

Edge Intelligence and the Industrial Internet of Things

Richard Hill and Hussain Al-Aqrabi ¹,
*Centre for Industrial Analytics, School of Computing and Engineering
University of Huddersfield, UK.*

Abstract. The Industrial Internet of Things is a key enabler of a new set of capabilities for the manufacturing industry. As more objects become inter-connected, there is a greater need to a) securely share data and b) intelligently manage and exploit the potential of data that is accessible close to its source. *Edge intelligence* is a consequence of the desire to “push” computation and storage to the limits of network infrastructure, in a way that supports intelligent visualisation and interaction with data within the context that it is generated. This chapter examines pertinent topics for potential consumers of IIoT, and by way of a case study in the manufacturing domain, explores the limits and potential of edge intelligence.

Keywords. Intelligence, Industrial Internet of Things (IIoT), analytics, visualisation, M2M

Introduction

Manufacturing has traditionally been a driver of automation to enable greater repeatability, accuracy, and enhanced optimisation of the consumption of resources required to create business value. The recent Industry 4.0 movement indicates that:

“the end-to-end digitisation of all physical assets and integration into digital ecosystems and value chain partners.”[1]

is a means by which new capabilities can be developed through better utilisation of data, and this has re-invigorated interest in automation by demonstrating the possibilities of collaborations between discrete Cyber Physical Systems (CPS) for hitherto undiscovered levels of optimisation. Examples include distributed manufacturing, where tacit production knowledge can be shared across multiple, dispersed sites in real-time; mobile logistics that permit more responsive and adaptive delivery scheduling to enhance product quality, customer satisfaction, and a reduction in operational costs; and healthcare that can both be tailored to the needs of individuals in their home environments, whilst also harvesting trends in treatment and care delivery to optimise and continuously enhance future healthcare services.

This is by no means an exhaustive list of the potential of systems being inter-connected by the internet, and new application domains that can benefit from such connected-ness are appearing continuously[2].

¹Corresponding Author: Richard Hill, University of Huddersfield, UK; E-mail: r.hill@hud.ac.uk.

The focus of this chapter is upon the industrial domain, where Industrial Digital Technologies (IDT) or Industrial IoT (IIoT) is prevalent. IIoT should be considered a subset of IoT; Ashton defines IoT as follows:

“a dynamic global network infrastructure with self-configuring capabilities based on standard and interoperable communication protocols where physical and virtual ‘Things’ have identities, physical attributes and virtual personalities and use intelligent interfaces, and are seamlessly integrated into the information network”[3].

As such IIoT is the specific application of IoT technologies to an industrial setting, which at present is predominantly within the manufacturing industry although this will change as broader acceptance of IIoT takes place[4].

The manufacturing domain is interesting for several reasons. First, there is a mature history of manufacturing being a driver of process automation to enable value-creation. Issues such as the formal design of embedded systems that can both acquire data and also control processes using that data have led to Supervisory Control and Data Acquisition (SCADA) systems that possess resilience and can be efficient to deploy.

Second, manufacturing systems have a particular focus upon each process having a specific, cohesive objective to create value. This means that the measures of quality, repeatability, accuracy, safety and cost are understood, and this assists the evaluation of whether a new technology is suited to the challenge at hand. Wireless Sensor Networks (WSN) have garnered considerable interest as they facilitate the unencumbered connection between systems to create new possibilities[5,6].

Third, manufacturing automation is a receptacle for the tacit knowledge that is required to produce a physical product or service. Traditionally, this has always been to support human operators, although with the increased capabilities of the technology this aspect is being steadily eroded. A key concept here is process intelligence; while systems have merely been automated to remove monotonous or unsafe work from the production staff, the intelligence or learning that takes place during the execution of the process has largely resided with the human operator. As learning systems become more prevalent and accessible, this balance is now being disrupted towards machines retaining, and in some cases developing, new process intelligence as a consequence of their operation.

Finally, the bulk of the development of process automation has largely taken place in the manufacturing industry. As the technology becomes more accessible, there is a democratisation of the knowledge that enables the broader application of CPSs, and this means that experience of the manufacturing domain is now spilling over into allied service industries such as logistics.

As consumers are exposed to the convenience and additional benefits of connected services, this shift in adoption can only accelerate, and such users will inform new approaches to manufacturing through their own interactions with IIoT enabled systems. Thus, the manufacturing industry is an influence upon and shall be influenced by, the customers of products and services that are a consequence the technologies used to create them.

The rest of this chapter is organised as follows. Section 1 shall briefly introduce the emergence of digital transformation, typically referred to as the fourth industrial revolution or Industry 4.0. Section 2 will then consider the demands that industrial organisations place upon IDT and how these challenges are shaping the introduction and development of learning systems for process automation. Section 3 explores how indus-

trial knowledge needs to be secured, and Section 4 describes a particular industrial case where IDT has been adopted and illustrates the potential for further development. Section 5 summarises some future research directions for edge intelligence and Section 6 concludes with some challenges and opportunities for IDT as the manufacturing industry moves towards a greater dependence upon intelligent systems.

1. Digital Transformation

At the time of writing, the digital transformation of business functions is generating considerable interest. Computational hardware is becoming both more capable and less expensive, and this technology exists within an environment where there is often ubiquitous access to communication networks. Many manufacturing organisations have optimised their production processes as far as the process technology has permitted them, and it is only through new technology, with new ways of working, that we can create a production scenario where more value is created.

From a historical perspective, the first major industrial revolution (now referred to as ‘Industry 1.0’) occurred during the 18th Century, and this period describes the introduction of steam power to mechanise production processes in order that new levels of productivity could be achieved.

The introduction of electricity as a power source, when combined with actuators such as electric motors, enabled another transformation during the early 20th Century, whereby the mass production of items became feasible, such as automotive vehicles (Industry 2.0).

This was followed by Industry 3.0 during the 1970s, where microelectronic systems enhanced the control of processes and automation, with some limited use of communication technologies to share data for reporting.

So far, each of these transformations has been labelled in retrospect. During 2011, the German government developed a digital transformation strategy for the industry, which has been described as ‘Industrie 4.0’ (Industry 4.0), or ‘Future Manufacturing’ in the UK. The key thrust that underpins the Industry 4.0 movement is the concept of Cyber Physical Systems (CPS), whereby production plant such as machines are enabled to exchange knowledge between each other so that new efficiencies can be realised. In

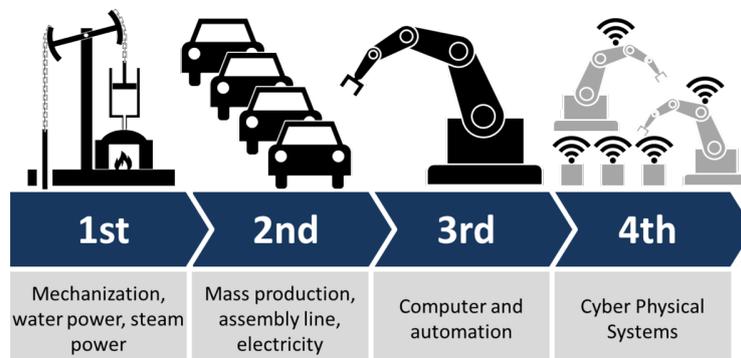


Figure 1. The “industrial revolutions”.

this vision of future manufacturing, the exchange of knowledge does not require human intervention, unlike previous industrial ‘revolutions’. Industry 4.0 advocates talk of fully integrated ‘digital value chains’, where all of the stakeholders can be considered, and have an input into the creation of value through the enhanced optimisation of activities and knowledge exchange [7]. Much of the literature pertaining to Industry 4.0 is based

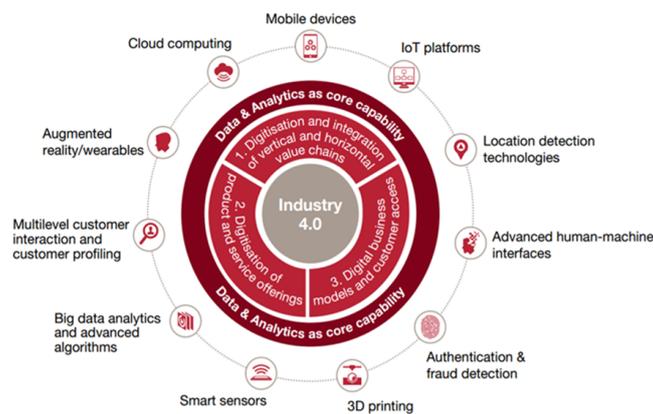


Figure 2. Overview of technologies that are relevant to the realisation of Industry 4.0.

upon a future vision of cooperation and collaboration at all levels [8,9,10,11,12], and it can be challenging for manufacturers to realise such projections when the technologies and approaches are still emerging and maturing. The collection of technologies described in Figure 2 illustrate the importance of data and how its creation and management is key to the potential benefits that are part of the Industry 4.0 vision.

One example of a customer-driven demand for enhanced manufacturing value-chains is that of *mass-customisation*, where consumers have a desire to acquire products that are tailored to their needs, yet are available in the timeframes and at the reduced costs of what were traditionally mass-produced items. This challenge (and opportunity) for the manufacturing industry is a useful context within which manufacturers can examine how they can adopt Industry 4.0 thinking in order to achieve and retain a competitive advantage.

2. Industrial Need

Whilst there is considerable interest in the technology that underpins the Industry 4.0 movement from academic research communities, industry early-adopters and commercial technology vendors, the primary driver for digital transformation is the desire to develop and grow a sustainable business. The potential of IDT to improve efficiencies is leading to a range of developments such as more agile supply chains, as well as the ability to closely-couple different organisational entities that can result in improved services for customers.

As such, it is the needs of industry that should frame the research activity that can provide the most tangible economic impact. In particular, there are three broad areas of research interest that are relevant to the industrial needs from IDT. First, there is a need for the secure exchange of data between smart objects in an increasingly vulnerable problem domain. The rapid introduction of devices that is being witnessed, is resulting in an exponential increase in the number of potential attacks from adversaries.

A second by-product of the accelerated adoption of smart entities through IDT means that there is already a considerable challenge faced by businesses when it comes to managing the outputs of the processes, and the processes themselves. It is a commonplace for manufacturing companies to be faced with a data deluge[13] without the introduction of IDT, such is the volume of data that is generated by production operations. Therefore, if the data can no longer be supervised by human actors, there is a need to automate this management, in order that such activities can be delegated to the systems themselves. Such systems require various degrees of intelligence to facilitate automation. However, in order to process such massive amounts of data, IIoT systems generally require cloud computing services which often experience large round-trip delays and poor Quality of Service (QoS) as a large amount of data needs to be transferred to centralized data-centres for computation[14].

Finally, there is a need to up-skill existing manufacturing personnel so that they can fully exploit the value of the data that is being collected, harvested and processed for new insight.

The need for data-literacy among all manufacturing personnel is becoming more important as digital transformation affects the demands that are placed upon employment roles. There is a foreseeable shift towards *knowledge workers* over traditional unskilled, manual labour, and this can only continue to increase as more IDT is embedded within the fabric of industrial production and its associated supporting services. For the purposes of this chapter, we shall focus on the first two categories of industrial need: security and intelligence. The up-skilling and development of staff, though essential to the successful adoption of IDT is thus beyond the scope of this discussion.

3. Securing Knowledge

Computer security is a fundamental issue for all industrial information systems, and at its most basic it is concerned with the processes and systems' ability to ensure that business process data is protected from leakage, loss and mis-use. However, the emergence of the possibilities of Industry 4.0 is tempered somewhat by a general reluctance from business organisations, particularly small and medium sized enterprises (SME), to engage with any technology that relies upon the use of networks, computation and storage facilities that transport data outside of the physical realm of the organisation.

Whilst there has been considerable investment in the development of secure communications and particularly, secure distributed systems, this does not allay the fears of SME businesses.

These fears are related to the potential introduction of a security vulnerability into a system that contains valuable intellectual property (IP). IP is central to the profitability and sustainability of SME businesses, and IP is generally contained within the combination of process equipment and the capabilities and experience for the human operators.

In the case of manufacturing plant that is not monitored by IDT systems such as sensors, computational and networking devices, the opportunities are mostly limited to nefarious human agents who must choose to release data outside of the business premises. Operations knowledge is valuable, though the value is diminished when real-time data from a process is absent.

It follows therefore, that the case for adopting IDT is challenging for SMEs to consider, as while there are the potential benefits of installing IDT, there is also the introduction of vulnerabilities that did not exist before.

A typical installation of IDT includes sensing, network connections (wired or wireless) and an embedded device to enable data to be captured and processed from sensors, before being sent via the network to a cloud based system that may reside off the premises.

The technological response to the forecasted increase in data generation and transport from IoT devices is to develop architectures whereby the computation is ‘pushed’ to the edge of the network, giving rise to *Edge* or *Fog* computing[15]. Such schemes do increase the potential for data to be processed local to a business process, but any significant analysis still relies on the use of encrypted communications to cloud-based systems in order to complete the system. Edge computing thus presents an opportunity to securely benefit from the introduction of IIoT technology, providing the computation and visualisation can remain within the confines of an organisation’s firewall.

Figure 3 illustrates how an edge computing architecture has been deployed for an industrial business that employs automated product inspection; further details are explored later in this chapter.

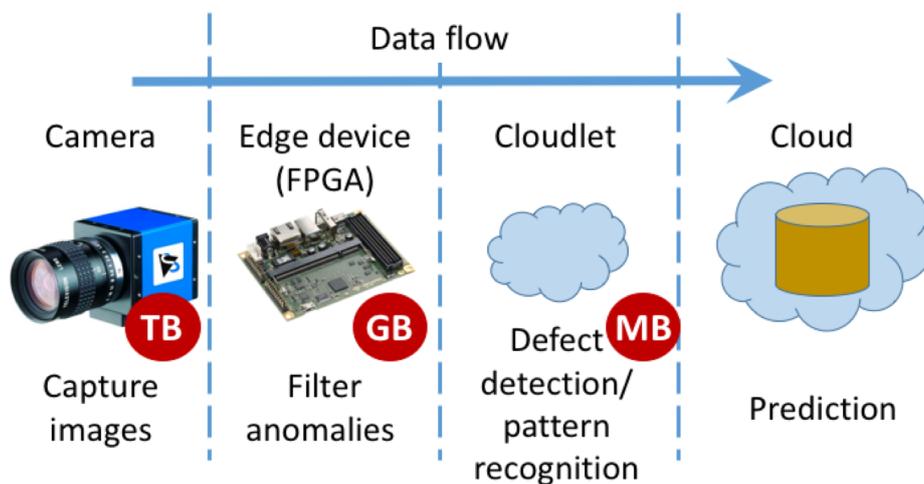


Figure 3. Tiers within an edge architecture for secure video stream data analytics[16]. Video data in the order of Terabytes (TB) is reduced to Gigabytes (GB) after object detection has taken place. This data stream is further reduced to Megabytes (MB) after pattern recognition processing is carried out by *cloudlet* hardware that remains within the physical boundary of the organisation. Analytics processing, or model creation and prediction are performed using rich cloud resources, which may be external to, or internal to the organisation.

Since Edge Computing (EC) is an example of distributed computing, there are some significant benefits for businesses. The architecture serves to reduce the volume of data

that is transported through the network by enabling data processing on edge devices. This reduces the load placed upon centralised cloud services. Such architectures demonstrate significant improvements in the Quality of Service (QoS) experienced by users, particularly with regard to lower network latency [17,18].

The localisation of computation and storage permits more agile analytics and reporting functions to be performed, and it also supports more sophisticated real-time control systems. However, unless there is a specific design feature to enable offline processing, EC architectures tend to assume almost constant network connectivity as it is typical to require at least some cloud resources for computation and storage.

In situations where physical security can be assured, a range of architectures are feasible. For example, if the manufacturing plant network is hardwired, and there is sufficient compute resource within the premises, private cloud instances are suitable, irrespective of whether any edge devices are employed.

In contrast, the situation is different if an organisation uses a remotely hosted cloud for retaining business transaction records. This is becoming more prevalent as businesses seek to reduce maintenance costs by storing general transactional data in remote clouds. A business may assess the risk of having transaction processes being disrupted as a matter of inconvenience, rather than actually ruining the business operation itself. This mindset is not the same when considering the potentially valuable operational level data that is produced by individual processes.

For example, a company may employ automated image recognition as a means of reducing defects in materials. This facility will produce specific datasets that enable the production process to be optimised to a much higher level, and as such might support a lower manufacturing direct cost that enables greater profit margins to be achieved. If a competitor had access to the dataset this might help them achieve a similar level of competitive performance. Conversely, a nefarious agent might seek to penetrate and poison the dataset (usually a machine learning (ML) model) so that the company loses control of its process through incorrect predictions.

In both scenarios, manufacturing organisations are seeking to ensure that the security of such processes is protected and therefore what is suitable for transaction processing and accounting is not necessarily satisfactory for higher value process IP.

In this case, the adoption of IDT should facilitate the monitoring and control of processes such that the IP is not released onto networks nor deposited in remote cloud services. This is a case for localised analytics processing, whereby EC devices perform the necessary computation and control, without requiring more substantial cloud resources. If a process is particularly sensitive, there may be a risk of internal adversaries stealing data from processes. One of the most effective ways of subverting security controls is to steal the credentials of an agent who is authorised to use a system.

Enterprise systems usually employ role-based permissions that enable users to have access to specific functionality according to the needs of their job. However, such systems can effectively provide excessive trust in a role, and in the case of IP protection the paranoia of losing value may well be justifiable.

One approach to avoid this is to utilise localised homomorphic encryption within the edge component. This means that data can be queried from a process, without a) disclosing the contents of the repository to the the agent of the query and b) without informing the edge node what the intent of the query is. This creates a “need to know”

environment where privacy is preserved, without inhibiting the normal functions that have to take place.

A similar situation arises if we consider mobile entities rather than fixed plant. The benefits of IDT become even more compelling if mobile objects can be used as both data sources and also a means to exercise physical control over an environment (as in CPSs). It is therefore likely that communication will be wireless (and thus prone to eavesdropping) and that valuable IP will be generated and shared as part of the normal functioning of the devices.

Again, this is an example where homomorphic encryption can assist the protection of the process. In terms of establishing trust between parties at the outset, the mobile scenario is more challenging than the fixed, enterprise system based approach. Hussain et al. [19] have developed and evaluated multi-party authentication schemes that are not restricted in terms of scalability which is essential for IIoT devices, where we expect to see an exponential increase in their adoption.

4. Manufacturing Case Study

We shall now consider the application of IDT, in this case IIoT equipment that is deployed within an edge architecture. A manufacturing company within the UK creates a product that requires visual inspection as part of an overall approach to managing the quality of the finished product.

The inspection process is performed manually by human operators, and while experience and training help the operators cope with high throughput volume, there are still variations in product quality that are not identified prior to being dispatched to a customer. Proprietary solutions exist that use cameras to automatically scan the product and identify any cause for intervention. However, such solutions are a) expensive to purchase at the outset, with an extended payback period, b) optimised to work only with a narrow selection of products, and c) are incapable of learning new variations once trained.

The organisation operates in a very competitive marketplace and is paranoid of breaches in security that disclose specific aspects of their manufacturing processes. In particular, specific knowledge of product defects gives rise to solutions that significantly increase material yields, and this is one example of how a manufacturer, suitably armed with such knowledge, can possess a competitive edge against rivals.

Figure 3 illustrates the physical ‘tiers’ of the edge architecture, and Figure 4 describes the conceptual representation that was used to evaluate the initial simulations for the case study, prior to implementation. It should be noted that the *Edge device* and *Cloudlet* nodes were combined into a single *Edge* node in the final architecture as the selected hardware was sufficiently capable, after the data reduction and pattern recognition techniques had been deployed [16,20,21]. The edge node accepts two discrete types of data input, as illustrated in Figure 4. The first is streamed video data from a digital camera, which constantly produces images at 20 frames per second (fps) of the product as it is produced. The second input is from a sensor network that monitors a variety of physical phenomena such as ambient temperature, humidity and air quality.

Data volume from the latter is limited and can easily be processed by simple embedded systems. In contrast, streamed video data arrives in large quantities, especially since there is a drive to increase the image resolution so that more defects can be de-

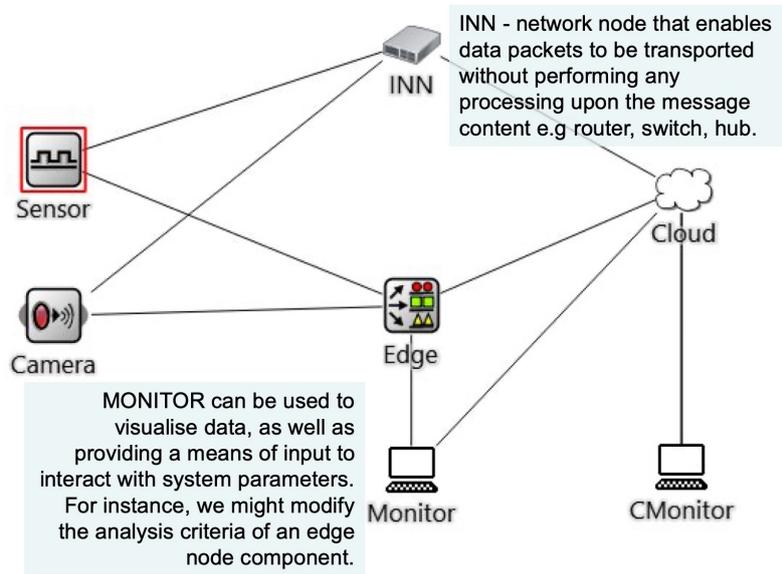


Figure 4. Logical arrangement of the components within an edge-inspired architecture for production defect analytics[16]. The architecture enables more localised analytics to occur near to where the data is being generated (*Monitor*) as well as the more traditional visualisation functions upon the data after processing (*CMonitor*) in a Cloud. The *INN* node permits comparisons in performance to be made and subsequent evaluation, through simulation, the results of which are described later.

tected during processing, rather than by the customer. Initial data reduction is achieved by selecting only a sample of the video frames for analysis, with the remainder being discarded by the edge node.

Video data places considerable strain upon IT networks, and it is a constant challenge to be able to minimise the adverse effects of such data upon the overall performance of a network. With the use of modern embedded systems hardware such as Field Programmable Gate Arrays (FPGA), it is now feasible to include more advanced data processing functions within network nodes, thus reducing the amount of data that requires transportation. Consequently, this removes a considerable potential load from the network infrastructure. Figure 5 illustrates the established approach to data capture and processing for the sensor inputs by way of the red path. Data from the sensors is transported via network infrastructure to a centralised cloud service for subsequent processing.

Alternatively, the blue path shows how data is processed and transported in an edge architecture. After data capture, the edge node performs analytics processing which enables redundant data to be discarded and only data that is of value to be transported any further. This means that a) only useful data is retained, and b) network load is reduced.

Figure 6 describes both arrangements with respect to the processing of video data. One practical advantage of an edge architecture is the ability to provide localised visualisation close to the process that generates the data. Recent interest and advances in the use of machine learning (ML) for performance improvement has led to increased dependence upon the harvesting of clean data for predictive model construction. Such data is usually obtained from a central repository (such as a cloud). This is impractical for large

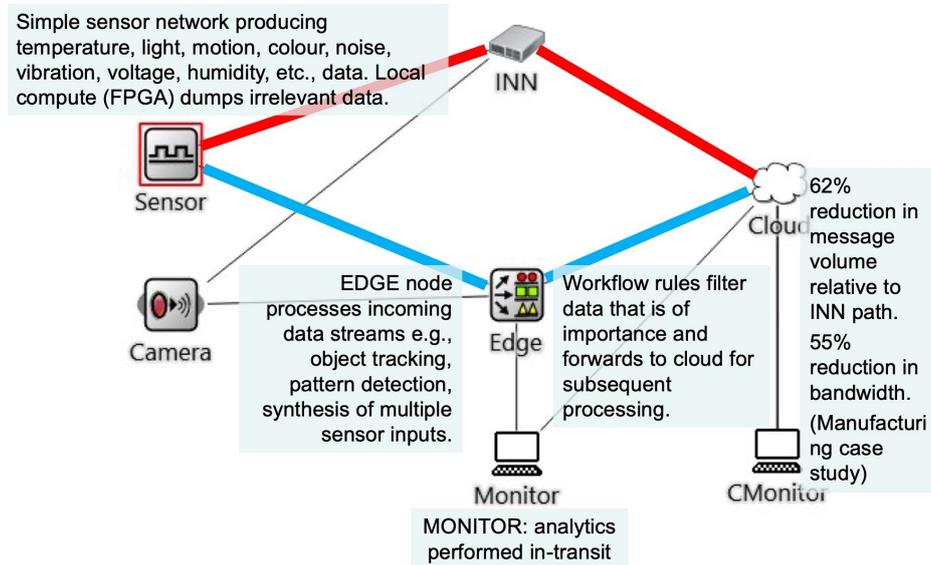


Figure 5. Comparative evaluation of sensor network data being processed by a) IT-centric cloud architecture, and b) edge computing architecture[16].

quantities of high-velocity data (for example video), and also if an organisation does not want to transport the data to a remote host. If the data (and to some extent, insight) is accessible for visualisation at the edge network node, an operator can assist and accelerate the training of an ML prediction model by providing verification of the data that is being presented.

An additional, significant benefit from this arrangement is also the intrinsic recording of tacit knowledge about a manufacturing process, which is typically acquired and retained by human operators. The capture and curation of such knowledge assist repeatability at the local level. However, once the data is encoded and retained within models and data repositories, it can then be shared amongst other processes, in other locations.

4.1. Simulation results

Both of the scenarios were simulated using Omnet++ v5.1 (<https://omnetpp.org>) using good principles[22] and a summary of the results are described in Figure 7. As per Figure 4, the simulation comprises three tiers: a data source tier that in this case includes a sensor and a video camera; an edge data processing tier that combines filtration and pattern recognition functions to process the data; and a centralised cloud tier that is used to perform data intensive tasks such as predictions, etc. The simulation contains the following elements:

- Sensor - process and environmental parameters are detected via sensors to provide a data source for the system. For the manufacturing case study this included simulated ambient temperature and humidity data;
- Camera - a digital video camera that provides a source of streamed video data for analysis;

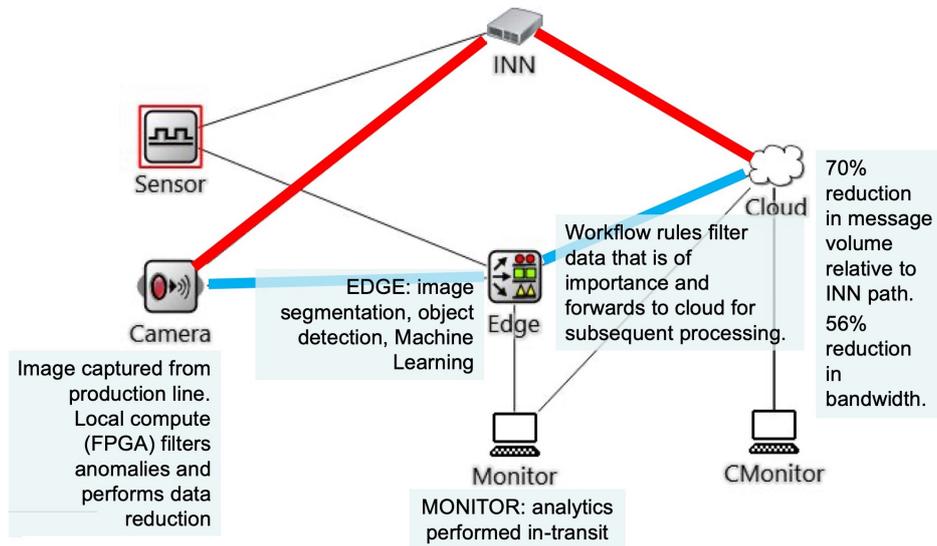


Figure 6. Different routing paths illustrating the relative routes of data transport and processing for conventional IT network (red path) as opposed to an edge architecture (blue path)[16].

- Intermediate Network Node (INN) - this is a traditional component in a network that is typically a router or switch;
- Edge - this device incorporates data processing, storage and networking functions to enable intelligence to be embedded into the node[15];
- Monitor - as referred to earlier, this device enables data visualisation to take place both local to the process as well as more conventionally by reviewing data that is stored in a central repository. In addition, this also provides an opportunity to interact with the data, such as for improving the accuracy of prediction models contained within the respective IIoT or network component.

Edge devices such as FPGAs can be added to existing SCADA hardware to augment their functionality in order to become IIoT entities. As the capability of hardware increases, there is an associated improvement in the ability of IIoT components to perform functions such as model training and outcome prediction. Most industrial manufacturing scenarios involve simple sensors which continuously produce data at regular intervals and IIoT devices which are attached to a system where a human operator monitors the feedback from sensors/IIoT to determine the safety of the system and to perform defect analytics. We investigate the use of in-transit analytics by considering two use cases for edge nodes.

In the first case, raw data from a sensor network is processed, albeit with constrained data storage capacity. For the second case, larger streams of video data are considered, where there is a time constraint for the processing of the data. It is important to understand that while it may be possible to perform image analysis with constrained hardware, the result of the analysis needs to be timely as the usefulness of any analysis may have diminished. Data such as temperature and humidity is straightforward to process in real-time and places minimal demands upon modern microcontrollers. The de-

mands on processing increase markedly when considering video data. Video data must be completed quickly to support timely visualisation, which is central to the provision of feedback[23,24].

Figure 7 shows a summary of the results of the simulations. When data processing occurs at the edge component, a 62% reduction in data transported to the cloud for subsequent analysis is observed. This is in contrast to a more conventional routing of data via the *INN* node to a cloud destination.

The potential data reduction will be influenced by the data that is collected; a sensor that is queried more frequently, or a collection of sensors, will produce data at a greater rate and will therefore offer more scope for data reduction. This variability will be largely dependent on the application scenario, although developments in sensor technology will also facilitate a reduction in data rates. For instance, ‘digital’ sensors that generate discrete, integer values, will not be as demanding upon a network connection as an analogue sensor producing raw data that requires subsequent signal conditioning.

From Figure 7, there is 55% less bandwidth consumption from the edge arrangement over that which employs an *INN* node, even though the identical volume of data was processed for both cases. In addition, the *Edge* scheme enables data processing to take place within the bandwidth constraints of the system, whereas the *INN* scenario illustrates that the bandwidth consumption is approaching the maximum available. For the video data,

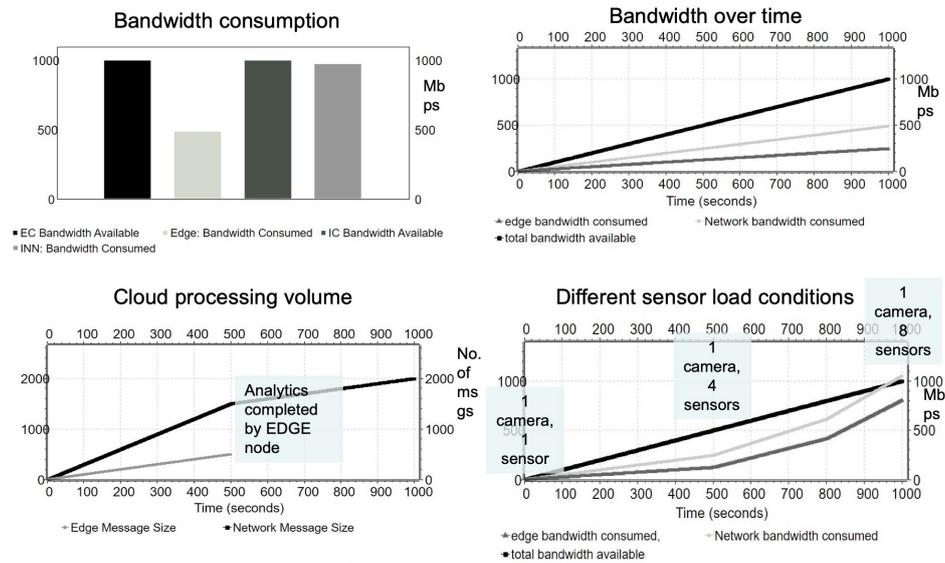


Figure 7. Summary of results from simulation runs. An edge-inspired network topology was compared with a conventional data-transport focused arrangement to identify the key performance characteristics[16]. In all cases, the *Edge* approach reduced demands upon the network infrastructure, while also providing the facility to visualise process-relevant data in the appropriate context.

time constraints are used to determine whether the stream is processed on the edge node, or it is transported to the cloud for offline processing. The simulations demonstrated that an edge based architecture had completed its data processing by $t = 500$ secs. Around 600 messages were processed by the *Edge* node as opposed to the 2000 messages that were handled by the *INN* node. Consequently, the *Cloud* node is still processing mes-

sages at $t = 1000$ secs, suggesting that an edge approach can yield a 70% increase in productivity.

In terms of total bandwidth consumption during the simulation, approximately 1GB is used between the *Edge* to *Cloud* nodes. This holds for one IIoT device (video camera) and one sensor and illustrates a 56% reduction in bandwidth consumption when compared to processing via *INN* and *Cloud* nodes.

4.2. Scaling Intelligence

Having established the potential of an edge-based approach to using localised intelligence to process sensor data, the next experiment investigated the implications of increasing the quantity data sources. Figure 7 shows the relative bandwidth consumption under different sensor load conditions. Three additional sensor nodes are added at $t = 500$ secs, and an additional four sensors (making a total of eight sensors and a camera) are added at $t = 800$ secs.

During the first 500 secs of the simulation, the edge bandwidth increases linearly but remains lower than that required by the *INN* node. Upon the addition of three sensor nodes, the bandwidth required by the *Edge* node increases at a steeper rate, albeit still less than that required by the *INN* arrangement.

Since the demands upon available network bandwidth are less using *Edge* rather than *INN*, we can observe from the simulation that it is not until $t = 900$ secs that the available network bandwidth is exceeded by the *INN* node. Bandwidth consumption of the *Edge* node is still increasing but remains within the network capacity. Given that the data sources produce an identical quantity of data, the use of an *Edge* component significantly reduces the strain placed upon the network by 54%.

5. Future research directions

As with any field that is developing rapidly, growth in performance is underpinned by architectures that are flexible and open to change. Whilst many manufacturing plant items are controlled by proprietary systems, the ITA architecture can help the inclusion of legacy systems whilst Industry 4.0 becomes embedded. An immediate challenge is an extent to which manufacturing plant suppliers can embrace open data communication protocols.

At present, the communication demands are relatively low-level, in terms of transporting data around a network. However, as M2M communication evolves towards *intelligence*, it is foreseeable that goal-based communication will be required as we imbue the capability for machines to negotiate and collaborate with each other autonomously. Another aspect of open architecture development is that of domain specificity; what influence does a sector of the industry, or a particular value chain have on the design of its data architecture for Industry 4.0?

5.1. Algorithms

As we employ edge based devices to perform more computation, there will be an increased demand placed upon the development of algorithms that are:

- computationally efficient (across single and multiple processor cores);
- energy aware;
- capable of reasoning;
- sufficiently robust and reliable to work with Operational Technology (OT);
- self-validating;
- adaptive.

Machine Learning and *Deep Learning* are of interest to the research community at present, particularly as one of the approaches to dealing with the impending data lakes that are generated by IIoT equipment. However, as we move to data stores that hold graphs as knowledge abstractions, whilst at the same time we increase the number of computational devices that are interconnected, we are moving away from the traditional scenario of learning algorithms being executed in a central location, upon a consolidated database.

One challenge to explore is how deep learning algorithms can be distributed across a number of possibly heterogeneous data stores. An associated challenge is that of dealing with high-velocity, streaming data. Often the ‘window’ of data is constrained by the timeliness of the visualisation, and therefore predictions are required from small amounts of data, rather than the larger quantities that suit ML techniques much better.

5.2. Design Abstraction

IIoT systems that include intelligence are inherently complex to design for resilient behaviour, particularly where there is physical actuation in a system such as with a CPS. One approach to managing such complexity is to use *agent* design abstractions to focus, whereby the emphasis of a system is expressed in terms of goals, actors and behaviours. Hill et al[25,26] explores this in the domain of agent-managed healthcare delivery in the community, leading to the use of Agent UML[27]. This has been developed to support Microservices Architectures[28], and is a promising area of research, especially when combined with automated model-checking for additional robustness.

5.3. IIoT management

IIoT implementation is already identifying not only a need to deal with large amounts of data, that is continuously streamed, but also there is a massive potential increase in the number of tasks and workflows that require governance. We are heading towards *super-complexity*, where the management and coordination of activities become more convoluted as more devices become network-enabled. In the context of this chapter, we have discussed and presented localised analytics as a prerequisite for effective visualisation.

5.4. Network infrastructure

The proliferation of IIoT devices that are networked wirelessly means that IIoT objects and networks will become increasingly fluid in their organisation. This will re-emphasise the need to create network designs that are both flexible and resilient. Software Defined Networking (SDN) is one approach that we can see greater use of as IIoT (and IoT) implementation increases. A key challenge is how SDN can be informed by the multiplicity of potential data sources in order that a network is correctly configured for its intended use at that time.

5.5. Visualisation

At present, we are examining the possibilities of improved visualisation, and how we can enhance the analytical processes applied to data to provide a clearer picture of insight for human comprehension[13]. The logical progression of this is for visualisation to be part of Machine-to-Machine (M2M) communication; visualisation would then be a function that would assist automated comprehension *between* IIoT and SCADA devices, the control of which can then be devolved. This would abstract humans from more of the detail to concentrate on higher-order decision making. Enhancing comprehension is not a trivial task by any means, but the ability to delegate some of this to interconnected devices has the potential to be positively disruptive.

5.6. Data Security

Edge computing systems maintain data security using data decentralisation, which means that if an adversary wants to breach the system, they would need to breach each one of a large number of decentralised computing nodes in order to collect the entire system's data. Whilst this assertion holds for the potential security breach of data stores; this does not protect against network sniffing where the streams of data are exposed[29]. Edge based architectures often use homomorphic encryption schemes as well as cryptographic hashing techniques to store confidential data at different distributed locations within the network. Due to the use of these techniques, even if attackers are able to gain access to secure databases, they will not be able to understand the data as it will be in an encrypted format.

This assumes that the hack is not based on imitating the credentials of a user who does have the permissions, in which case the encryption will not be a barrier to data theft. Multiparty authentication schemes[30,31,32], are one approach to addressing this, and we expect to see more work in this area as IIoT devices (particularly mobile devices) are deployed more widely.

6. Conclusions

The manufacturing value-chain is rich with demands, constraints and opportunities to optimise and enhance performance. Digital transformation places an emphasis upon the management of the data that is produced by each process, and the interaction between processes, and this is pertinent for the productive realisation of Industry 4.0.

Such a digital transformation is facilitated by an informed use of data in ways that have not been possible traditionally. Analytics, or the ability to model and predict outcomes using machine intelligence, together with visualisation that is relevant to a specific context, requires devices to be interconnected in ways that enable important knowledge to be shared in real-time.

A practical consideration for IIoT is what has to be done with existing production control and monitoring hardware, that existed before the advent of IIoT technologies. Such equipment presents challenges for the interfacing of SCADA and legacy IT infrastructure in order to realise the free (and secure) exchange of data between machines.

As networks have evolved, there has been a tendency to organise data, and the collection of data, into hierarchical systems that clearly delineate between *Operational Tech-*

nology (OT) and IT system architecture. This has come about due to the need to have control systems (SCADA) that perform a constrained set of activities in a robust way, since the control of plant may cause danger including loss of life. On the other hand, systems that process transaction data are usually more comprehensive and include far more functionality, have far more users interacting with them and as a result do not receive the same attention to their design in the formal sense.

Emerging approaches to software engineering such as *Microservices Architecture* (MA)[33,34] potentially offer a more rigorous route forward, which is essential if software systems are to interact with CPSs such as IDT and in the case of this chapter, the IIoT.

Edge intelligence, when combined with IIoT devices, can serve to augment and enhance pre-existing IT infrastructure that is used by the manufacturing industry. The ability to not only capture, but process, visualise and interact with the data in the context within which it is generated, provides enormous potential for the automation of finer-grained control, and its subsequent optimisation, of production processes.

References

- [1] PWC. Industry 4.0: building the digital enterprise. Technical report, Price Waterhouse Cooper, 2016.
- [2] Gartner. Gartner says 6.4 billion connected ‘things’ will be in use in 2016, up 30 percent from 2015. techreport, Gartner, November 2015.
- [3] K. Ashton. That “internet of things” thing. *RFiD Journal*, 2009.
- [4] P. Friess S. Woelflé H. Sundmaeker, P. Guillemin. Vision and challenges for realising the internet of things. In *Cluster of European Research Projects on the Internet of Things, CERP IoT*, 2010.
- [5] Nik Bessis, Fatos Xhafa, Dora Varvarigou, Richard Hill, and Maozhen Li, editors. *Internet of Things and Inter-Cooperative Computational technologies for Collective Intelligence*. Number 460 in Studies in Computational Intelligence. Springer, 2013.
- [6] J. Wan A. Vasilakos C. Lai S. Wang X. Li, D. Li. A review of industrial wireless networks in the context of industry 4.0. *Wireless Networks*, pages 1–19, 2015.
- [7] T.D. Oesterreich and F. Teuteberg. Understanding the implications of digitisation and automation in the context of industry 4.0: a triangulation approach and elements of a research agenda for the construction industry. *Comput. Ind.*, (83):121–139, 2016.
- [8] M. Rusch E. Hofman. Industry 4.0 and the current status as well as future prospects on logistics. *Comput. Ind.*, (89):23–24, 2017.
- [9] B. Bigliardi F. Galati. Do different npd project characteristics lead to the establishment of different npd networks? a knowledge perspective. *Technol. Anal. Strateg. Manag.*, 10(29):1196–1209, 2017.
- [10] E. Brad S. Brad, M. Murar. Design of smart connected manufacturing resources to enable changeability, reconfigurability and total-cost-of-ownership models in the factory-of-the-future. *Journal of Production Research*, 6(56):2269–2291, 2018.
- [11] S.B. Gershwin. The future of manufacturing systems engineering. *International Journal of Production Research*, pages 1–14, 2017.
- [12] E.D.F.R Loures L.F.P. Ramos Y. Liao, F. Deschamps. International journal of production research. *Past, present and future of Industry 4.0-a systematic literature review and research agenda proposal*, 12(55):3609–3629, 2017.
- [13] J. Wan D. Zhang A. Vasilakos X. Rong F. Chen, P. Deng. Data mining for the internet of things: Literature review and challenges. *International Journal of Distributed Sensor Networks*, 11(8), 2015.
- [14] G.S.S. Chalapathi V. Chamola, C. Tham. Latency aware mobile task assignment and load balancing for edge cloudlets. In *2017 IEEE International Conference on Pervasive Computing and Communications Workshops (PerCom Workshops)*, pages 587–592, Kona, 2017. IEEE, IEEE.
- [15] S. Marusic M. Palaniswami J. Gubbi, R. Buyya. Internet of things (iot): A vision, architectural elements, and future directions. *Future generation computer systems*, 7(29):1645–1660, 2013.

- [16] Richard Hill, James Devitt, Ashiq Anjum, and Muhammad Ali. Towards in-transit analytics for industry 4.0. In *2017 IEEE International Conference on Internet of Things (iThings) and IEEE Green Computing and Communications (GreenCom) and IEEE Cyber, Physical and Social Computing (CPSCom) and IEEE Smart Data (SmartData)*, 2017 IEEE International Conference on Internet of Things (iThings) and IEEE Green Computing and Communications (GreenCom) and IEEE Cyber, Physical and Social Computing (CPSCom) and IEEE Smart Data (SmartData), pages 810–817. IEEE, 6 2017.
- [17] Flavio Bonomi, Rodolfo Milito, Jiang Zhu, and Sateesh Addepalli. Fog computing and its role in the internet of things. In *Proceedings of the First Edition of the MCC Workshop on Mobile Cloud Computing*, MCC '12, pages 13–16, New York, NY, USA, 2012. ACM.
- [18] Adam A. Alli and Muhammad Mahbub Alam. Secoff-fciot: Machine learning based secure offloading in fog-cloud of things for smart city applications. *Internet of Things*, 7:100070, 2019.
- [19] Hussain Al-Aqrabi and Richard Hill. A Secure Connectivity Model for Internet of Things Analytics Service Delivery. *2018 IEEE SmartWorld, Ubiquitous Intelligence & Computing, Advanced & Trusted Computing, Scalable Computing & Communications, Cloud & Big Data Computing, Internet of People and Smart City Innovation (SmartWorld/SCALCOM/UIC/ATC/CBDCOM/IOP/SCI)*, pages 9–16, 2018.
- [20] Muhammad Usman Yaseen, Ashiq Anjum, Omer Rana, and Richard Hill. Cloud-based scalable object detection and classification in video streams. *Future Generation Computer Systems*, 80:286–298, 3 2018.
- [21] Muhammad Usman Yaseen, Muhammad Sarim Zafar, Ashiq Anjum, and Richard Hill. High performance video processing in cloud data centres. In *Proceedings - 2016 IEEE Symposium on Service-Oriented System Engineering, SOSE 2016*, pages 180–189. Institute of Electrical and Electronics Engineers Inc., 5 2016.
- [22] J. M. García-Campos, D. G. Reina, S. L. Toral, N. Bessis, F. Barrero, E. Asimakopoulou, and R. Hill. Performance evaluation of reactive routing protocols for vanets in urban scenarios following good simulation practices. In *2015 9th International Conference on Innovative Mobile and Internet Services in Ubiquitous Computing (IMIS)*, pages 1–8. Institute of Electrical and Electronics Engineers Inc., 10 2015.
- [23] D. Zhang D. Li C. Zhang S. Wang, J.Wan. Towards the smart factory for industrie 4.0: A self-organized multi-agent system assisted with big data based feedback and coordination. *Computer Networks*, 101:158–168, 2016.
- [24] D. Li C. Zhang S. Wang, J. Wan. Implementing smart factory of industrie 4.0: An outlook. *International Journal of Distributed Sensor Networks*, 12(1), 2016.
- [25] Richard Hill, Simon Polovina, and Martin D. Beer. From concepts to agents: Towards a framework for multi-agent system modelling. In F. Dignum, V. Dignum, S. Koenig, S. Kraus, M. P. Singh, and M. Wooldridge, editors, *Proceedings of the Fourth International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 05)*, pages 1155–1156. Utrecht, The Netherlands, July 25-29 2005. ACM Press.
- [26] Martin Beer, Richard Hill, Wei Huang, and Andrew Sixsmith. An agent-based architecture for managing the provision of community care: The inca (intelligent community alarm) experience. *AI Communications*, 16(3):179–192, 2003.
- [27] Martin Beer, Wei Huang, and Richard Hill. Designing community care systems with auml. In *IEEE International Conference on Computer, Communication and Control Technologies (CCCT2003)*, 2003.
- [28] Richard Hill, Dharmendra Shadija, and Mo Rezai. Enabling community healthcare with microservices. In *The 16th IEEE International Conference on Ubiquitous Computing and Communications (IUCC 2017)*, Guangzhou, China, December 2017. IEEE Computer Society.
- [29] Hussain Al-Aqrabi and Richard Hill. Dynamic Multiparty Authentication of Data Analytics Services Within Cloud Environments. *2018 IEEE 20th International Conference on High Performance Computing and Communications; IEEE 16th International Conference on Smart City; IEEE 4th International Conference on Data Science and Systems (HPCC/SmartCity/DSS)*, pages 742–749, 2018.
- [30] X Xu, L Zhu, Y Liu, and M Staples. Resource-oriented business process modeling for ultra-large-scale systems. In *2nd international workshop on ultra-large-scale software-intensive systems*, pages 65–68, 2008.
- [31] Richard Hill, Hussain Al-Aqrabi, Philip Lane, and Hamza Aagela. Securing manufacturing business intelligence for the industrial internet of things. In Xin-She Yang, Simon Sherratt, Nilanjan Dey, and Amit Joshi, editors, *Fourth International Congress on Information and Communication Technology*, volume 2 of *Advances in Intelligent Systems and Computing*, Singapore, 1 2019. Springer Singapore.

- [32] Hussain A. Al-Aqrabi, Lu Liu, Richard Hill, and Nick Antonopoulos. *A multi-layer hierarchical inter-cloud connectivity model for sequential packet inspection of tenant sessions accessing BI as a service*, pages 498–505. Proceedings - 16th IEEE International Conference on High Performance Computing and Communications, HPCC 2014, 11th IEEE International Conference on Embedded Software and Systems, ICESS 2014 and 6th International Symposium on Cyberspace Safety and Security, CSS 2014. Institute of Electrical and Electronics Engineers Inc., 8 2014.
- [33] J Lewis and M Fowler. *Microservices: a definition of this new term*. 2014.
- [34] Dharmendra Shadija, Mo Rezai, and Richard Hill. *Microservices: Granularity vs. performance*. In *UCC 2017 Companion - Companion Proceedings of the 10th International Conference on Utility and Cloud Computing*, pages 215–220. Association for Computing Machinery, Inc, 12 2017.