have enough resources to train machine learning models, they need to receive the trained models from the IT layer, but the IT layer needs data from OT to train the models, hence the necessity to have a continuous integration and to adopt DevOps techniques between IT and OT emerges and becomes concrete.

The continuous integration and DevOps also enable the full software components' life-cycle management, scaling, and re-deployment. However, managing ML-powered apps differs from managing regular apps, as ML-powered apps are typically made up of a variety of software components and artifacts, as well as one or more Machine Learning trained models. This pushes us to employ MLOps strategies and techniques to manage and orchestrate the entire life-cycle of the machine learning models, so that more accurate monitoring solutions are needed for ML-powered apps in general and for the trained models in particular. Because of this, monitoring represents the single most crucial component of the machine learning process in production. Monitoring enables us to identify any potential problems, such as data and concept drift, model staleness, performance degradation, and other issues. These issues can be countered by an effective monitoring process of the following metrics: *input metrics*: data distribution, features range of value, etc.; *output metrics*: accuracy, precision, recall, F1-score, etc.; *operational metrics*: IO, CPU, Memory, throughput and latency of the ML model end point, etc.

### 3.1. The OT level features

As discussed the OT features are strictly related to edge devices and are related to low level problems to enable IoT.

#### 3.1.1. Connectivity and communication

Enabling communication between different edge devices and between devices that compose the Industry 4.0 informatic sensing infrastructure. Therefore represent one of the most important backbones of Industry 4.0 applications. In particular, this OT level feature includes all solutions and strategies that enable the data flow between the various components of the architecture and define common instruction set to interact with edge devices. For example, the OT level includes the communication protocol infrastructure, the data format adopted to exchange information, and also all the communication technologies adopted to enable connectivity. Moreover, the OT layer also enable the connectivity between the edge devices and management solution within the IT layer and the Cloud.

#### 3.1.2. Data standardization and enrichment

Since the sensing layer of nowadays Industry 4.0 applications is typically composed by devices heterogeneous in both nature and data acquire. Such condition poses a challenge for future uses of such data. Due to its proximity with data sources, the OT layer also includes the necessary procedure of data standardization and their improvement via metadata in order to simplify and foster the development of high level applications.

#### 3.1.3. Caching

Data memorization can be implemented within different levels of an Industry 4.0 infrastructure. However, within each level it exhibits different characteristics. Within the OT level, data memorization is performed by limited storage resources that are deployed within the sensing networks. As a result, in this layer the data is mainly cached to then be forwarded toward the IT layer which can then elaborate the sensed information. However, by means of OT caching, it is possible to implement simple yet effective strategy to reduce network resource consumption (i.e. data compression) and also increase data availability, by decoupling the IT layer for edge devices that presents periodic downtimes.

#### 3.1.4. Devices orchestration and application specific features

Due to its close proximity to edge devices, both sensors and actuators, the OT is also encharged to implement specific features tailored to: devices nature, applications requirements, and users requirements. For example, for production machines it can result useful enabling rapid reconfiguration of the sensing resources (more samples within each timeframe, more precision, etcetera) to better adapt to different production stages. Similarly to reconfiguration, the OT layer also implements the diverse process that devices might exhibit for firmware updates.

### 3.2. IT level features

Contrary to the OT, the IT is designed to implement high level functionalities to support the development of new applications, training machine learning models, while also providing effective tools for the management of a whole production plant.

#### 3.2.1. Data ingestion layer

Also the IT layer disposes of a specific set of mechanisms and strategies to gather data acquired via edge devices connected to the production machines. However, within the IT layer such procedure is achieved via a proper ingestion layer which is designed to enable further elaboration on the data gathered. Moreover, such ingestion layer is not targeted only to connect to devices but also to more complex data sources that Industry 4.0 architectures might present, such as OT instances, databases, remote data lakes, and even enterprise and business applications.

#### 3.2.2. Data storing

Contrary to the OT level, the memorization of sensed information within the IT level is not solely designed to optimize network usage and provide a simple data filtering but also to support the application of such information. In particular, within the IT level the data are stored within data lakes which enable long term persistence, fault tolerance strategies, and reliable endpoints for Industry 4.0 applications.

#### 3.2.3. Data processing

While within the OT layer the data might be quasi-elaborated via data standardization and enrichment procedures, the IT layer integrates more sophisticated processes that allow the consumption of such data. More specifically, the data processing within the IT layer includes the further standardization of the data, since different data sources might adopt different formats, and mechanisms that enable data analysis. For example, the IT level implements mechanisms for querying the data lake, implementing data aggregation, and so on. Finally, within the IT layer, the data processing is also supported on both batched and streamed data.

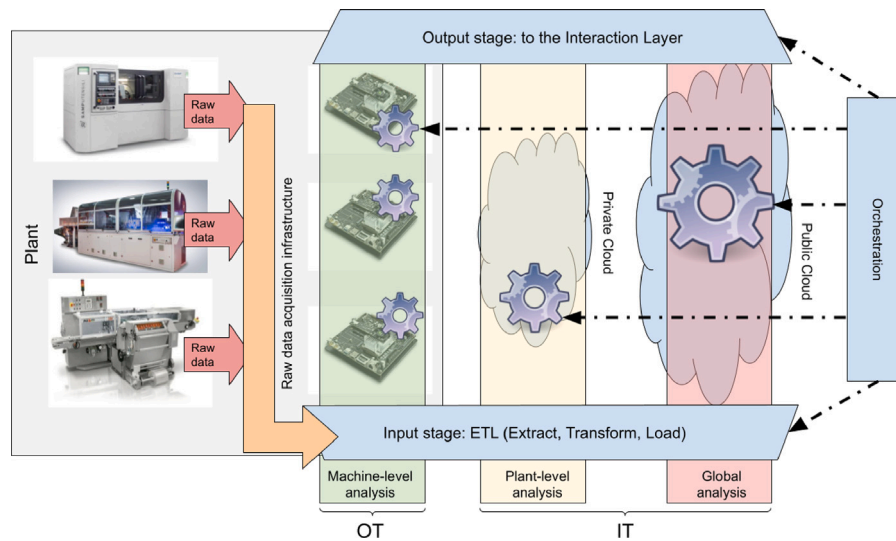


Fig. 1. The IT and OT levels in a smart factory.

### 3.2.4. Data visualization

Another aspect that differentiates the IT layer from the OT layer is the integration of solutions for data visualization. In particular, the IT layer can also include User Interfaces that allow users to directly interact with the data or utilize the data processing mechanisms and APIs disposed by the adopted solution to extrapolate information of interest. Moreover, due to the heterogeneity of the potential data sources, the IT level solutions typically supports diverse views, which can also be customize, that allow end-users, data analyst and even business stakeholder to visualize the data stored within the local data lake with at different level of abstraction.

### 3.3. Combining IT and OT for DevOps and MLOps functionality

The logical partition of the Industry 4.0 architecture in the two layers: IT and OT, enables modern platforms to not only better organizing the process of acquisition, storing, elaboration and utilization of production data, but also define solutions that support DevOps and MLOps within the production plants perimeters.

In fact, combining DevOps and MLOps allows the management of ML-powered applications, which typically consist of a software application component and an ML model that powers the latter. Such applications need three things in order to be supported: *a comprehensive view of the data generated within the manufacturing facility; a sandbox where DevOps products can be safely developed, tested, and polished;* and finally, *a location with enough computing power to train, test, version, and store the ML models.* In this context, the IT layer covers all the requirements for the supervision of multiple data sources, hence multiple OT instances, and integrates all the required function to reorganize, process, and analyze the data received. Moreover, since it is logically separated from the edge nodes, it defines a component in which develop DevOps and MLOps functionalities.

On the other hand, the OT layer, which is instantiated on top of each production machine, is a component with a limited view on the overall production plant but is able to meet all specific requirements of the machine sensing devices and gracefully update parameters, reconfigure devices, or even automatically update software components running on edge nodes. Therefore, the OT is able to actively manipulate all edge node in order to effectively and autonomously deploy and implement the strategies and ML model finalized in the IT layer. As a result, via the IT/OT architectural approach, SMEs can easily integrate Industry 4.0 functionalities while also gradually developing and improving their private IT ecosystems. Another reason for separating the layers is that, when compared to pure IT companies (e.g., social media, e-commerce,

and so on), SMEs cannot produce large amounts of data, so the BIG data collected is insufficient to perform extensive and in-depth analysis, such as using Deep Learning algorithms, which have proven to be very successful in recent years, requiring only big data. To address this issue, we must switch from a model-centric to a data-centric approach. The latter promotes data enrichment, refinement, cleaning, and cleansing. The separation between the two layers is essential, with each layer attempting to enrich the collected data as much as possible using the information and context in which it exists. E.g OT layer adding the context information and the IT layer adding extensive statistics.

## 4. The Bi-Rex Big Data platform

The Bi-Rex Big Data platform emerged from the integration of the outcomes from the Big Data for Manufacturing (BD4M) and Dynamic Edge computing for Plant MONitoring (DEEPMON) projects that were financed by the research and innovation grants, co-funded by the Italian Ministry for Economic Development. In these projects, the universities of Ferrara and Bologna collaborated with a pool of world leading industrial partners, such as Bonfiglioli Riduttori, SACMI, Aetna Group, Philip-Morris International, Poggipollini, etc., to develop Big Data solutions of general applicability and data-driven applications of specific applicability to each Industry 4.0 domain. More specifically, BD4M addressed the requirements at the IT level while DEEPMON focused on OT layer features.

Within those projects, academic and industrial researchers worked closely together by adopting an open innovation approach, which in the manufacturing sector represents a remarkably innovative approach to leveraging collaborative intelligence of various industries. This is in contrast to the close innovation approach, which has always been the way to go in industries, particularly in the Italian one.

The Bi-Rex Big Data platform represents a comprehensive vertical solution capable of addressing the manufacturing industry needs to support business applications such as Product Lifecycle Management (PLM), Customer Relationship Management (CRM), and Enterprise Resource Planning (ERP). In addition, the Bi-Rex platform introduces an innovative solution to enable adaptive analytics at both the IT and OT levels.

### 4.1. BD4M

The BD4M project is designed to overcome the limitations of the main analytic solutions available on the market by creating a general purpose Big Data platform for Industry 4.0 applications. In particular, it

integrates a series of tools and plug-and-play services that support the 4 main operations that Big Data applications exert on data: collection, storage, processing, visualization and analysis.

Data collection is the comprehensive set of features and tools that allows BD4M to gather data from heterogeneous sources such as sensors, remote storage, the Cloud or other platforms, and allocate it within the local storage service. To achieve this BD4M leverages on common messaging solutions designed for Big Data applications. In particular, BD4M implements both Kafka and MQTT brokers, which allows the platform to connect to remote sources and transfer data in a publish-subscribe fashion.

BD4M complete the ingestion process by also integrating specific tools for storage the gathered data. In particular, BD4M disposes of an internal data lake that allows for long-term storage of the data providing a unique endpoint for Big Data applications to acquire data of interest. Since BD4M ingestion process can interact with a highly heterogeneous set of data sources, the platform also support the contextual storage of the collected data in order to optimize the use of storage resources.

The platform further supports the applications by also providing a set of features and capabilities to process and analyze the data collected. More specifically, it enables data manipulation and validation from various types of machines at multiple levels of abstraction (sensor, single machine, production line and finally multi-line or multi-plant) via querying or more sophisticated known algorithms like map reduce. In addition, the platform supports these operations on both structured and unstructured data.

Finally, BD4M also includes end-user decision-making tools, allowing the definition of comprehensive applications directly within the platform. It also offers an internal customizable Web-based dashboard service that allows raw and processed data to be visualized in a human-readable format.

#### 4.2. DEEPMON

DEEPMON is designed to mainly tackle the typical issues of any industrial OT layer. DEEPMON consists of an edge computing architecture deployed close to industrial machines or production line on an edge device at OT level. It is responsible of reading data from field, enriching those data according with a company-provided data model with a specific ETL module, storing and visualizing the processed data on the edge. In addition, it provides an application service that enables the OT-IT communication, this service is a messaging mechanism based on Kafka. DEEPMON has the pivotal role of enabling DevOps and MLOps operations on the edge node.

The logical architecture depicted in Fig. 2 has been implemented using the Siemens Industrial Edge<sup>1</sup> (Siemens IE), more in details, Siemens IE Simatic Edge (nanobox pc) as edge node device, and IE Management deployed on a VM in the private Cloud of Bi-Rex as a IT component. The Simatic Edge is an edge node directly installed on the machine/production line that runs the Siemens Runtime Environment. This system, despite being proprietary, has some points of openness, for example it is based on Docker Engine and it is able to run third-party and self-developed applications. The IE Management is the component at IT layer responsible of handling all the edge devices, deploying field connectors and applications. It also manages all the security and accounting aspects of the solution.

The primary OT level task is the field data connectivity. To read data from the shop floor is not a trivial task since the reading process can be very heterogeneous, because every machine can have its own specific protocol and data model. We developed an OPC-UA connector since it is one of the most used field communication protocol, and we tested it by reading data from a real machine, the DMU 65 Monoblock (DMG

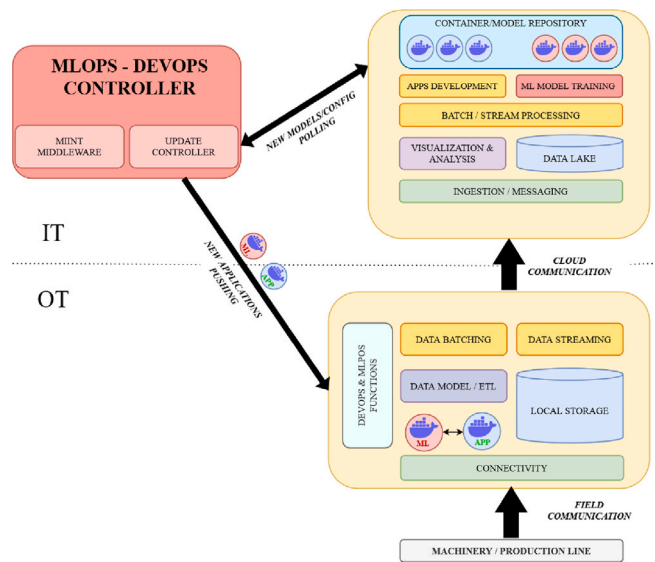


Fig. 2. Logical architecture for cloud-to-edge DevOps and MLOps functions.

Mori<sup>2</sup>), installed in the Bi-Rex tech lab. Once the data has arrived on the edge node, it is delivered to the edge microservices using MQTT. The first service that processes and transforms data is a ETL module. This service is responsible of enriching and integrating the raw data read according a configurable data model. The ETL module is a dockerized application developed by DataRiver s.r.l.<sup>3</sup> called MOMIS (Magnotta et al., 2018). MOMIS is a data integration service that consumes the raw data via MQTT, enriches them with meta information through a configuration file and then sends the transformed data to the Storage Service via HTTP/REST and to the IT layer via Kafka. MOMIS provides a contextual meaning and standardization to the simple raw data coming from the field.

We developed and deployed on the OT edge node a Storage Service to provide degree of consistency and persistence closer to the field. In this way other modules can consume the data received and perform some action such as visualization and execution of machine learning models. This helps in finding errors or failures from machines, avoiding the propagation of errors to the upper layer. Another crucial role of the edge Storage Service is to provide data persistency to tackle edge node connection discontinuity. During periods of disconnection, the Storage Service save and store data so that it can be transmitted to IT when connectivity is restored. The Storage Service is composed by a No-Sql DB and a HTTP front-end written in node that exposes all the database CRUD operations via HTTP/REST. The node server makes the entire systems independent from the specific database and this allow to change its implementation without any consequence on the other services. We adopted MongoDB as proof of concept inside the Storage Service.

Finally, DEEPMON provides an interface that enable the dynamic and run-time deployment of new services or configuration from IT layer directly to the OT node.

#### 5. Automatic deployment of DevOps products and ML models in the Bi-Rex platform

We specifically designed the Bi-Rex platform to support the automatic deployment of new or updated services and ML modules. This

<sup>1</sup> <https://siemens.mindsphere.io/en/industrial-iot/industrial-edge>

<sup>2</sup> <https://it.dmgmori.com/>

<sup>3</sup> <http://www.datariver.it>

functionality is provided by an innovative component, called MLOPS-DEVOPS Controller (MDC), which cooperates with DEEPMON and BD4M. MDC enables the automatic and dynamic service and ML model deployment by including and exploiting the UC and MIINT middleware modules. The latter acts as a bridge with state-of-the-art industrial edge devices, such as those produced by Siemens.

Powell et al. (2022) proposed a framework for further classifying the existing literature on ZDM. The cooperation between BD4M and DEEPMON with the MDC and UC module enables the Bi-Rex solution to be compliant with Zero Defect Manufacturing. Based on this framework, our applied solution focuses on the product, process, and people at the strategic, tactical, and operational levels, both single-stage and multi-stage, using various technologies such as AI, Big Data, and ML/DevOps.

Using the Bi-Rex platform will allow companies like Bonfiglioli Riduttori<sup>4</sup> and SACMI<sup>5</sup> to employ data-driven and machine learning applications to optimize their operations and increase efficiency. Bonfiglioli Riduttori, which manufactures gearboxes, wants to use data to develop and deploy a predictive service to forecast possible faulty gearboxes without the necessity of passing them through a testing machine. This would optimize and speed up the production process, which can be very costly. SACMI, a manufacturer of ceramic production plants, wants to use ML-based applications to optimize the Continuous-line<sup>6</sup> and assist operators in minimizing defects and rejects during the process of finding the optimal configuration of the production line. This would lower operating costs and make the production process more sustainable.

With the help of the Bi-Rex platform, SMEs should be able to quickly and effectively integrate Zero Defect Manufacturing (ZDM) solutions into both new and pre-existing production processes. This cuts down on the time and resources needed to produce high-quality products “first-time-right”. The Bi-Rex solution can keep track of the product from the raw materials all the way through to the finished, marketable product. The MLOPS-DEVOPS Controller can be used to deploy and update current services and ML models, which can then be used to automatically and dynamically adjust the process parameters.

### 5.1. BD4M and DEEPMON composition for an IT/OT architecture

Thanks to their design, BD4M and DEEPMON represent two effective components to define a comprehensive IT/OT architecture for Industry 4.0. BD4M covers all the requirements required by an IT layer since it is able to supervise multiple data sources, hence multiple OT instances, and integrates all the required function to reorganize, process, and analyze the data received. On the other hand, DEEPMON, which is instantiated on top of each production machine, is a component with a limited view on the overall production plant but it is able to meet all specific requirements of the machine sensing devices and gracefully update parameters, reconfigure devices, and even automatically update software components running on edge nodes.

Hence the cooperation between these two platforms defines a comprehensive vertical solution mainly enabled by the BD4M ingestion layer. In particular, thanks to the BD4M ingestion layer, diverse DEEPMON instances, which are related to different machines, are able to forward data to a shared data lake, where both consumers and developers can access, manipulate, and utilize the collected data for both DevOps and machine learning, creating innovative apps increasing the value to the business.

The communication between the two platform is achieved by a message-oriented middleware which decouples various data sources from the consumer, making our general architecture fitting most of the

work machines and the several monitoring and actuation tools present in the different industrial environments. More specifically, we opted for Apache Kafka, an open source messaging broker that enables the communication between multiple producers and consumers. This tool supports the durable retention of messages and permits to handle a huge amount of data. Moreover, since both BD4M and DEEPMON have been designed as primary stand-alone platforms, and thanks to the adoption of a message-oriented communication middleware, the composition between the two can be achieved by a fully distributed architecture, for example, by deploying BD4M in Cloudlets if the computational resources within the production plant are insufficient.

### 5.2. MLOPS-DEVOPS controller and automatic deployment workflow

While the upstream of information is managed by the BD4M ingestion layer, in order to define a comprehensive and vertical solution we developed a component, called MLOPS-DEVOPS Controller (MDC), which realizes the monitoring, (re)configuration, and (re)deployment of software components implementing the adaptive data analytics, leveraging methodologies and tools from the MLOps and DevOps worlds. This design has been fostered by the variety of different environments which characterize SMEs.

In turn, MDC is composed by two main modules, the Update Controller (UC) and the MIINT Middleware. In addition, this service also contains a container registry that provides an internal reliable storage for Docker images and ML models for supporting the UC and MIINT Middleware operations. MDC is the component that enables the automatic and dynamic deployment of services, new or updated, and ML models directly on the DEEPMON edge nodes without the need of human assistance. Indeed, it plays a crucial role of enabling Devops functionalities on the edge devices within the OT layer. It is responsible to trigger the automatic deployment of a new service, to update an already deployed service, or even to automatically deploy a new configuration for the edge running services. All these operations are typically edge platform dependent, and this could result in a big constraint in an even more heterogeneous industrial scenario. To overcome this limitation, MDC includes and exploits the MIINT Middleware modules that exposes Standard API to perform the above mentioned operation in fully platform-agnostic manner (Venanzi et al., 2021). By giving more detail, MIINT middleware is a service oriented architecture (SOA) middleware typically deployed at IT level that has the pivotal role of abstracting the underneath edge layer implementation by exposing a set of REST Standard APIs that allows to automatically and dynamically manage any industrial edge platform. MIINT middleware is composed by microservices and it is exploited to deploy new services and ML models on the DEEPMON edge nodes, as well as to update the current deployed ones. In order to do so, MIINT middleware includes and interacts with a management component provided by SIEMENS; the Industrial Edge Management. MIINT acts as a bridge with the SIEMENS edge solution by interacting with the SIEMENS Edge Management. In the two following sub sections we detail better these two components and how they work.

#### 5.2.1. Update controller

The Update Controller (UC) is a module of the MLOPS-DEVOPS Controller (MDC) component deployed into IT layer, it is responsible for checking the presence of an updated version of a service or Machine Learning (ML) model and deploy it on the DEEPMON edge node. The Update Controller plays a crucial role as a monitoring system to apply MLOps principles in our architecture, more precisely, its job is to continuously keep in check the BD4M container repository for a new version of a ML model or service. These new ML models and/or service updates are generated by BD4M architecture, that receives data from DEEPMON and checks for possible model drifts or performance degradation that can often occur in manufacturing scenarios.

<sup>4</sup> <https://www.bonfiglioli.com>

<sup>5</sup> <https://www.sacmi.com>

<sup>6</sup> <https://www.youtube.com/watch?v=Bxu9fXBzOQ>

When BD4M detects a model drift, that can simply occur by changing the type of the used material, or a drop in model performance, it re-trains the model with the new data and makes the new model available. Once the UC detects this new model or service updated version, it downloads the model/service version and triggers the MIINT Middleware module for deploying the new model/service and updating the DEEPMON behavior on the edge node. In summary, the UC has the job of monitoring the BD4M repository and checking if a new service version or ML model is available for the system, and then it invokes the MIINT Middleware module for updating the edge node ML processes, and services.

5.2.2. Siemens industrial edge adoption in Bi-Rex solution

Due to the widespread adoption of SIEMENS products adopted for IIoT in SMEs, we integrated Industrial Edge platform, which allows to manage both the production lines and the gathered raw data stored within the cloud service with ease. SIEMENS Industrial Edge solution is composed by two main entities, the Industrial Edge Management (IEM) and the Industrial Edge Device (IED), running the SIEMENS Industrial Edge Environment. Usually, in a typical SIEMENS Edge Deployment there is an instance of IEM and one or more instances of IEDS, where IEM has visibility of all IEDs and it is responsible of piloting them, deploying field connectors and application, and handling security and updating activities. While, on the other hand, the IEDs directly interacts with industrial machinery, runs edge services, and manages all the sensing and actuating tasks. The Bi-Rex solution has been deployed on the SIEMENS Edge solution stack, in particular, DEEPMON is runs on a IEDs and it exploits all the functionalities provided by SIEMENS to gather data from the field and for running edge services and ML-powered applications, while MDC includes IEM to manage all the plant's nodes and it triggers the automatic and dynamic deployment of services and ML models.

Typically, both of these two entities are managed, and all the tasks are piloted, by a human operator from a Web Graphical User Interface (GUI) without leaving space to any kind of automation and dynamicity. To tackle this strong constraint, SIEMENS developed a set of REST APIs to automate the piloting of IEM and IEDs and exclusively provided them to us asking for beta testing for the project. These APIs directly and programmatically interact with IEM and IED by enabling the remote and dynamic piloting and deploying of services and configurations. We used them to develop and deploy our vertical Bi-Rex solution, and finally we reported some feedback to SIEMENS.

More in detail, this set of REST APIs enables three categories of functionalities: Authentication, Device Management, and Service/Application management. The Authentication category provides all those set of functionalities to log the user and permit him to perform all other operations, to invoke any other API the user must be authenticated. The Device Management APIs enable an authenticated user to fully control the status and the life-cycle of the edge devices. For example, these functionalities allow to create, delete, and activate the edge devices. Finally, the Service/Application subset of APIs allows to un/deploy services and applications on the devices, to update a service/application to a new version, to reconfigure a running service, and to check the applications/services status.

5.2.3. MLOPS-DEVOPS workflow

To summarize, The BD4M platform receives enriched data from DEEPMON, it analyzes them and checks if there is some ML model drift or performance degradation. If so, BD4M create a new service version or re-trains the ML model of the ML-powered application and publishes it on its repository. The UC, inside the MDC, detects the repository update, downloads the image and triggers the DEVOPS or MLOPS operation on the MIINT middleware. The MIINT middleware module transforms the received request and invokes the remote piloting REST APIs on the IEM. Finally, the deployment API sends the command to IEM to deploy the new service/model on the IED by

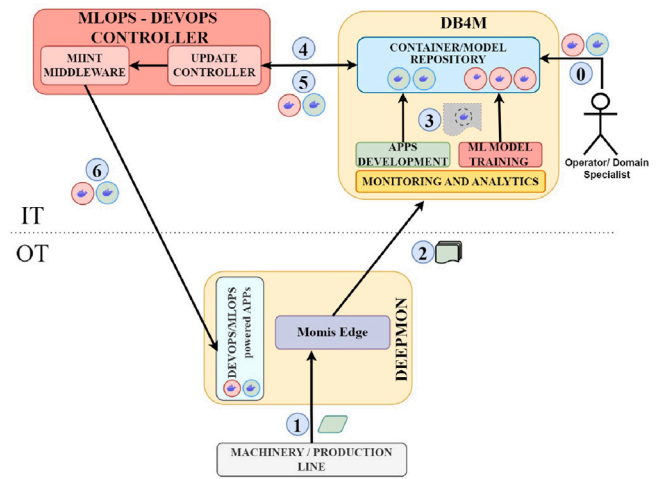


Fig. 3. MLOPS-DEVOPS Workflow.

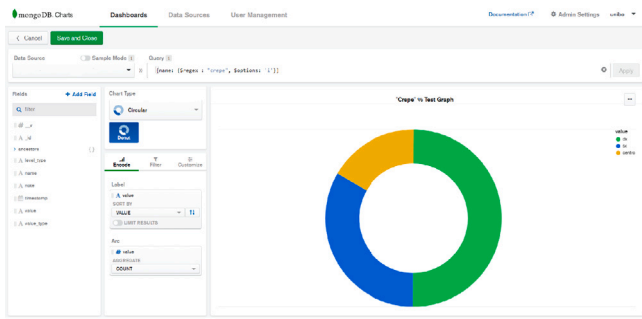
updating DEEPMON processes. In addition to these steps, the shop floor operator and/or the domain specialist can develop a new ML model or Service image and upload it directly into the repository. The whole MLOPS-DEVOPS workflow is depicted in Fig. 3. In particular, the figure highlights that the full workflow has six distinct steps:

1. DEEPMON on the Siemens IED reads data from the field machinery
2. DEEPMON's Momis service enriches the raw data with meta-data of company's datamodel
3. BD4M's services analyze the enriched data received and if it detects a model drift, BD4M re-trains the ML model and saves it on the repository
4. The UC continuously monitors the BD4M repository for a new service version or ML model
5. If UC detects a new service version or ML model, it downloads it and triggers the MIINT middleware
6. MIINT middleware pilots the IEM by invoking the deployment request for a new deployment on the DEEPMON IED

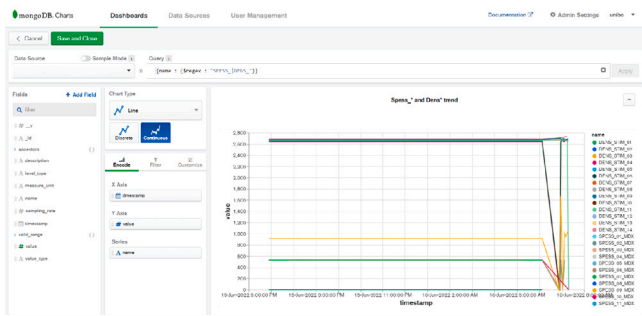
In addition to the above described six steps, we added step 0 to underline where the human factor is involved in the loop. It is worth to be noticed that the operator can trigger the MLOPS-DEVOPS workflow at any time by uploading a new model/image on the repository, or he/she can star-up the whole process by doing it at the beginning. From the point of view of Industry 5.0 pillars, we want to underline that our model helps the operators to have a central role in the factory operation without having to be expert technicians in virtualization technologies. Our MIINT Middleware and UC components can take charge of change detection and automatically expose new updates and features in the deployment (see steps from 3 to 6 in Fig. 3), allowing OT-level workers to be agnostic with respect to knowledge of the tools used internally by our platform. Furthermore, our deployment goes toward a sustainable industry thanks to the aim of ZDM to have a zero-waste production plant. Ultimately, we respect the resilience principle of the I5.0 transition by implementing a deep knowledge of the draft in plant operations that could lead to unplanned downtime, mitigated by the ML continuous analysis of the production plant's working parameters.

6. Bi-Rex solution implementation, experimental tests, and results

In this section, we show the potential of the integrated Bi-Rex solution. Then, we will present our testbed arrangement and the extensive testing of our implementation. We tested the scaling capability of our middleware in order to understand the load peak that the Siemens Industrial Edge platform can manage.



(a) Pie chart query builder

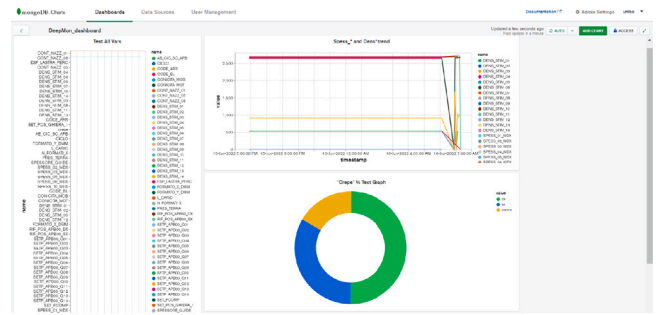


(b) Line chart query builder

Fig. 4. Query builder features.



(a) Two graphs dashboard



(b) Three graphs dashboard



(c) Full dashboard of a plant

Fig. 5. Full output dashboards.

6.1. Bi-Rex platform visualization capabilities

In the next paragraphs, we will show the advanced monitoring functions of the Bi-Rex platform. We reiterate that the platform is capable of managing the variations that occur on the production lines under monitoring in order to correct or improve the behavior of running applications on Siemens devices and on the edge of the network. Our platform is in line with the concept of the closed-loop control system, reiterated by the Industry 4.0 transition, which aims to minimize human interventions and try to bring production lines as close as possible to the definition of zero defect manufacturing.

Fig. 4 shows the great possibilities that the powerful query builder offers to the final users of the Bi-Rex platform. By means of queries written in MongoDB Query Language (MQL),<sup>7</sup> users can selectively create graphs based on the data received by the devices at the edge of the network.

Instead, Fig. 5 shows two examples of complete dashboards built with the query builder. Through the visualization tool used, it is possible to continuously update the graphs in order to detect any data drifts and then prepare a new model or a new configuration to be updated automatically on the devices responsible for gathering and tagging the data coming from work machines.

In the next sections, we will test the operation of the application and configuration update loop based on the needs that arise during the production cycle itself, without the need to carry out cumbersome post-analysis evaluation or even to have unexpected downtime of the production lines.

6.2. Middleware performance tests

We arranged a testbed deployment composed of an Ubuntu virtual machine running our Miint Middleware (Venanzi et al., 2021) that receives the requests for deployment of new services on the Siemens

on premise installation. We used the IaaS OpenStack<sup>8</sup> to provision the virtual machine running the Miint Middleware with the following characteristics:

- Operating System: Ubuntu 20.04 LTS,
- Hard Disk: 150 GB,
- RAM: 16 GB,
- CPU: 4vCPU.

The Miint Middleware component is the access point for all the clients that want to interact with the edge ecosystem, in our case the Siemens Industrial Edge. It provides a REST API interface that masquerades the direct access to the edge environment. Our deployment allows decoupling of the callers from the particular subsystem that manages the gathering of work machines data at the edge of the network. In this way, the data collection subsystem becomes completely independent

<sup>7</sup> <https://www.mongodb.com/>

<sup>8</sup> <https://www.openstack.org/>



from the caller, and we can in the future replace it with whatever technology we want simply by changing the business logic executed by the controller following the receipt of a particular REST call.

The use case we tested covers the steps depicted in Fig. 3. The deployment we used in the tests includes a physical Siemens edge node on OT level, namely Siemens SIMATIC IPC227E, and a VM running the Miint Middleware on IT level. This is the minimal deployment for running our deployment performance test. The actual deployment of the presented architecture can manage many node devices at OT level, while at IT level, BD4M and MLOPS-DEVOPS Controller are components based on the Service Oriented Architecture (SOA) principle and they can be deployed in a distributed fashion on many VMs or even in the cloud. As stated in Section 6.1, from the data sets stored in the BD4M platform, some drifts may emerge from the normal operation of the working machines at full capacity. Through the graphic display of the trend of the parameters under observation and the execution of ML algorithms that analyze the trend of the data coming from the work machines, the Bi-Rex platform can establish control thresholds and take corrective actions when these are exceeded. In these cases, it is essential that the platform acts in a targeted and fast way to correct the gathering of data or the tagging of the same at the OT layer, where the applications run.

We consider the case of an out-of-scale parameter and the need to update an application that acts as a connector at the OT layer, so that the data can be read at a higher frequency, solving the problem found at the IT level. We want to test the scaling capability of the Miint Middleware in receiving and forwarding simultaneous requests coming from the Update Controller in order to update running applications on the edge of the network. Based on the type of request, our middleware service executes corresponding commands on the Siemens Industrial Edge Management (IEM). In this test, we simulated the requests for deployment of a connector on the Siemens Industrial Edge Device (IED) with updated configurations with respect to the older version. We measured the time between the deployment request and the REST response stating that the IEM took charge of the request. We send asynchronous requests, i.e., we send requests to deploy a new service without being aware of any other requests managed by the Siemens ecosystem at that moment. In this way, we want to simulate a typical asynchronous and sequential request submission model. To send asynchronous REST requests we built a *Collection* with the **Postman** tool<sup>9</sup> and scripted the asynchronous execution of the calls using a Postman CLI, i.e., **Newman**,<sup>10</sup> and a script executed with **node.js**.<sup>11</sup> The *Collection* contains a POST REST request that provides a *compose.yml* file used by Siemens to deploy the corresponding service on the device.

Fig. 6 shows the different trends of the round trip time belonging to the requests sent with different delays from each other. With round trip time we refer to the time needed by the middleware to process the service deployment request and to deliver it to the Siemens Industrial Edge platform, plus the time to deliver back the ACK answer. More in detail, the round trip time is calculated from when the script sends the request to when it receives the answer from the IEM. The round trip time is actually the time required for the IEM to take charge of the execution of the service deployment request. Each series is identified by a different delay between the requests. For each series, we sent 20 asynchronous requests and we repeated the test ten times. So each point in the graphs is the average round trip time of ten tries. The abscissa axis shows the timestamp regarding the sending of each request, while the ordinate axis shows the round trip time

As we can see in Fig. 6, our Miint Middleware is scalable and capable of managing all requests sent to it with the selected delays. We tried lowering the delay from 5 to 0.25 s between consecutive

requests. The response of the Siemens subsystem occurs with good linearity within 30 and 31 s. In Fig. 6, particularly for use cases where the delay is 5 s and 2 s, we find peaks that are configured in the range of 2.5% of the total time (see in Fig. 6 the values at timestamps 30, 35, and 40 in the case of using 5 s of delay, and for timestamps equal to 32, 34, and 36 in the case of using 2 s of delay). Given the general purpose nature of the devices in question, we think that these values of variance are entirely negligible and dependent on the processes that the Siemens devices are performing at that particular time and the requests queue load on the Siemens Industrial Edge Management.

In the face of a low delay between successive requests, we noticed a drop in some requests from Siemens Industrial Edge as it was unable to manage them. More precisely, under a delay of a second between subsequent requests, Siemens Industrial Edge drops some of them. We discovered the dropping of the requests using the packet capture mechanism on the machine running the middleware, by executing the `tcpdump`<sup>12</sup> command. We found that our middleware forwards all the packets to the Siemens IEM, but not all the requests are executed, so we can state that the drop is done from the Siemens counterpart.

In Fig. 7 we show the number of requests performed successfully (y-axis) for each group of 20 requests sent at related delays (x-axis). By increasing the delay between subsequent calls, which take place asynchronously anyway, the success rate increases, settling at 100% of the calls when the delay increases beyond one second.

## 7. Conclusions

To accelerate the adoption of ZDM, with significant advantages in terms of efficiency and sustainability, it is crucial to provide adaptive Big Data solutions designed for industrial applications. In particular, supporting SMEs in the process of revolutionizing their infrastructure represents a key step to enable local economies thriving.

With this goal in mind, the Bi-Rex consortium developed a comprehensive and vertical Big Data solution that addresses the data analytics needs at both the IT and OT layers. Our solution also includes a mechanism for safely and automatically reconfiguring, retraining, and redeploying analytics-oriented software components at the industrial edge, enabling a more agile development-to-production process.

In accordance with the classification proposed by Powell et al. (2022), the Bi-Rex Big Data platform represents a single-stage and multi-stage solution based on Big Data Analytics and Machine Learning technologies that can be applied at the strategic, tactical, or operational level for long, mid, or short term decisions.

Our proposal aligns with Powell et al.'s (2022) vision of ZDM as a multi-faceted approach that considers the product, process, and people at various stages of the value chain, and it provides a possible implementation of the reference architecture for ZDM proposed in Maganini et al. (2020). More specifically, the Bi-Rex Big Data platform aims at advancing ZDM with a particular focus on the zero waste value chain and "first-time-right" concept that follows the principle of feedback loop described in Myklebust (2013) for enabling Dev/MLOps operations and providing machine or production line reconfiguration features. The proposed solution includes DEEPMON, which provides data standardization and integration, as well as storage and visualization services, enabling DevOps and MLOps operations on the node. BD4M, on the other hand, offers IT support, analysis features, and end-user decision-making tools, and is also responsible for triggering Dev/MLOps operations, making the solution scalable. The Bi-Rex Big Data platform aims to overcome the global challenges of adopting ZDM by providing SMEs with an easily configurable, all-in-one solution that can facilitate the adoption of ZDM in their environments and reduce the digital divide.

<sup>9</sup> <https://www.postman.com/>

<sup>10</sup> <https://github.com/postmanlabs/newman>

<sup>11</sup> <https://nodejs.org/>

<sup>12</sup> <https://www.tcpdump.org/>

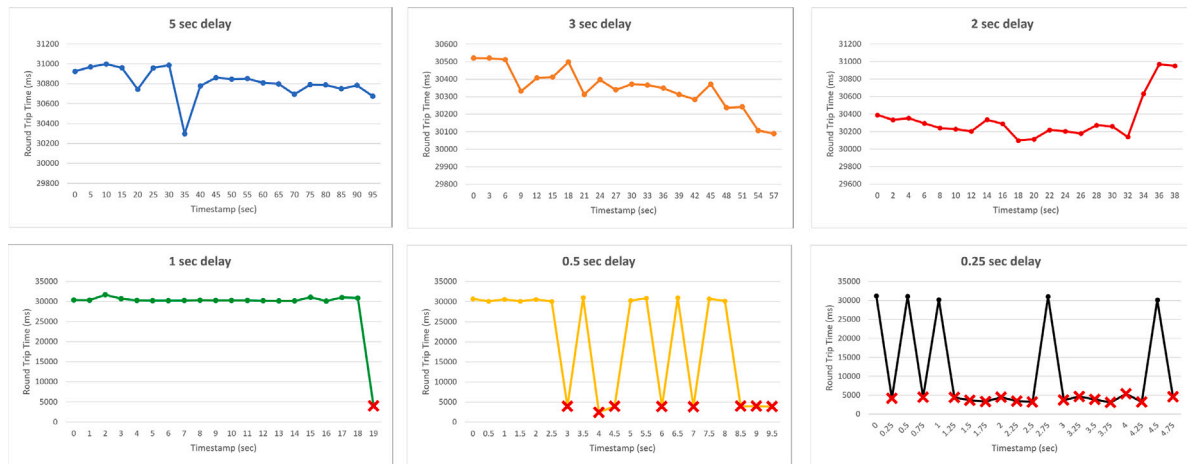


Fig. 6. Trends of Miint Middleware responses for different delays.

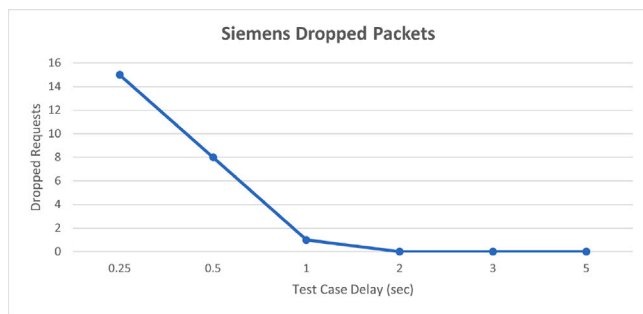


Fig. 7. Dropped packets rate.

**CRedit authorship contribution statement**

**Riccardo Venanzi:** Software, Writing – original draft, Data curation, Writing – review & editing. **Simon Dahdal:** Software, Data curation, Writing – review & editing. **Michele Solimando:** Software, Data curation, Writing – review & editing. **Lorenzo Campioni:** Writing – original draft. **Alberto Cavalucci:** Writing – original draft, Software. **Marco Govoni:** Project administration, Software, Writing – original draft. **Mauro Tortonesi:** Conceptualization, Methodology, Investigation, Supervision, Writing – review & editing. **Luca Foschini:** Conceptualization, Supervision, Methodology. **Loredana Attana:** Conceptualization, Project administration, Supervision, Validation. **Matteo Tellarini:** Conceptualization, Project administration, Supervision, Visualization, Validation. **Cesare Stefanelli:** Conceptualization, Supervision, Resources, Funding acquisition.

**Declaration of competing interest**

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

**Data availability**

Data will be made available on request.

**Acknowledgment**

All authors approved the version of the manuscript to be published.

**References**

Amaral, A., Peças, P., 2021. SMEs and industry 4.0: Two case studies of digitalization for a smoother integration. *Comput. Ind.* 125, 103333.

Bolla, R., Bruschi, R., Burow, K., Davoli, F., Ghrairi, Z., Gouvas, P., Lombardo, C., Frances Pajo, J., Zafeiropoulos, A., 2021. From cloud-native to 5G-ready vertical applications: An industry 4.0 use case. In: 2021 IEEE 22nd International Conference on High Performance Switching and Routing. HPSR, pp. 1–6. <http://dx.doi.org/10.1109/HPSR52026.2021.9481826>.

Boyes, H., Hallaq, B., Cunningham, J., Watson, T., 2018. The industrial internet of things (IIoT): An analysis framework. *Comput. Ind.* 101, 1–12.

Breque, M., De Nul, L., Petridis, A., 2021. Industry 5.0: Towards a Sustainable, Human Centric and Resilient European Industry. European Commission Directorate General for Research and Innovation Publications Office, <http://dx.doi.org/10.2777/308407>, URL: <https://data.europa.eu/doi/10.2777/308407>.

Caccamo, C., Eleftheriadis, R., Magnanini, M.C., Powell, D., Myklebust, O., 2021. A hybrid architecture for the deployment of a data quality management (DQM) system for zero-defect manufacturing in industry 4.0. In: Dolgui, A., Bernard, A., Lemoine, D., von Cieminski, G., Romero, D. (Eds.), *Advances in Production Management Systems. Artificial Intelligence for Sustainable and Resilient Production Systems*. Springer International Publishing, Cham, pp. 71–77.

Caiazzo, B., Di Nardo, M., Murino, T., Petrillo, A., Piccirillo, G., Santini, S., 2022. Towards zero defect manufacturing paradigm: A review of the state-of-the-art methods and open challenges. *Comput. Ind.* 134, 103548.

Corradi, A., Di Modica, G., Foschini, L., Patera, L., Solimando, M., 2021. SIRDAM4.0: a support infrastructure for reliable data acquisition and management in industry 4.0. *IEEE Trans. Emerg. Top. Comput.* 1.

Corradi, A., Foschini, L., Giannelli, C., Lazzarini, R., Stefanelli, C., Tortonesi, M., Virgili, G., 2019. Smart appliances and RAMI 4.0: Management and servitization of ice cream machines. *IEEE Trans. Ind. Inform.* 15 (2), 1007–1016.

Gerrikagoitia, J.K., Unamuno, G., Urkia, E., Serna, A., 2019. Digital manufacturing platforms in the industry 4.0 from private and public perspectives. *Appl. Sci.* 9 (14).

Gokalp, M.O., Kayabay, K., Akyol, M.A., Eren, P.E., Koçyiğit, A., 2016. Big data for industry 4.0: A conceptual framework. In: 2016 International Conference on Computational Science and Computational Intelligence. CSCI, pp. 431–434. <http://dx.doi.org/10.1109/CSCI.2016.0088>.

Han, H., Trimi, S., 2022. Towards a data science platform for improving SME collaboration through industry 4.0 technologies. *Technol. Forecast. Soc. Change* 174, 121242.

Hofmann, E., Rüsçh, M., 2017. Industry 4.0 and the current status as well as future prospects on logistics. *Comput. Ind.* 89, 23–34.

Jaskó, S., Skrop, A., Holczinger, T., Chován, T., Abonyi, J., 2020. Development of manufacturing execution systems in accordance with industry 4.0 requirements: A review of standard- and ontology-based methodologies and tools. *Comput. Ind.* 123, 103300.

Kreuzberger, D., Kühn, N., Hirschl, S., 2022. Machine learning operations (MLOps): Overview, definition, and architecture. <http://dx.doi.org/10.48550/ARXIV.2205.02302>, URL: <https://arxiv.org/abs/2205.02302>.

Magnanini, M.C., Colledani, M., Caputo, D., 2020. Reference architecture for the industrial implementation of zero-defect manufacturing strategies. *Procedia CIRP* 93, 646–651, 53rd CIRP Conference on Manufacturing Systems 2020.

Magnotta, L., Gagliardelli, L., Simonini, G., Orsini, M., Bergamaschi, S., 2018. MOMIS dashboard: a powerful data analytics tool for industry 4.0. In: Proceedings of 25th International Conference on Transdisciplinary Engineering. TE2018, pp. 1074–1081. <http://dx.doi.org/10.1109/WoWMoM.2019.8793033>.

- Masood, T., Sonntag, P., 2020. Industry 4.0: Adoption challenges and benefits for SMEs. *Comput. Ind.* 121, 103261.
- Moeuf, A., Pellerin, R., Lamouri, S., Tamayo-Giraldo, S., Barbaray, R., 2018. The industrial management of SMEs in the era of industry 4.0. *Int. J. Prod. Res.* 56 (3), 1118–1136.
- Myklebust, O., 2013. Zero defect manufacturing: A product and plant oriented lifecycle approach. *Procedia CIRP* 12, 246–251, Eighth CIRP Conference on Intelligent Computation in Manufacturing Engineering.
- Nikishechkin, P.A., Chervonnova, N.Y., Nikich, A.N., 2020. An approach of developing solution for monitoring the status and parameters of technological equipment for the implementation of industry 4.0. *IOP Conf. Ser. Mater. Sci. Eng.* 709 (4), 044065.
- Powell, D., Eleftheriadis, R., Myklebust, O., 2021. Digitally enhanced quality management for zero defect manufacturing. *Procedia CIRP* 104, 1351–1354, 54th CIRP CMS 2021 - Towards Digitalized Manufacturing 4.0.
- Powell, D., Magnanini, M.C., Colledani, M., Myklebust, O., 2022. Advancing zero defect manufacturing: A state-of-the-art perspective and future research directions. *Comput. Ind.* 136, 103596.
- Sahal, R., Breslin, J.G., Ali, M.I., 2020. Big data and stream processing platforms for industry 4.0 requirements mapping for a predictive maintenance use case. *J. Manuf. Syst.* 54, 138–151.
- Venanzi, R., Cavalucci, A., Foschini, L., Bellavista, P., 2021. MIINT: Middleware for IIoT platforms integration. In: 2021 IEEE Global Communications Conference. *GLOBECOM*, pp. 1–6. <http://dx.doi.org/10.1109/GLOBECOM46510.2021.9685653>.
- Villalonga, A., Beruvides, G., Castaño, F., Haber, R.E., 2020. Cloud-based industrial cyber-physical system for data-driven reasoning: A review and use case on an industry 4.0 pilot line. *IEEE Trans. Ind. Inform.* 16 (9), 5975–5984.
- Zheng, P., Wang, H., Sang, Z., Zhong, R.Y., Liu, Y., Liu, C., Mubarak, K., Yu, S., Xu, X., 2018. Smart manufacturing systems for industry 4.0: Conceptual framework, scenarios, and future perspectives. *Front. Mech. Eng.* 13 (2), 137–150.