


Contents lists available at [ScienceDirect](https://www.sciencedirect.com)

Journal of Manufacturing Systems

journal homepage: www.elsevier.com/locate/jmansys

An end-to-end big data analytics platform for IoT-enabled smart factories: A case study of battery module assembly system for electric vehicles

Sinan Kahveci^a, Bugra Alkan^{b,*}, Mus'ab H. Ahmad^c, Bilal Ahmad^a, Robert Harrison^a

^a Warwick Manufacturing Group (WMG), University of Warwick, Coventry CV4 7AL, United Kingdom

^b School of Engineering, London South Bank University, London SE1 0AA, United Kingdom

^c ATS Automation Solutions Ltd., 11 Westleigh Business Park, Winchester Avenue, Leicester LE8 4EZ, United Kingdom

ARTICLE INFO

Keywords:

Big data analytics
Data visualisation
IoT
Industry 4.0
Smart manufacturing

ABSTRACT

Within the concept of factories of the future, big data analytics systems play a critical role in supporting decision-making at various stages across enterprise processes. However, the design and deployment of industry-ready, lightweight, modular, flexible, and cost efficient big data analytics solutions remains one of the main challenges towards the Industry 4.0 enabled digital transformation. This paper presents an end-to-end IoT-based big data analytics platform that consists of five interconnected layers and several components for data acquisition, integration, storage, analytics and visualisation purposes. The platform architecture benefits from state-of-the-art technologies and integrates them in a systematic and interoperable way with clear information flows. The developed platform has been deployed in an electric vehicle battery module assembly automation system designed by the Automation Systems Group at the University of Warwick, the UK. The developed proof-of-concept solution demonstrates how a wide variety of tools and methods can be orchestrated to work together aiming to support decision-making and to improve both process and product qualities in smart manufacturing environments.

1. Introduction

The conceptual basis of Industry 4.0 has risen from the requirement to turn existing production systems into self-aware and self-learning systems to add capacity and control to maintenance planning, anticipate failures, and adjust systems to new requirements and unanticipated changes [1,2]. The paradigm of Industry 4.0 provides the complete integration of computer systems and industrial automation. It allows manufacturing systems to learn from production data through artificial intelligence (AI) and machine learning (ML) algorithms, and enables production efficiency and autonomy to be enhanced while making possible fully customised systems [3–7]. From the standpoint of the Industry 4.0 paradigm, the evolution of smart factories adopts the following core design principles: interoperability and interconnection, information transparency and virtualisation, decentralisation and autonomous decisions, real-time capability, technical assistance and service orientation, and modularity [8–10].

In smart factories, big data analytics plays a vital role in offering improved productivity, product quality, and process safety, as well as

economic and environmental resilience of production systems [11,12]. In big data analytics systems, Internet-of-Things (IoT) facilitates the deployment of smart sensors capturing real-time production data [13], cloud computing allows networked data collection, data handling and off-site analytics [14], and AI/ML facilitates flexibility in decision-making and supports autonomous decisions [15].

Despite its importance, the practical deployment of IoT-enabled industrial big data analytics systems has received little attention in the literature [16]. This paper proposes and demonstrates an end-to-end IoT-enabled big data analytics platform incorporating five layers with several components for data acquisition, data integration, data storage, data analytics, and data visualisation. The proposed solution allows for an integrated system that facilitates decision-making at multiple levels of the management process, and has been implemented in an electric vehicle battery module assembly system deployed by the Automation Systems Group (ASG) in Warwick Manufacturing Group (WMG), the University of Warwick, the UK. This paper contributes to the academic literature and industrial practice in several ways. First, it provides a secure, interoperable, resilient, and scalable reference architecture for

* Corresponding author.

E-mail address: alkanb@lsbu.ac.uk (B. Alkan).

¹ [orcid=0000-0002-5994-4351]

<https://doi.org/10.1016/j.jmsy.2022.03.010>

Received 30 October 2021; Received in revised form 22 February 2022; Accepted 13 March 2022

Available online 31 March 2022

0278-6125/© 2022 The Authors. Published by Elsevier Ltd on behalf of The Society of Manufacturing Engineers. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

IoT-based end-to-end big data analytics platforms in cyber-physical production systems. A particular emphasis is placed on the use of readily available industry standards for interoperability and open-source technologies to design and develop cost efficient enterprise systems for industrial applications as evidenced in the case study. Last but not least, it introduces techniques for multi-tier processing and analytics of streaming machine data to generate actionable insights in real-time.

The rest of the article is structured as follows. Chapter 2 reviews the existing solutions enabling big data analytics in smart factory settings. Chapter 3 presents an end-to-end IoT-enabled big data analytics platform architecture. Chapter 4 represents the implementation of the proposed reference architecture on an electric vehicle battery module assembly automation system as case study. Chapter 5 provides a discussion on the solution and outlines the future work. Finally, Chapter 6 concludes the paper.

2. Existing solutions enabling big data analytics in smart factories

Diverse manufacturing automation and computing services providers offer a variety of solutions for managing big data in smart factory settings. These solutions are becoming more common as big data analytics platforms expand. This section examines existing data acquisition systems, IoT-enabled cloud and edge computing solutions, machine learning-based data analytics tools and visualisation options in smart factory environments.

2.1. Data acquisition

Data acquisition in Industry 4.0 big data analytics systems enables the collection of data from field devices for purposes such as data storage, visualisation, and analytics. OPC Unified Architecture (OPC-UA) and Modbus are two common data communication protocol examples used in the acquisition of data in smart manufacturing settings [17]. These protocols can collect production data in real-time or in batches (or in some instances both). Data communication protocols such as OPC-UA, Modbus, and MTConnect, for example, can allow big data analytics systems to gather production data in real-time, whilst software utilities such as Apache Hadoop can collect data in both real-time and batch-oriented modes [18]. It should be noted that IoT gateways are frequently used to realise data gathering and integration in IoT-enabled big data analytics systems. These gateways provide a variety of important services, including protocol translation, encryption, data processing, management, and filtering, as well as wireless networking of legacy and distant industrial equipment [19]. Furthermore, IoT gateways may be configured to connect with end sensor nodes or I/O devices using protocols such as MQTT, Constrained Application Protocol (CoAP), Hypertext Transfer Protocol (HTTP), and many others. Currently, the majority of shop floor communication protocols are proprietary, which poses a challenge for information technology (IT) / operational technology (OT) integration required to promote data inseparability and connectivity [20]. Therefore, there is a need to develop methods to make use of open communication standards and promote seamless integration between devices and systems.

2.2. IoT-enabled cloud platforms

Sensors are now widely accessible and reasonably priced, allowing industrial equipment, machinery, and other devices to produce massive amounts of data [21]. To make use of these data sources, industrial activities must first be linked to the digital world. In other words, to analyse raw production data, it must be acquired, connected to a network, and stored digitally. In today's manufacturing, GE's Predix, ABB's Ability, Siemens's MindSphere, Schneider Electric's EcoStruxure Platform and Honeywell's Forge are a few examples of commercially

available industrial IoT solutions which can help industries to utilise the big data. Furthermore, cloud service providers like Amazon Web Services (AWS), Microsoft Azure, Google Cloud, Oracle, and IBM's IoT platforms offer capabilities that may be utilised for a variety of industrial applications. The majority of these cloud platforms include distributed computing, big data analytics solutions, tools for data and device management, machine-to-machine (M2M) communication capabilities, and a wide variety of supporting services. In smart manufacturing, field devices can be connected to IoT cloud platforms using several methods, including "plug and play", open communications standards for industrial automation, such as the OPC UA, and publish-subscribe network protocols, for example MQTT [22]. Naturally, cloud-based solutions offer scalable compute power, data storage, and a wide range of services to meet individual requirements. Also, industrial-grade IoT solutions prioritise device safety as well as cyber-security which is a critical component of the industrial cloud computing infrastructures [23]. On the other hand, in most cases, they introduce dependencies on external connectivity, proprietary technologies, limited support for industrial protocols, and custom implementation [24–26].

2.3. Edge computing

Edge computing is a distributed computing paradigm that brings processing and data storage closer to data sources, and it is becoming more popular as a means to perform on-premises real-time data analytics [27]. It allows automation components to be directly linked to the industrial cloud services from the edge, hence providing a better performance than the traditional approaches. The edge interfaces can provide M2M communication, real-time data monitoring, filtering, data analytics, and even machine learning capabilities. Since edge computing is happening in close proximity to data source, on-premises process monitoring can be quickly done through soft sensing models or anomaly detection algorithms and other predictive models. Nowadays, several IoT service providers including AWS, Siemens, and Honeywell started offering edge computing solutions for industrial applications. But existing solutions are still in early days for solving a large fraction of industrial problems [28].

2.4. Machine learning solutions

Today, a plethora of AI/ML tools are accessible both commercially and open source. Within the context of industrial cyber-physical production systems, many of these tools are offered by cloud service providers to the industry as a machine learning as a service (MLaaS). Popular MLaaS choices at the time of writing include Amazon SageMaker, Microsoft Azure ML, Google Cloud AI/ML, and IBM Watson Machine Learning. These technologies incorporate powerful machine learning algorithms that are extremely scalable and capable of data cleansing, outlier removal, automatic feature selection, regression and classification, and many other functionalities. Furthermore, most commercially available MLaaS solutions offer predictive analytics capabilities using deep learning packages/libraries and big data computing with software platforms like Apache Spark, Hadoop, and others. As an alternative to commercial solutions, there are also several freely available machine learning tools, libraries and frameworks, including but not limited to Apache Mahout, Colab, Cortex, Oryx 2, H2O.ai, BigML, Accord.NET, Apache SystemML, TensorFlow, Gradio, PyTorch Lightning, Scikit-learn, Spark MLlib, etc. These tools, like commercial ones, may be used either through cloud services or on-premises, depending on the user preferences. It is possible to access some of these libraries via proprietary machine learning systems. Both Apache Spark MLlib and H2O.ai, for example, are featured in Microsoft Azure HDInsight, a fully-managed big-data cloud service. On the other hand, commercial deep learning software solutions are available from a variety of suppliers, including the IBM PowerAI platform, SkyMind Intelligence Layer (SKIL), NVIDIA DGX-1 Software, Intel Nervana

platform, and others. Additionally, open-source frameworks like H2O.ai, Caffe, Tensorflow, Chainer, Deeplearning4j, Flux, Apache SINGA, MXnet, Keras, BigDL, Theano, and Microsoft Cognitive Toolkit (CNTK) give massive data analytics capabilities via continually updated deep

learning models. Cloud systems such as Predix and MindSphere also provide application services for big data analytics, predictive maintenance, and other tasks in the industrial automation area. For instance, MindSphere offers application programming interfaces (APIs) for

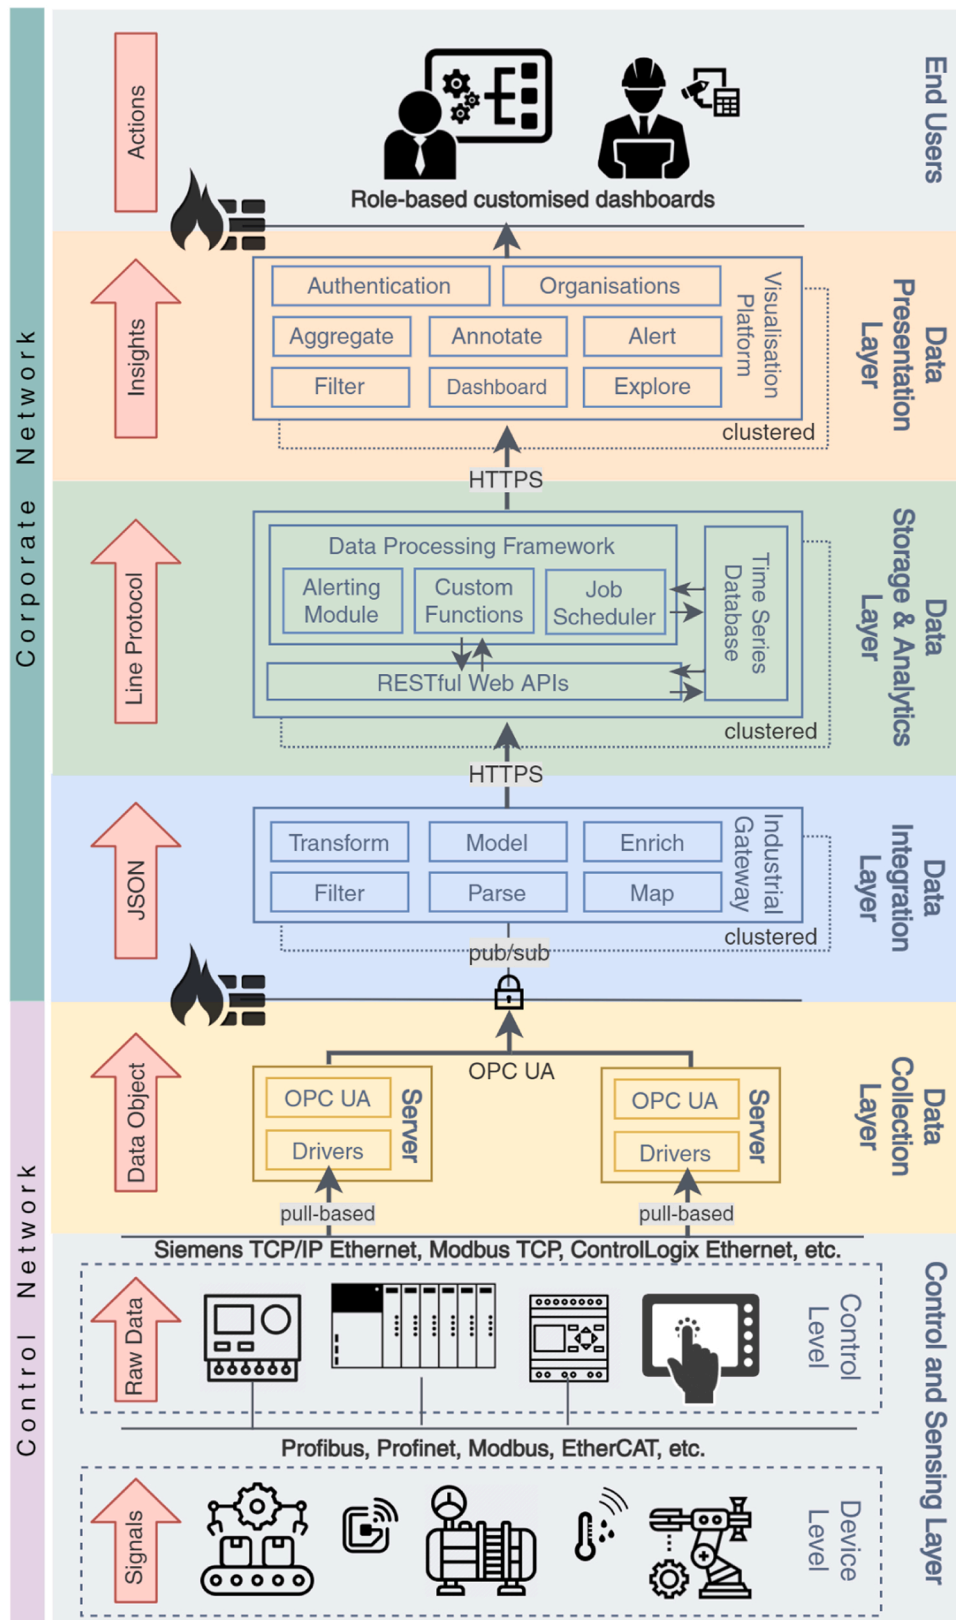


Fig. 1. High-level architecture of the proposed IoT-based end-to-end big data analytics platform.

purposes such as anomaly detection, outlier detection, dimensionality reduction, etc. In addition to pre-build APIs, many clients also prefer to develop problem-specific applications within the cloud service. Despite their vast potential, there are still barriers to using machine learning in smart factories. The transition to big data analytics involves a focus on adopting the appropriate machine learning models and tuning their hyperparameters precisely [29]. Furthermore, the volume and quality of collected data, as well as the efficacy of data pre-processing, are all critical factors in machine learning solutions' performance; hence, big data platforms with structured data models and efficient data management are required [30].

2.5. Data visualisation

There are a large number of data visualisation and analytics software and platforms available both on-premises and on the cloud. Each of them has its advantages and disadvantages. It is inaccurate to assert that one tool is superior than another without considering the specific needs of a certain use case. Increasingly, cloud-native data visualisation and analytics solutions take the advantage of serverless computing and can automatically scale without any infrastructure to manage. Amazon QuickSight by AWS and Power BI on Microsoft Azure are two good examples to create insights from big datasets. Alternatively, Chronograf, Grafana, Kibana, and Splunk can be utilised on-premises, on-demand, or in hybrid environments. For custom dashboards which requires hands-on development, a variety of free and commercial libraries are also available such as D3.js, CanvasJS, and Chart.js [31–33]. It is crucial to make better use of the overwhelming amount of OT level data by improving the accuracy and speed of decision making process. Correct decisions are expected if data presentation and visualisation delivered in a way that capitalises human perception. Moreover in case of poorly designed dashboards, end users could show less interest in the solution despite of its reliability and effectiveness. Therefore, the graphics must be designed carefully in a smart way that attracts the audience. This can be achieved by providing visual functions to end users enabling them displaying the right information in an appropriate format.

3. An IoT-based end-to-end big data analytics platform for Industry 4.0

The proposed solution incorporates a data-oriented approach to create a secure, interoperable, resilient, and scalable reference architecture. In this way, it is aimed at providing a practical platform for businesses from large organisations to small and medium sized enterprises (SMEs) to deliver new experiences, improve processes, products, and services at low cost. As can be seen in Fig. 1, the architecture consists of five major layers: (i) control and sensing layer, (ii) data collection layer, (iii) data integration layer, (iv) data storage and analytics layer, (v) data presentation layer.

3.1. Control and sensing layer

The "things" in IoT refers to a broad range of sensors, actuators, tags, machines, controllers, and devices. A common quality of these is that they are great sources of data that is needed for condition monitoring and performance analysis. This layer involves all kind of sensing and control activities at shop floor that generate raw data within industrial equipment and the surrounding environment. It is considered as the first stage of the value chain. One of the most prominent challenges of this layer is to develop a seamless method for supporting data collection process.

In order to address this challenge, several methods have been devised. First, the signals were transmitted to individual tags without making any processing to calculate even the simplest metrics. This method introduced an advantage in the low level of control system since it does not rely on any code alteration. However, more network traffic

and compute load would have to be addressed in the following higher layers. Also, the details of the control logic had to be understood by other layers to make use of the received tags, such as calculating the cycle time in a higher layer. For the rest of the platform, this introduced a strong dependency on changes of the low level control logic. Nevertheless, this method can still be used when there is no access granted to the control system that is already in its operational phase. Second, an external edge device was added that is only dedicated to collect, pre-process, and perform some basic calculations on the system's raw data. This method helped to eliminate the shortfalls mentioned in the first method. However, it clearly requires extra hardware and additional networking capability on shop floor as an additional cost driver. Last, basic modifications to control code were made when necessary in order to fully utilise the signals that are already available in controllers. For instance, when there was need to calculate basic process metrics such as cycle time, idle time, and fault time of a station, these were handled in the control level since they don't require high calculation power unlike more sophisticated ones such as overall equipment effectiveness (OEE). The third method provided the most optimum performance with a clear isolation of control logic from the rest of the system, and hence, to develop a more flexible and reusable solution at lower cost.

3.2. Data collection layer

An industrial production line represents a sophisticated formation of machines and controllers. For a versatile and customised setup, it may require equipment such as programmable logic controllers (PLC), energy measurement devices from different vendors. In these kind of setups, connectivity requirements are more complicated because of proprietary protocols. Furthermore, the same protocols usually are not capable of meeting advanced data encryption and authentication/authorisation needs due to operating on resource-constraint devices.

In this layer, low level OT protocols are upgraded into a platform independent, firewall-friendly, industry standard communication protocol, the OPC-UA, to address interoperability and security requirements of a solid machine-to-enterprise integration. To achieve this, a multi-node client/server implementation was done with embedded drivers to communicate with the bottom layer. Consequently, it helped to reduce the burden of dealing with a wide variety of industrial protocols as well as allowed to utilise the full potential of the OPC-UA from advanced security features (e.g. authentication using X.509 certificates) to event-driven publish/subscribe communication mechanism. The importance of introducing this layer lies in the ability to optimise the amount of data that is to be ingested while providing data encryption and secure endpoints for inter-layer communication. For instance, the data coming through pull-based data collection pattern every 200 ms is published from the server only when there is a change in its value. Accordingly, subscribers get updates in an object payload along with metadata. Lastly, to achieve high-availability, server instances were deployed with network redundancy. They evenly share the workload in normal working condition. But when a server fails, the system fail-overs to the other active instance automatically.

3.3. Data integration layer

IoT extends existing industrial OT and IT networks. The protocols that are used in a traditional OT environment are likely to be different from modern IT environments. This leads to an integration gap between the two. For example, payloads that are sent from the OT level via OPC-UA comes with a generic object model, its associated type system, and metadata. In the platform, Representational State Transfer - Application Program Interfaces (RESTful APIs) are used to access the data storage layer by taking the advantage of Hypertext Transfer Protocol Secure (HTTPS). To support the convergence of OT and IT systems, communication protocol and message format need to be compatible with the contract of the RESTful APIs, hence the database behind.

A software-gateway prototype has been implemented in this regard for handling two major workloads. The former is protocol-related operations ranging from protocol transformation and validation to exception handling, and the latter is message-related operations including data transformation, mapping, parsing, enrichment, and filtering. Two individual data flows were developed within the gateway. One is for handling PLC data and the other for energy monitoring data. These data flows are control logic agnostic and can be reused for similar systems straightforward. Furthermore, it is always possible to create new data flows on the gateway if required. Consequently, the OT level communication protocol is upgraded to the native web protocol, HTTP/S. Payloads are parsed, filtered unnecessary metadata, enriched where necessary, mapped, and transformed to a common data model. This solution has helped to address a tremendous challenge of establishing such integration in a robust, secure, and smooth manner. Fig. 2 shows an example of how the data coming from the industrial equipment looks like in an OPC UA Server, then how it is transformed to the common data model of the platform represented in JavaScript Object Notation (JSON) format, and finally how it is stored in the time series database as a text-based format, the InfluxDB line protocol.

3.4. Data storage and analytics layer

The data collected from industrial equipment are often latency sensitive, non-relational, semi-structured, streaming a large number of transactions with relatively small payloads. To satisfy these specific requirements, the platform adopts an open source, time series optimised database technology called InfluxDB for storing all industrial metrics and events. It allows flexible data modelling capability, horizontal scalability and high availability with clustering support. In this way, the platform meets heavy read and write requests in a cost effective way beyond the limitations of a single node deployment.

InfluxDB relies on a text-based format that is called "line protocol" to write data points. It represents a data point with a measurement name, a tag key-value set, a field key-value set, and finally a timestamp as shown below.

```
<measurementName>[,< tagkey>=<tagvalue>][, < tagkey>=<tagvalue>]] <fieldkey>=<fieldvalue>[, < fieldkey>=<fieldvalue>] [<timestamp>]
```

The example below shows how an actual data point for energy monitoring data looks like in the database. It has a unique measurement name to differentiate it from other types of data that are collected such as PLCs. The measurement name is followed by a tag key-value set including data about that particular metric, then a single field key-value for the actual value of the metric, and lastly, a timestamp of the metric's occurrence.

```
Example: energy,productionline=IML,area=LegacyLoop,station=Station-01,device=PM5500,mtype=Current,tagname=CurrentA value=0.60 1567770220400
```

Tags are indexed automatically, hence queries on tags are exceptionally fast. Therefore, tagging data points is a great way to be able to run high-performance queries for gathering field key-values. The platform supports this performance optimisation by allowing users tag (i.e., data point enrichment) industrial metrics and events in the data integration layer. Moreover, since leveraging a schema-less database, one can add new tag key-value and field key-value pairs on the fly at any time.

In this layer, alongside the time series database, a data processing framework was implemented to process streaming data in real-time. It allows users to write custom scripts for both pre-processing to perform advanced analytics before shipping the data to the database, and post-processing to run machine learning algorithms for detecting anomalies as well as ETL (extract-transform-load) jobs. User-defined scripts are orchestrated by the job scheduler component. The data processing framework also incorporates an alerting component to trigger actions based on certain conditions defined in custom-written scripts.

All communication with and within this layer is made by secure web APIs, which enables seamless integration between the database, data processing framework, data integration layer, data presentation layer along with any other 3rd party applications and clients. This ensures that there is no direct connectivity to the database. To allow secure communication among all these components, Transport Layer Security (TLS) is enabled and the HTTPS protocol is used by default. Three options are available for authenticating with the database web API: basic authentication, query parameters in the URL or request body, and using JWT tokens. Lastly, authorisation is enforced by leveraging built-in capabilities of the database.

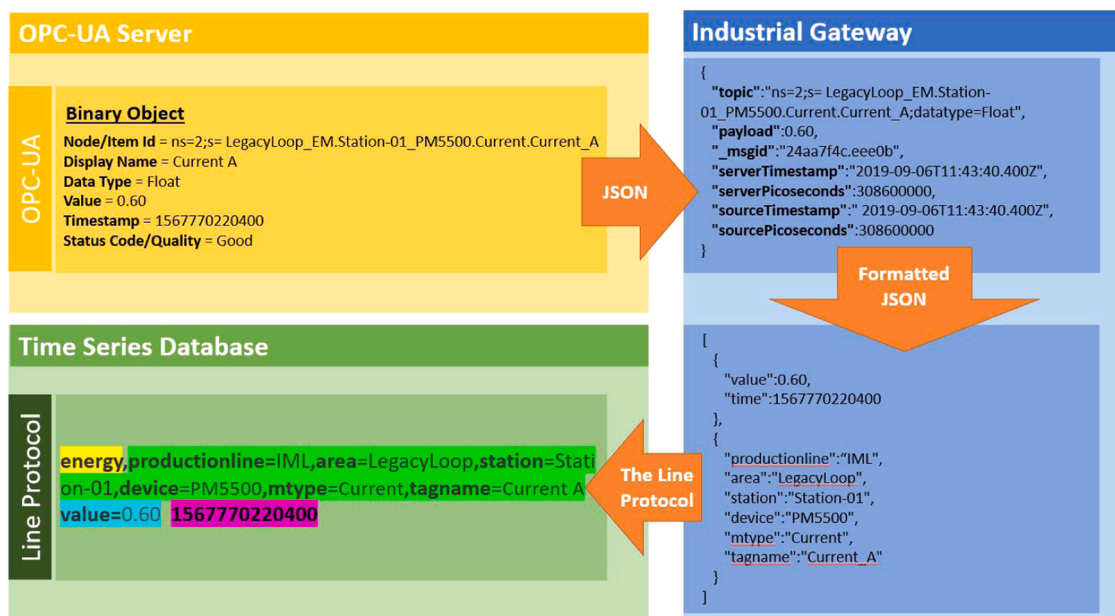


Fig. 2. Representation of data transformation in the data integration layer.

3.5. Data presentation layer

The real value of IoT is in the data. Having said that, only collecting and storing data do not offer much help in creating actionable insights. In order to add value through improved processes and operations, this layer comprises an open source, multi-platform, data visualisation web application.

The platform offers responsive and highly interactive web user interfaces that are supported by a wide variety of visualisation options with a built in user authentication and authorisation system. Apart from default password authentication, the platform allows users to access dashboards through identity federation. Dashboards can be designed and developed, targeting different end-user roles. For example, while technicians are able to monitor and get notifications about low level metrics from machines and abnormal changes in processes through mobile devices on site, back office engineers and managers can monitor higher-level key performance indicators (KPIs), identify patterns, and trends in data. Furthermore, rules for alerts can be defined visually and notifications can be sent to 3rd party systems including Slack, ServiceNow, and PagerDuty.

Consequently, ad-hoc queries were written to explore data. Dynamic and reusable dashboards were developed for most important metrics to address different roles along with alerts based on user-defined thresholds. Results were integrated with systems in the wider organisation in order to help one understand data and make informed decisions.

4. Case study

4.1. Description of the system

A full-scale demonstrator called IML (Integrated Manufacturing and Logistics) was installed at WMG, the University of Warwick to support the research and development activities for emerging manufacturing automation methods, tools and technologies. Some of the research works involving IML can be found in [34–40]. The system has been configured to assemble battery packs for electric vehicles. The pack assembly consists of 6 modules that, in turn, either accommodate 18650 or 26650 form-factor cylindrical cells, Fig. 3. The IML is comprised of eight stations that are divided into five zones as follows: (i) a launch station (ii) a legacy loop that employs a conveyor to move modules through its four stations, (iii) a stand-alone station for pack spot welding, (iv) another stand-alone station for inspecting the spot welding quality, and finally (v) a disassembly station. Automated guided vehicles (AGVs) move products between these five zones according to manufacturing execution system (MES). Fig. 4 below shows the layout and the zones of the IML. And finally, the IoT-based end-to-end big data analytics platform, which has been introduced in the previous section, has been implemented at the IML to carry out this case study.

The IML features state-of-the-art control system and automation equipment from leading vendors, such as Siemens, Bosch Rexroth, Rockwell Automation, ABB, Schneider Electric, Mitsubishi, Festo and




Fig. 3. Computer aided drawings of 18650 form-factor cylindrical cell modules.

SMC. The IML's strong support for multi-vendor manufacturing automation leads to an inevitable challenge for interoperability of communication protocols. For instance, Siemens ET200s PLCs installed in the IML require proprietary S7 communication protocol for accessing PLC data, while Schneider PLCs need Modbus TCP. This applies to all other automation equipment shown in Fig. 4. In the initial setup, the IML does not have any data ingestion pipeline that collects and transforms data real-time to a unified metadata model, stores in a purpose-built database engine, facilitates monitoring and analytics. The entire experience was built on local historians that allow raw data collection locally for a limited period of time. Ad hoc queries were performed on large-scale offline datasets, which limits drastically the potential of generating actionable insights at scale and pace. Therefore, a robust solution is needed in order to collect, store, monitor, and analyse the operating condition of the IML in near realtime for improving production line performance, increasing product quality, and reducing production cost. All this functionality must be secure to prevent malicious attacks, highly available to perform mission-critical operations, and scalable in a way that can accommodate the increasing demand in data collection, processing, and storage without affecting the overall system performance. Lastly, loose coupling, flexibility in deployment, and cost-effectiveness are the major requirements of the desired solution.

4.2. Data collection phase

The KPIs for IML have been categorised into two groups: operational KPIs, and energy KPIs (aka e-KPIs). The raw data required to calculate operational KPIs are gathered from stations' PLCs, excluding the zone 5 that has one manual station for disassembly. These data are: CycleTime, IdleTime, Status, FaultTime, and PartType. Each PLC update the values of CycleTime, PartType and IdleTime at the end of each cycle, whereas the value of Status gets updated at the beginning ($status = 1$, i.e., running) and the end ($status = 0$, i.e. back to idle) of each cycle. In case of any malfunction occurs in the station, $status$ value gets changed to 3, and when the malfunction gets resolved (i.e., $status$ value becomes either 1 or 0) then the value of $FaultTime$ gets updated too. Operational KPIs are calculated with the help of simple modifications made in the control code as depicted in Fig. 5. These KPIs included but not limited to throughput, takt-time, OEE, resource utilisation, changeover time, return-on-assets, resource downtime rate, downtime to operating time, capacity utilisation, first pass yield, maintenance cost per unit, overtime rate, on-time delivery, work-in-progress and scrap rate. On the other hand, the raw data required for e-KPIs come from energy monitoring devices (i.e., PM800 and PM5500 by Schneider Electric). Two PM5500 devices are installed on two individual stations (i.e., station 1 and station 3) out of four in zone 2. Also one PM8000 device is installed for the overall legacy loop as shown previously on Fig. 4. These data are: current, voltage, active power, reactive power, power factor, etc. The raw energy data is then used to measure e-KPIs including: energy efficiency, energy cost per unit, lean energy indicator, energy quality, energy in saturation, etc. in the upper layers of this implementation. Taking the full advantage of the proposed IoT-based platform, the interoperability of communication protocols are provided along with meeting the strict security, scalability, and availability requirements. Fig. 5 shows an example of a custom developed routine that has been deployed in various PLCs.

4.3. Data integration and storage phase

In order to integrate the collected raw data with the time series database, two pre-developed data flows were used within the industrial-gateway prototype. The first one is for the operational data that are gathered from PLCs and the second one is for the energy data that are obtained from energy monitoring devices. Leveraging the flexibility of the proposed platform, a seamless integration was accomplished between OPC UA servers and the time series database by transforming data

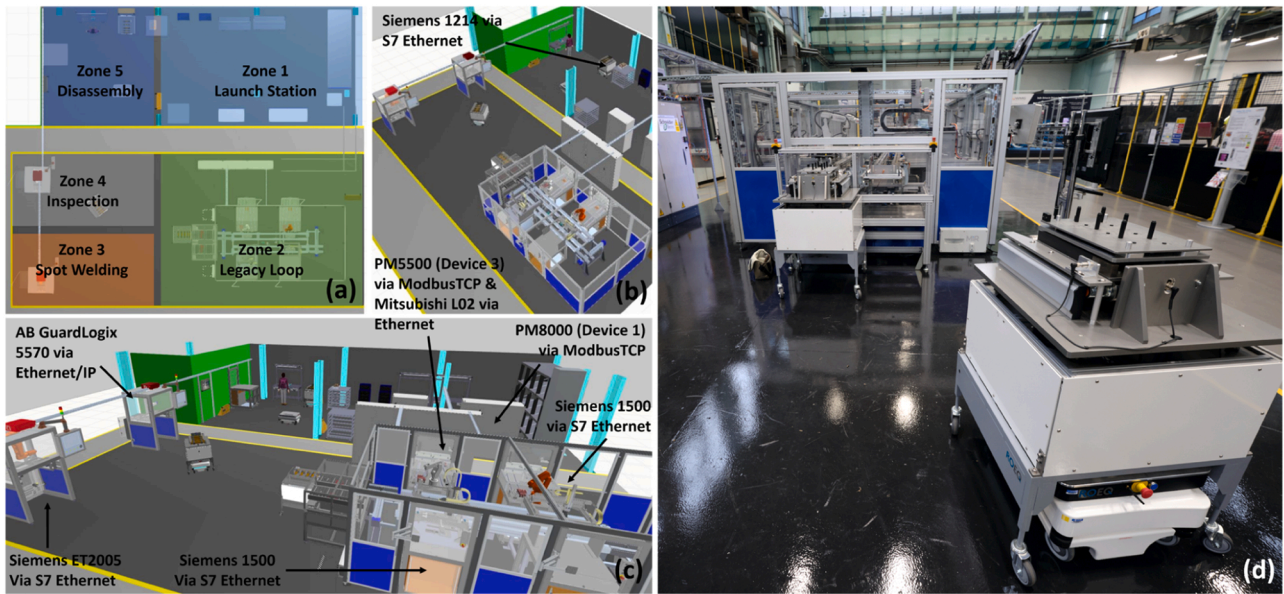


Fig. 4. Integrated Manufacturing and Logistics (IML) demonstrator: a) The IML layout, b and c) 3D models of the IML rig, d) An AGV is carrying battery cells to the legacy loop battery module assembly machine.

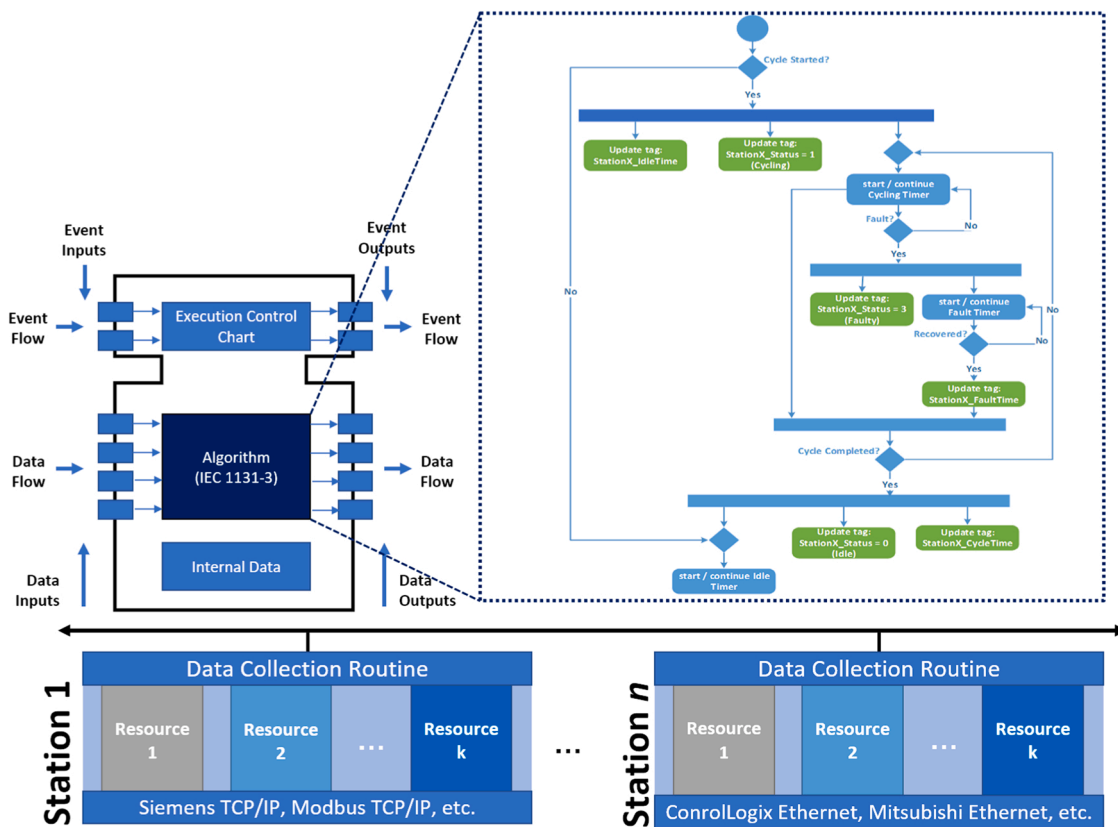


Fig. 5. Example custom developed function block for a data collection routine.

packages on-the-fly from generic and hefty object format to human readable, common data model represented in JSON first, and then to the database's native text-based format. Similarly, the communication protocol was upgraded from OPC UA to HTTP/S, allowing to access the database through highly scalable and secure RESTful web APIs. In total, 67 OPC UA tags (44 tags for energy monitoring data and 23 tags for PLC data) were successfully integrated with 200 ms scan rate for the

monitored items. With a single database node (4 cores vCPU, 32 GB RAM, 1000 IOPS) running on locally attached solid state drives (SSDs), the platform managed to respond up to 250,000 writes and 25 moderate queries per second. Lastly, to create automated actions based on certain thresholds before shipping any data to the database, custom scripts were developed and deployed to the data processing engine of the proposed platform. This helped to accelerate the reaction and response time to

anomalies in the overall system.

4.4. Data analytics and visualisation phase

The dashboards were developed to visualise operational and energy KPIs at three levels, targeting different groups of end-users by providing different details of information. These are: (i) IML level operational KPIs such as current production and takt time, (ii) zone level operational and energy KPIs for zone 2 (i.e., the legacy loop) such as energy cost per part and OEE, (iii) the last level is for showing both operational and energy KPIs for each of the four stations of zone 2, such as cycle time, utilisation and energy cost per shift. Fig. 6 shows the e-KPI dashboards visualising energy data at the workstation level within the IML. These dashboards were created using Grafana by querying and aggregating operational metrics and events that are stored in the time series database. Apart from the operational KPIs, they provide real-time energy information such as energy cost, hourly power and energy consumption, weekly, monthly

and annual energy cost predictions as well as machine level energy factors such as current, voltage, power factor, etc. Moreover, the developed dashboards can be used to visualise predictive analytics capabilities provided by integrated custom Python scripts. System anomalies, e.g., unusual behavioural patterns, outliers in the system response, machine breakdowns, etc., can be detected through the use of machine learning algorithms, and visualised in real-time and can be broadcast as a system alert on dashboards along with IT systems in the wider organisations. Fig. 7 shows an example of visualisation on both shop-floor display screen and a hand-held device.

5. Discussion and future work

The purpose of this work is to provide an industry-ready, secure, scalable, modular, and cost efficient end-to-end solution architecture in order to create a real-time interoperable big data analytics and visualisation platform for IoT-based smart factories. With well-defined



Fig. 6. Station level e-KPI dashboards for first three workstations.



Fig. 7. The legacy loop station level dashboard (on both hand-held device and display screen).

interconnected five layers, the proposed platform defines major components, methods and best practices to help guide the development of innovative solutions for delivering new experiences, improving processes, products, and services. The paper combines theory and experiment as illustrated in the case study, thus clearly shows how device level signals and raw data can be transformed into actionable insights in an industrial environment.

It has been seen that understanding numerous manufacturing-related performance parameters and defining KPIs is not straight-forward and requires a systematic approach for success. Therefore, the solution needs to be flexible to accommodate changes as it goes through the different phases of the implementation. In the control and sensing layer, the authors followed the method of making basic modifications to the control code for fundamental calculations when necessary. Even though it helped to isolate the control logic from the upper levels of the system and provide a more flexible solution with optimum performance, it has been appreciated that the other two methods can also be valid options for use cases which do not allow access to the control logic. Examples of such cases are in pharmaceutical industry, where any changes on the control logic requires tough re-validation process.

The platform addresses interoperability and interconnection challenges of the evolution of smart factories in two layers: data collection and data integration. The former is a central node that serves as a bridge between the control network and the corporate network by removing the complexity of various industrial communication protocols. However, because of its limited ability to meet complex message and protocol-related requirements, the latter was also introduced to support the convergence of OT and IT systems. This also allowed to create a more modular and scalable system.

It has been experimented that defining a common data model for

similar workloads significantly reduces consistency problems in distributed systems. It supports development activities to generate insights in data analytics and visualisation layers as well as increases integration capabilities across different environments. The authors developed and deployed custom scripts to process data both before shipping to the database and after for post-processing. In this way, the platform did not only enable to facilitate much faster responses to events, but also allowed to define behavioural patterns and detect anomalies, leveraging machine learning algorithms. It has been proved that implementing a data processing engine in the platform maximises the value of streaming data and accelerates research in machine learning on the existing industrial data.

Building the presentation layer using open source web based data visualisation technologies significantly reduced development time. By taking the full advantage of responsive and interactive web user interfaces, the solution successfully delivered multi-platform innovative user experiences for different roles with different level of details.

In future work, the following points could be considered: (1) More data points from the shop floor should be introduced to extend the solution to include decision support systems and allow users to take actions and informed decisions to improve KPIs. (2) The current solution was only deployed on an on-premises cluster. Hybrid or cloud-based deployments should be studied. (3) To better maintain the data storage layer, data retention policies need to be studied and data down-sampling techniques should be applied.

6. Conclusion

This paper proposed a reference architecture for end-to-end, secure, interoperable, resilient, and scalable big-data analytics platform. The proposed solution has been implemented in an electric vehicle battery module assembly system as a case study, allowing for an industry-ready, integrated system that facilitates decision-making at multiple levels to improve production line performance, identify bottlenecks, increase product quality and reduce production cost. As a prime example of smart manufacturing systems, specific layers of the proposed platform, including their individual components, were defined in this implementation, as were data models and information flows from data collection through both batch and real-time data analytics and visualisation in order to provide actionable insights. The challenges tackled during the development and deployment of the platform were discussed along with alternative solutions where possible. It has been proved that utilising a well-defined multi-tier platform helps to adopt core design principles of smart factories for Industry 4.0 in cost-effective manner. As future work, it is planned to extend the capabilities of the proposed platform by (i) introducing more data points to include decision support systems and improve KPIs, (ii) considering cloud-based or hybrid deployments, and (iii) applying data retention policies as well as data down-sampling methods for better maintenance of data.

CRedit authorship contribution statement

Sinan Kahveci: Research Methodology (Data Collection, Gateway Design, M2M Communications, Data Modelling and Visualization), Solution Architecture, Software Design and Development, Data curation, Writing – original draft. **Bugra Alkan:** Conceptualization of this study, Research Design, Research Methodology (AI based Data Analytics and Data Visualization), Software, Writing – original draft, Writing and editing of manuscript. **Mus'ab H. Ahmad:** Research Methodology (Data Collection, M2M Communications and Data Visualization), Software Development, Data Creation, Writing – original draft. **Bilal Ahmad:** Funding acquisition, Project Management, Manuscript Review. **Robert Harrison:** Funding acquisition, Project Management.

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.

References

- [1] Rojko A. Industry 4.0 concept: background and overview. *Int J Interact Mob Technol (IJIM)* 2017;11(5):77–90.
- [2] Ji W, Wang L. Big data analytics based fault prediction for shop floor scheduling. *J Manuf Syst* 2017;43:187–94.
- [3] Lasi H, Fettke P, Kemper H-G, Feld T, Hoffmann M. Industry 4.0. *Bus Inf Syst Eng* 2014;6(4):239–42.
- [4] Tao F, Qi Q, Liu A, Kusiak A. Data-driven smart manufacturing. *J Manuf Syst* 2018;48:157–69.
- [5] Alexopoulos K, Sipsas K, Xanthakis E, Makris S, Mourtzis D. An industrial internet of things based platform for context-aware information services in manufacturing. *Int J Comput Integr Manuf* 2018;31(11):1111–23.
- [6] Dhuiieb MA, Laroche F, Bernard A. Context-awareness: a key enabler for ubiquitous access to manufacturing knowledge. *Procedia CIRP* 2016;41:484–9.
- [7] Lenz J, Pelosi V, Taisch M, MacDonald E, Wuest T. Data-driven context awareness of smart products in discrete smart manufacturing systems. *Procedia Manuf* 2020;52:38–43.
- [8] Hermann M, Pentek T, Otto B. Design principles for industrie 4.0 scenarios. 2016 49th Hawaii international conference on system sciences (HICSS). IEEE; 2016. p. 3928–37.
- [9] Sahal R, Breslin JG, Ali MI. Big data and stream processing platforms for industry 4.0 requirements mapping for a predictive maintenance use case. *J Manuf Syst* 2020;54:138–51.
- [10] Wang J, Xu C, Zhang J, Zhong R. Big data analytics for intelligent manufacturing systems: a review. *J Manuf Syst*.
- [11] Nguyen G, Dlugolinsky S, Bobák M, Tran V, García AL, Heredia I, et al. Machine learning and deep learning frameworks and libraries for large-scale data mining: a survey. *Artif Intell Rev* 2019;52(1):77–124.
- [12] Fang P, Yang J, Zheng L, Zhong RY, Jiang Y. Data analytics-enable production visibility for cyber-physical production systems. *J Manuf Syst* 2020;57:242–53.
- [13] Jeschke S, Brecher C, Meisen T, Özdemir D, Eschert T. Industrial internet of things and cyber manufacturing systems. In: *Industrial internet of things*. Springer; 2017. p. 3–19.
- [14] Xu X. From cloud computing to cloud manufacturing. *Robot Comput Integr Manuf* 2012;28(1):75–86.
- [15] Monostori L. Ai and machine learning techniques for managing complexity, changes and uncertainties in manufacturing. *Eng Appl Artif Intell* 2003;16(4):277–91.
- [16] Kabugo JC, Jämsä-Jounela S-L, Schiemann R, Binder C. Industry 4.0 based process data analytics platform: a waste-to-energy plant case study. *Int J Electr Power Energy Syst* 2020;115:105508.
- [17] Mahnke W, Leitner S-H, Damm M. OPC unified architecture. Springer Science & Business Media; 2009.
- [18] Vavilapalli VK, Murthy AC, Douglas C, Agarwal S, Konar M, Evans R, Graves T, Lowe J, Shah H, Seth S, et al. Apache hadoop yarn: yet another resource negotiator. In: *Proceedings of the 4th annual symposium on cloud computing*; 2013. p. 1–16.
- [19] Wang G, Nixon M, Boudreaux M. Toward cloud-assisted industrial iot platform for large-scale continuous condition monitoring. *Proc IEEE* 2019;107(6):1193–205.
- [20] Lopez F, Shao Y, Mao ZM, Moyné J, Barton K, Tilbury D. A software-defined framework for the integrated management of smart manufacturing systems. *Manuf Lett* 2018;15:18–21.
- [21] Boyes H, Hallaq B, Cunningham J, Watson T. The industrial internet of things (iiot): an analysis framework. *Comput Ind* 2018;101:1–12.
- [22] Ye X, Jiang J, Lee C, Kim N, Yu M, Hong SH. Toward the plug-and-produce capability for industry 4.0: an asset administration shell approach. *IEEE Ind Electron Mag* 2020;14(4):146–57.
- [23] Ani UPD, He H, Tiwari A. Review of cybersecurity issues in industrial critical infrastructure: manufacturing in perspective. *J Cyber Secur Technol* 2017;1(1):32–74.
- [24] Ferrari P, Flammini A, Rinaldi S, Sisinni E, Maffei D, Malara M. Impact of quality of service on cloud based industrial iot applications with opc ua. *Electronics* 2018;7(7):109.
- [25] Zhou J, Cao Z, Dong X, Vasilakos AV. Security and privacy for cloud-based iot: challenges. *IEEE Commun Mag* 2017;55(1):26–33.
- [26] Mohiuddin I, Almogren A. Security challenges and strategies for the iot in cloud computing. 2020 11th international conference on information and communication systems (ICICS). IEEE; 2020. p. 367–72.
- [27] Zhou Z, Chen X, Li E, Zeng L, Luo K, Zhang J. Edge intelligence: paving the last mile of artificial intelligence with edge computing. *Proc IEEE* 2019;107(8):1738–62.
- [28] Nain G, Pattanaik K, Sharma G. Towards edge computing in intelligent manufacturing: past, present and future. *J Manuf Syst* 2022;62:588–611.
- [29] Sharp M, Ak R, Hedberg Jr T. A survey of the advancing use and development of machine learning in smart manufacturing. *J Manuf Syst* 2018;48:170–9.
- [30] Lu Y, Xu X, Wang L. Smart manufacturing process and system automation-a critical review of the standards and envisioned scenarios. *J Manuf Syst* 2020;56:312–25.
- [31] Peddoju SK, Upadhyay H. Evaluation of iot data visualization tools and techniques. In: *Data visualization*. Springer; 2020. p. 115–39.
- [32] Protopsaltis A, Sarigiannidis P, Margounakis D, Lytos A. Data visualization in internet of things: tools, methodologies, and challenges. In: *Proceedings of the 15th international conference on availability, reliability and security*; 2020. p. 1–11.
- [33] Chapman C. A complete overview of the best data visualization tools, Dosegljivo. (<https://www.toptal.com/designers/data-3334>).
- [34] Yao F, Alkan B, Ahmad B, Harrison R. Improving just-in-time delivery performance of iot-enabled flexible manufacturing systems with agv based material transportation. *Sensors* 2020;20(21):6333.
- [35] Chinnathai MK, Alkan B, Vera D, Harrison R. Pilot to full-scale production: a battery module assembly case study. *Procedia CIRP* 2018;72:796–801.
- [36] Chinnathai MK, Al-Mowafy Z, Alkan B, Vera D, Harrison R. A framework for pilot line scale-up using digital manufacturing. *Procedia CIRP* 2019;81:962–7.
- [37] Chinnathai MK, Alkan B, Harrison R. A novel data-driven approach to support decision-making during production scale-up of assembly systems. *J Manuf Syst* 2021;59:577–95.
- [38] Ahmad M, Ahmad B, Harrison R, Alkan B, Vera D, Meredith J, Bindel A. A framework for automatically realizing assembly sequence changes in a virtual manufacturing environment. *Procedia CIRP* 2016;50:129–34.
- [39] Chinnathai MK, Alkan B, Harrison R. Convertibility evaluation of automated assembly system designs for high variety production. *Procedia CIRP* 2017;60:74–9.
- [40] Kohr D, Ahmad M, Alkan B, Chinnathai MK, Budde L, Vera DA, et al. Proposing a holistic framework for the assessment and management of manufacturing complexity through data-centric and human-centric approaches. *Complexis* 2018:86–93.