IoT reliability: a review leading to 5 key research directions

Samuel J. Moore¹ · Chris D. Nugent¹ · Shuai Zhang¹ · Ian Cleland¹

Received: 27 February 2020 / Accepted: 22 July 2020 / Published online: 7 August 2020
© The Author(s) 2020

Abstract

The Internet of Things (IoT) is rapidly changing the way in which we engage with technology on a daily basis. The IoT paradigm enables low-resource devices to intercommunicate in a fully flexible and pervasive manner, and the data from these devices is used for decision-making in critical applications such as; traffic infrastructure, health-care and home security, to name but a few. Due to the scarce resources available in these IoT devices, being able to quantify the reliability of them is a critical function. This report presents a detailed evolution of the area of reliability measurement, followed by an in-depth review of the state-of-the-art for quantification of reliability in the IoT, revealing the many challenges associated with this task. From this in-depth review, a set of key research directions for IoT reliability is determined. Despite the critical nature of the research area, at this current moment, this study is the first detailed review available in the area of assessing IoT reliability.

Keywords IoT · Reliability · Review · Research Directions · Security

1 Introduction

Computing has been the fastest growing field of the last century. Computing systems now pervade the fabric of our everyday lives. We cannot make a purchase from a store, withdraw money from our bank accounts or visit a hospital without interacting with a computing system. Computing systems are now relied upon for many mission-critical systems, such as aircraft control systems, military systems and nuclear power plants. With the criticality of our computing systems in mind, it is vital that there are methods in place to ascertain the reliability of such systems. One of the fastest-growing fields within computing is the Internet of Things (IoT). The IoT is expected to grow to an immense size over the next number of years. In 2011 Cisco predicted that there would be 50 billion devices connected to the IoT by 2020 (Evans 2011). These huge claims have also triggered predictions of monetary investments reaching into the trillions by 2020 (Rayes and Salam 2016). While these

numbers suggest a truly rapid growth in IoT, there are still many research challenges which must be solved for IoT to become fully integrated into our day-to-day lives. These barriers include trust, security, interoperability, reliability, scalability, performance, availability and mobility (Al-Fuqaha et al. 2015; Wang 2018; Ahmed et al. 2017; Saini 2016). These areas represent significant research challenges that must be addressed if we are to allow IoT to become the ubiquitous technology that it has set out to become (Sicari et al. 2016). If the vision of IoT is to be fully implemented in our homes, cities and workplaces then we will be trusting intelligent systems to make thousands of decisions daily that will have profound impact on our lives, through applications such as; home security (Ghorbani and Ahmadzadegan 2017), providing healthcare services to patients (Da Li et al. 2014) and monitoring critical traffic infrastructure (Singh et al. 2014).

The IoT is tasked with considering devices which may be extremely constrained by nature (Chiang and Zhang 2016). Considering that the IoT will be responsible for managing key infrastructure such as traffic lights, critical health systems and home security, it is easy to appreciate how the impact of unreliable IoT infrastructure may affect the decision-making of the system in a potentially severe or fatal manner (Fekade et al. 2017). The reliability issue does not end at the device and hardware layer either, there is also the consideration of the reliability of the network layer. This can often be difficult to determine due to the

This research is supported by the BTIIC (British Telecom Ireland Innovation Centre) project, funded by BT and Invest Northern Ireland.

✉ Samuel J. Moore
moore-s34@ulster.ac.uk

¹ School of Computing, Ulster University, Shore Road, Newtownabbey BT37 0QB, Northern Ireland

heterogenous nature of the devices connected to it and how they transmit data, often wirelessly over lossy links. Beyond the data transmission, there is also the issue of actuation to be considered. This raises an important question, how can we be assured that decisions are taken by the system based on robust information, given the challenges at the lower layers of the architecture? There must also be mechanisms in place to determine the accuracy of the decision-making models that determine the actuation of the system. Incorrect decision-making at this level could potentially be life-threatening for end-users, making this a key research issue (Sato et al. 2016). The vulnerabilities of IoT devices are becoming a prominent issue in the consumer and government industries. In October 2018, the UK government issued a set of guidelines describing minimum standards for smart-home devices, in order to protect the consumer (DDCMS 2018). This is demonstrative that IoT system reliability is something that will need to be addressed comprehensively in order for the technology to fully mature. If we are able to successfully quantify the reliability of our IoT infrastructure and which applications can avail of its service, this will then allow us to use the quantified reliability metric to reason about the fitness for purpose of our critical IoT infrastructure.

This report represents a novel in-depth study of reliability in relation to IoT, from first examining the fundamentals of reliability in engineering; these principles are then applied to reliability in computing and IoT. An exhaustive literature review, which to the best of the authors' knowledge is the first of its kind, is also included in this report. This is followed by a summary of the current state-of-the-art research into reliability of IoT. Finally, this report proposes a novel set of research directions which are crucial in the advancement of IoT reliability, based upon the findings of the in-depth study.

The rest of this paper is organised as follows: Section 2 describes the method and process used to perform this research. Section 3 is an overview of the meaning and origins of reliability engineering and how it applies to computing. Section 4 is a detailed view of reliability in IoT, and the challenges that make the quantification of reliability difficult in IoT. Section 5 is a detailed literature review of current research into quantifying reliability in IoT systems. Section 6 then discloses the five key research directions gleaned from the in-depth study conducted in this work, followed by the conclusion of the work in Sect. 7.

2 Methods and processes

At the time of writing, to the best of the authors' knowledge, this work represents the first work to review and summarise the practice of quantifying and measuring reliability in the IoT. This indicates that, in general, the area of reliability

quantification in the IoT is as yet under-researched, and would benefit from an in-depth review. Moreover, in order to guide the efforts of future researchers who may apply the knowledge herein to their own research projects, it would be of benefit for this in-depth review to summarise the key knowledge points, and then synthesise them into a list of key research directions. Crucially, for a piece of work such as this, there must be a process in place in order to ensure that the research is carried out in a robust and reliable manner. As such, this section details the three-stage methodology that was used to perform this research. This methodology aims to support the research in achieving the following contributions to knowledge:

1. Define reliability in the IoT through researching state-of-the-art in other well-established disciplines, such as biomedical engineering and engineering.
2. Using this reliability information from contribution 1, apply it to the scope of the IoT, detailing the key requirements for IoT reliability
3. Perform an in-depth literature review surveying the state-of-the-art in IoT reliability quantification
4. Analyse the current state-of-the-art research and derive key research directions for the field.

2.1 Defining reliability in the IoT

Given the nascent nature of this field, as described earlier in this section, it becomes necessary to formulate a definition for the IoT. While reliability is a mature field in the related disciplines of engineering (Bradley 2016) and software engineering (Xie et al. 2004), there is no agreed upon definition for the IoT. As such, it becomes necessary to review the core principles of reliability in these other domains. Highly cited and relied upon publications were selected from the industries of engineering, biomedical engineering and computing which described the practice of reliability definition and quantification. From these works, a clear definition of reliability is then developed. This core definition of reliability is then taken and applied against the backdrop of the IoT. This ensures that the application of any reliability definition is firmly grounded in highly cited academic and scientific articles.

The definition of IoT reliability is then further expanded, and viewed in an end-to-end sense, using the core physical architecture of the IoT as a structure to discuss the applications of reliability in the IoT.

2.2 Selection of works

With reliability then firmly defined in a wider sense, then narrowed into an IoT-specific scope and definition, works were then selected to demonstrate the current state of the

art in performing reliability analysis in the IoT. In order to ensure that the works included in this review were of significant scientific standard, the following steps were applied in sourcing and including the literature:

1. Wide searches were conducted in the following databases; IEEE Xplore, ACM Digital Library and Google Scholar (including Springer and Elsevier).
2. Search terms used were “IoT”, “Reliability”, “Reliable”, “Trust”, “Dependability”, “Quality” (in any combination) in the title of the work.
3. Prior to 2010, IoT was a seldom-used term in research literature, as such, works before 2010 were excluded before to ensure currency.
4. Works were then excluded which did not address some method of quantifying, measuring or directly aiming to assess or improve reliability in the IoT.

2.3 Analysis and research directions

Once the works were selected, each of them were reviewed and a detailed summary presented detailing their contributions to the field of IoT reliability research. This information is then synthesised into benefits and shortcomings, which are analysed and presented in a summary table.

Finally, the analysis of each of these surveyed works was used to derive a set of research directions for the research area. Given that, at the time of writing, this review paper is the first to specifically study reliability in the IoT, this research output represents a trusted and evidenced roadmap to improving reliability in the IoT.

3 The science of reliability

Reliability, at a fundamental level, is concerned with the study of failures (Fries 2006). More specifically, it is concerned with how failures are caused, how they can be addressed and how they can be prevented. There are many misconceptions regarding what reliability actually represents. It is not as simple as testing and re-testing a device until relative satisfaction is reached. Reliability can be represented by a formal definition which includes four key requirements. Fries (2006) defines these requirements of reliability by stating that devices must be able to; perform a required function, perform without failure, perform under stated conditions and operate for a specific period of time. Therefore, the specification for reliability requires that we fully identify the expected conditions of use, what constitutes proper function and what constitutes a failure (Mavrogiorgou et al. 2018). The remainder of this section will cover the key areas in establishing reliability engineering. First, a definition is presented to illustrate the difference between

two often misused terms; quality and reliability. Then basic failure patterns are discussed, and how they impact new services. Finally, a description of reliability as it applies to the field of computing is presented, alongside standard metrics used to quantify reliability in computing.

3.1 Quality versus reliability

The terms “quality” and “reliability” are often misunderstood. Sometimes, these terms are used interchangeably, however, there are important distinctions between the two terms. Both terms exist to describe a characteristic of a product or system. Fries (2006) determines that the main difference between these two terms as being the temporal nature of quality. The term “quality”, as defined in ISO 9000 (ISO 2015) is the “ability to consistently provide products and services conforming to their requirements”. In this definition of quality, there is no stipulation of a time period for which these requirements must be met or continue to be met in the future. A quality test, therefore, reflects only a snapshot of a particular time at which quality requirements are either met or not met. Reliability, however, refers to the performance of a system or product over a specific window of time. This is an important distinction between the two terms, especially when it comes to assuring continued adequate performance of a device or system over time. In essence, we can evaluate quality in the IoT, however, this will not offer us any assurance in the continued successful operation of the deployment. A product or system can be designed and released to a very high standard of quality, however, this will not provide any information with regards to how often the product fails. Moreover, we cannot use quality to ascertain the probability that the system or product will be operating without fault at a given time. Quantified reliability measures, on the other hand, allow us to ascertain vital pieces of information regarding the up-to-date operational state of the IoT deployment. These pieces of information can include, how often a device fails, the average interval between failures, the average time taken to repair a component and the probability that a component will need to be replaced by a certain date.

3.2 Common failure patterns

There are three main patterns of failure which are defined in the field of reliability engineering; infant mortality, constant failure rate and wearout failure (Bradley 2016). Infant mortality refers to failures that occur predominantly early in the lifecycle of a product and gradually are neutralised over time. Wearout failure patterns are observed when a device begins to exhibit an exponentially higher number of errors compared to a previously consistently low number of errors. This failure pattern indicates that a device is nearing the end of its useful life period (Fries 2006). The constant failure

rate describes a pattern where the number of errors within a given period of time remains constant. For example, for a device to have a constant failure rate we would expect that the same total count of errors to occur in each calendar month, though these do not necessarily need to occur at the same times each month. These three error patterns are characterised in Fig. 1, the dotted vertical lines mark the beginning and the end of the useful life period where the failure pattern is at a constant rate.

3.3 Reliability in computing

There are many ways of assessing reliability in computing. The most appropriate method may depend upon the nature and function of the system being assessed. Reliability should be a quantitative measure which broadly represents the ability for a computer system to perform its intended function (Xie et al. 2004). Xie et al. (2004) outline several key metrics that help to define reliability in computing. Mean Time to Failure (MTTF), closely related to Mean Time Between Failures (MTBF) is the expected lifetime that the system will operate normally before a failure occurs. The failure rate function, also known as the hazard function, is a metric that helps to define the rate of system aging. The failure rate is the probability that a device will fail within a specified window of time. The failure rate function when used to evaluate hardware would be expected to follow an exponential distribution, which thereby allows us to reason about the aging and deterioration of the hardware. When used in software, however, the failure rate would remain constant because software does not age or deteriorate physically. Maintainability, according to Xie et al. (2004), is a metric that represents the probability that a failed system can be returned to normal operation within a given period of time. Availability is a metric representing the probability that a system will be operating as normal in a given period of time. Availability and maintainability are closely related, however, they differ in one key aspect: availability concerns the period of time in which a system is expected to be operating

normally, whereas maintainability concerns the period of time in which a fault has occurred. Of course, beyond this technical definition, maintainability also is concerned with the continuing and ongoing operation of a system - this may relate to such tasks as meeting new requirements, refactoring and restructuring code and other maintenance tasks which contribute towards maximising the useful life period of a system.

Within the study of computing reliability there are four key areas which have different approaches in establishing reliability; hardware, software, network and system (Xie et al. 2004). These different areas each have reliability analysis methods that are uniquely suited to the requirements and issues in the area. Hardware reliability concerns the reliability over time for the physical components of a computer system, such as CPU, disk and sensors. These components are prone to wear and tear, and therefore we would expect the reliability to reduce over time for these components. Software components, on the other hand, should not be subject to physical wear and tear, therefore we would not expect to see a decay in reliability over time (Mavrogiorgou et al. 2018). Network reliability concerns the network performance over a given period of time, which is determined by a blend of hardware and software. Systems reliability is a combination of all the components combined, and there are specialised techniques to analyse this.

Reliability quantification methods in computing are well established and understood, as presented in this section. With the onset of the novel IoT paradigm, however, it is important that we formulate IoT-centric reliability quantification methods. These methods must suit the unique nature and constraints of IoT, which is discussed in the following section.

4 Reliability in the Internet of Things

The definition for IoT is fundamental in understanding the problem of reliability within the paradigm. The definition of IoT is often under-represented and ill-defined (Atzori

Fig. 1 The three common failure patterns; infant mortality, constant (steady) failure rate and wearout

et al. 2017). Often, the IoT is crudely defined as being able to add internet connectivity to every-day devices, in effect allowing “*your toaster to talk to your fridge*”. While this statement is true for some part of the IoT, it does not encompass the whole paradigm of the IoT. A useful starting point in defining the IoT paradigm is by considering the key components of IoT. These components are; sensing, actuating, communication, services and applications

(Rayes and Salam 2016). These four components can then be mapped to an architecture for the IoT, as presented in Fig. 2. Sensing and actuating are carried out at the lowest layer of the architecture, also referred to as the device layer. The next layer up, the edge layer, enables the communication between the devices and the application layer. Typically, this communication is enabled by semi-capable devices behaving as hubs, collecting data from the sensors

Fig. 2 The four layers of the IoT architecture; cloud layer, service management layer, fog layer, and device layer

and relaying it into the cloud and sending commands to the actuators as necessary.

With the key components of IoT in mind, we can now form a full definition of IoT, which is a paradigm which enables interconnectivity in anything and everything to create monitoring and control infrastructure which can be used in applications to enrich everyday user experience (Rayes and Salam 2016).

To begin understanding the problem space of reliability in IoT, it is best to use the architecture, as presented in Fig. 2, as a reference point. We can then observe reliability issues in each of the layers of the architecture and understand how they contribute to the problem.

4.1 Device reliability

From a device perspective, that is the sensors and actuators, the first problem we can observe is the highly constrained nature of these devices (Al-Fuqaha et al. 2015). These constraints concern battery, memory and computational capacity (Kouicem et al. 2018). Battery is a concern for IoT applications, because often the application layer is unaware of the remaining battery left on the device thereby making it difficult to determine when the device requires a battery replacement (Shi et al. 2016). This battery life concern is further compounded when we consider that devices may be located in places that are physically difficult or dangerous to reach to replace. The memory and CPU constraints on the devices limit the device's ability to store complex encryption methods, meaning that IoT devices must rely on lightweight encryption to protect the data being transmitted by the device (Rayes and Salam 2016; Alaba et al. 2017).

Another issue evolves from the constrained nature of the devices when it comes to updating the limited firmware of these low-powered sensors. It is impractical, due to the lack of power and implications on battery life for the device, to connect to a cloud service routinely and check if new firmware needs to be downloaded and installed on the device (Chiang and Zhang 2016; Yaqoob et al. 2017; Allhoff and Henschke 2018). This leads to a scenario where devices could potentially be operating with outdated firmware, thereby leaving them vulnerable to security breaches.

The sensors and actuators that are used in the IoT are often deployed in remote and distant locations, and can often be subject to harsh environmental conditions such as heat, freezing temperatures, mechanical wear, vibration, and moisture (Rayes and Salam 2016). As discussed in Sect. 2 in this report, there is a need to determine the "useful life" period of a device, so that we can determine when the device needs to be retired. This useful life will shorten if the device is employed in a harsh environment, therefore, we could expect to see great variances of device lifetime for identical

devices deployed in different environments, which results in the system reliability being difficult to manage.

Another concerning aspect with regard to device reliability in IoT, is the propensity for sensors to "fail-dirty" (Karkouch et al. 2016, 2017). This phenomenon concerns a scenario where a sensor continues to send erroneous readings after having suffered a failure. This is a well-known, yet little understood, problem that is pervasive in IoT environments. In particular, this issue is hard to diagnose because the sensor appears to be operating normally. The impact of a false reading being sent in an IoT environment can be critical, when we consider that actuation often has physical impact on human lives.

4.2 Communication and network reliability

Mobility is one of the key expectations of an IoT network (Al-Fuqaha et al. 2015), whereby users of the network can dynamically move between applications while the device onboarding and identification happens seamlessly in the background. Global addressing, however, is a difficulty in IoT (Rafferty et al. 2018), given that manufacturers do not co-ordinate to provide globally unique identifiers for all IoT devices. This means that the responsibility of assigning unique identification resides within the IoT network itself. When we consider that IoT devices are expected to be mobile, this creates a problem given that the device ID might differ across different networks, meaning that we might lose traceability of the device. This then introduces a reliability concern when it comes to tracking or auditing the device as it moves through different IoT applications.

Internet Protocol (IP) is the current de-facto standard for communication and identification in traditional networks. IP in its current state is, however, not well suited to the IoT (Tsai et al. 2014). Introducing new protocols into this problem space will require these new protocols to mature quickly, which is not always easy. This problem is exacerbated further when we consider the implications of unique addressing. IPv4 has a 32-bit length address, which creates room for 4.3 billion addresses, keeping in mind the predictions of 50 billion devices discussed previously in this paper, it becomes clear that IPv4 is not suitable to fulfil the vision of IoT. This problem is further compounded by the fact that IPv4 ran out of addresses in 2010 (Evans 2011). As such, it becomes necessary to implement a protocol with suitable addressing space, such as IPv6, which boasts an address space of 128 bits, allowing space for

$$3.4 \times 10^{38}$$

addresses. This new addressing space, however, creates problems for constrained devices, not all of which are

capable of handling the overheads required for the address (Al-Fuqaha et al. 2015).

A remedy to this large address overhead is offered by the 6LoWPAN protocol (Kushalnagar et al. 2007). 6LoWPAN is able to compress the header size of the IPv6 packets in order to make them compatible with the IEEE 802.15.4 standard (Al-Fuqaha et al. 2015), and thus better suited to the IoT. These new and emerging standards to cope with the new requirements of the IoT contributes to the creation of a landscape of disparate standards and protocols among IoT devices and deployments targeted for communication for constrained devices in IoT networks. Given the lightweight and constrained nature of some of these protocols, not all of them feature quality of service (QoS) guarantees, meaning that the reliability of the network connection becomes harder to assess.

Karkouch et al. (2016) indicate in their study that due to the scarce resources and intermittent communication, the network is liable to drop readings, or produce unreliable readings. This notion that readings can be dropped due to the inherent nature of IoT networks is an area for concern, especially given that IoT infrastructure is often responsible for managing mission-critical applications (Fekade et al. 2017).

4.3 Application layer reliability

The application layer of the IoT paradigm is not subject to the same constraints of either the network or the device layer of the architecture. It is important to note that in many cases the reliability of the application layer is a function of how reliable the lower layers of the architecture are. If anomalous data is sent from the device through the network into the application layer, this will reduce the reliability of the application. In this regard, it is important that the application layer has sufficient anomaly detection techniques to eradicate errors and maintain the reliability of the application. Given that IoT networks feature a heterogeneous range of constrained devices transmitting many pieces of information in different formats, this task can be difficult (Abeshu and Chilamkurti 2018).

While the application layer doesn't suffer from the physical constraints of the device layer, there is still a need to manage the reliability of the applications that are being deployed. A study Moore et al. (2019) which observed the impact of anomalous data on classification in the IoT application of human activity recognition found that some classifiers were considerably more vulnerable to errors than others, and that the preparation method of the data can also make the application more vulnerable to failure. With this in mind, developers must make a conscious effort to establish and understand the reliability of applications being hosted in IoT infrastructure, in order to prevent critical errors from creeping into the system.

4.4 Toward an effective solution for reliable IoT systems

The issues of reliability at the three levels of the architecture in IoT combine to create a vulnerable landscape for IoT, which often leads to anomalous data being generated and sent through the network. This notion demonstrates a strong need for effective, quantifiable reliability measures that will allow us to reason about the fitness for purpose of our IoT systems. The issue of anomalous data is highly problematic to the IoT vision, given that actuations will be made on the basis of this data which could, in the most severe cases, threaten human lives (Fekade et al. 2017). Therefore, it is essential that any framework that aims to assess and quantify reliability in IoT must be able to detect the presence of anomalies in the system. Once reliability has been quantified, this opens up a new opportunity to further enhance the robustness of the system by placing a human in the loop (HITL). HITL is an essential ingredient for the future of reliability in IoT, and was identified as a key future research area in Stankovic (2014). The HITL paradigm opens up opportunities for detecting and resolving reliability issues in critical IoT infrastructure. The use of a human observer brings an element of domain expert knowledge to the application, which allows the human to synthesise information presented by the system with their own expert knowledge to come to an informed conclusion about the reliability of the system. Furthermore, humans enable the assessment of ground truth, such as a true temperature value, which can help to verify a machine reading. This concept is relatively novel in IoT research and, to date, no reliability studies have opted to use a HITL method to assist in reliability assessment in combination with classical reliability models. This novel combination would be a step towards a new and effective solution.

This section reviewed the key concepts and requirements for IoT reliability. The vulnerable landscape which the IoT occupies presents a clear requirement for the research community to design and implement frameworks and solutions which might aid in assessing and understanding reliability for our key IoT infrastructure.

5 Current research toward IoT reliability quantification

The definition of reliability, as discussed in the previous sections, has a strong element of quantification associated with it. Reliability, as defined in Sect. 3, is not a subjective science, and therefore mechanisms aiming to assess reliability should be objective and quantifiable in their nature. There is also a heavy focus within reliability in defining and using

metrics to assess the reliability of components and systems. Research in the area of IoT reliability has been conducted to enhance reliability at various levels of the IoT architecture. This section summarises the research available in the areas of device reliability, data quality, network reliability and anomaly detection, all of which represent key areas for improvement of IoT reliability.

5.1 Device reliability

Several authors researching IoT device reliability integrated classical reliability metrics into IoT-centric solutions. Reliability, failure rate, availability, and MTTR were quantified by Zin et al. (2016). The work proposed a probabilistic model for measuring reliability in connected IoT devices positing that the failure structures of IoT devices adhere to a certain probability distribution. The authors define the reliability measure $R(t)$ as being the probability that the device is operating correctly at time interval $[0, t]$. This probabilistic function allows estimation of the expected time to failure, availability and reliability for a given IoT device. Meanwhile, Mavrogiorgou et al. (2018), included Mean Time to Repair (MTTR), MTTF, MTBF, and availability metrics in their work, which proposed a mechanism for capturing the reliability of heterogeneous IoT devices. This mechanism considered both known and unknown device types and sought to differentiate between which devices were reliable and which were not, with the goal of collecting data from the reliable ones and discarding data from unreliable devices. The mechanism consisted of four stages: devices recognition, specifications classification, reliability estimation and reliability validation. Using this mechanism, the authors were able to build a ranking of connected fitness devices based upon their reliability results from known reliability metrics. Lastly, Kim (2016) used reliability, failure rate and recoverability in their study which proposed a weighted model to quantifying reliability in the IoT. The model consisted of four quality criteria; functionality, reliability, efficiency and portability. Metrics were defined within these criteria which were assigned weights so that the model could provide a total score for the quality of the IoT application. The model was then evaluated in a virtual environment and scores were produced for each of the metrics. This model provides weighting, however, each criterion was weighted evenly in this experiment. These classical metrics provide a useful starting point in the quantification of IoT reliability, but have not yet matured in capability and cannot attest to reliability across all levels of the IoT architecture.

Moving away from the classic well-defined reliability metrics, some non-standard reliability metrics have been designed and implemented in recent studies. Saini (2016) presented a model to evaluate trust factor and reliability over a period of time (ROPT) for IoT systems. Due to the notion

that identical IoT sensors might be deployed in drastically different environments (i.e., exposed to varying levels of humidity, temperature, wind) these identical IoT sensors might exhibit different expected lifetimes. The author proposed that ROPT is calculated for every individual device and gateway in the IoT system in order to gain a full understanding of how reliable the system is. The author also presented a trust factor rating scale allowing us to reason on how some IoT applications require higher levels of trust, i.e. defence systems, and therefore higher levels of availability. This study only uses one metric to determine the reliability of the system, and cannot represent the entire picture of reliability in the IoT. Li et al. (2012) also proposed three non-standard reliability metric definitions to observe real time quality of data collected from devices in IoT environments. The study validates the implementation of these metrics by applying them to two real-world open source datasets. The three metrics defined were; currency, availability and validity. Implementing the metrics onto real-world datasets validated that it was possible to calculate these metrics in real-time, but this was not able to attest to the effectiveness of the applied metrics in identifying data quality issues in IoT.

A more complete framework for managing quality and reliability is proposed by Sicari et al. (2016) and Sicari et al. (2014). This architecture is designed to quantify the security and quality of individual devices in IoT applications. The model used NOS (Networked Smart Objects) to extract metadata from IoT nodes in a network (Rizzardi et al. 2016). The parameters extracted from a security perspective were confidentiality, integrity, privacy and authentication. The collected parameters for quality were accuracy, precision, timeliness and completeness. Each parameter was attributed an index score ranging from zero to one, which reflected the effectiveness of the node with regard to that parameter. The model was tested using Raspberry Pis and sensors from a meteorological station and it was successfully able to calculate the specified parameters. This model concerns the data quality characteristics of IoT nodes though the quality metadata, which is not sufficient in describing the holistic reliability of an IoT system. The security metadata provides some insight into how secure a given node is in an IoT system, this could be enhanced by adding anomaly detection.

The research presented in this section is valuable in aiding the understanding of how reliable and prone to failure the devices in our IoT infrastructure are. These pieces of research help form an understanding of how some of this information can be quantified, using metrics like availability, MTBF and MTTR. Nevertheless, the quantification of hardware reliability is only one step in the solution to overall IoT reliability. These research studies are unable to attest to reliability at the network level or make an assessment about the likelihood of the system providing anomalous data or falling victim to a spreading threat.

5.2 Network reliability

Beyond being able to reason about the fitness of our IoT devices, we must also be able to attest to the reliability of the network infrastructure that forms the backbone of IoT communication. Generally speaking, there are two forms of network reliability studies which are discussed in this section; studies for enhancing QoS in networks, and studies aimed at quantifying reliability metrics for networks. This section presents the current state-of-the-art research in IoT networks reliability.

A novel IoT network QoS metric was proposed by Maalel et al. (2013) in their work, which designed a lightweight and energy efficient routing protocol to enhance and measure reliability in IoT applications, specifically emergency applications. Emergency applications in the IoT require a rapid response for alarms that have been raised. The work proposed a mechanism called AJIA (Adaptive Joint Protocol based on Implicit ACK) for packet loss and route quality evaluation. The mechanism relies upon the broadcast nature of the protocol, where messages are broadcast to all nearby nodes. The nearby nodes can therefore “overhear” the message being sent. This overhearing function is used rather than traditional ACK messages to ensure reliability of the message being sent. The links between nodes are then evaluated with a metric called Link Quality Indicator (LQI), which uses the history of packet loss in the link to determine the reliability of that particular path. Other QoS metrics, such as delay throughput, and packet loss, were quantified by Kamyod (2018). This work employed Riverbed’s Optimized Network Engineering Tools (OPNET) to observe these network reliability parameters in a smart agriculture scenario. These parameters were monitored so that they might provide some information as to how reliable the overall end-to-end IoT system was. The study found that increasing the number of nodes in the network saw longer packet delays and significantly longer transmission times and packet loss. Brogi and Forti (2017) proposed a general model for a QoS-aware IoT infrastructure, based on the fog computing paradigm. The model allows IoT applications to generate QoS profiles in order to request certain QoS characteristics from the Things it interacts with. Each communication link in the IoT system has an associated QoS profile, which allows the model to determine the potential latency and bandwidth for an application to things communication. The model only considers latency and bandwidth, which is a limited subset of QoS characteristics which would not fully represent the reliability of the network at a given point in time.

Further IoT network QoS metrics, embedded in a management framework, were examined in a study by Al-Masri (2018), which presented a microservices QoS management framework (mQoS M) for use in Industrial IoT (IIoT), which is a QoS-aware middleware that monitors the behavior of

microservices in order to determine the “best” microservice amongst all discovered microservices. This information can then be used by IoT architects to decide if they wish to integrate the microservice. This framework monitors the following parameters; response time, throughput, availability, reliability and cost. The model presents a useful step towards generating a situational awareness of the IoT system with regards to reliability and performance, however, it has not been scaled up beyond microservices in an IoT environment.

An approach of reliability modelling using Generalised Stochastic Petri Net (GSPN) was proposed by Li and Huang (2017). This approach theorised mathematical models at edge nodes to provide statistics on the performance of IoT devices. The metrics calculated were time consumption, response time, failure rate and repair times. These metrics only speak to the performance of the device to edge layer and offer a very limited view of network performance which does not present a holistic view of IoT reliability. A gateway redundancy model was proposed by Sinche et al. (2018). This work made use of redundancy at both the ISP (Internet Service Provider) level and Gateway (edge node) level. This model tested three cases, an IoT infrastructure with no redundancy, an IoT with gateway redundancy and an IoT with gateway and ISP and gateway redundancy. The model was tested using a physical IoT testbed, wherein the devices were communicating using the I2C bus protocol. RTT (return trip time) was used as the performance metric to determine the effectiveness of the model. The results shown in the study found that the model which did not use the redundancy approach saw the RTT increase by 14% during fault conditions, whereas the redundancy models resulted in only a 1% increase in RTT. This study considers reliability at the network and cloud level only. Therefore, it does not consider the reliability of the physical devices, or their propensity to fail at any given time. This study also does not consider the heterogenous nature of IoT communication protocols. Alam (2018) presented a framework to handle reliability issues in IoT based on the TCP (Transmission Control Protocol). There are three components to the framework; the reliability calculator, the reliability controller and the reliability handler. The framework uses delay to determine the failure-state of the IoT system. If high levels of delay are observed by the reliability calculator, the reliability controller will attempt retransmission and the reliability handler will initiate a broadcasting mode and enter a power-saving state. This framework only deals with the delay QoS metric in IoT, thus it cannot represent the full state of reliability in the network.

The research presented in this section shows that while some attempts have been made to enhance reliability in IoT networks, both by enhancing the network’s QoS and by monitoring and quantifying network reliability, there is currently

not a research approach which successfully combines device and network reliability into one framework.

5.3 System reliability

Some research has also been conducted to evaluate IoT reliability at a system level. These approaches are at a high level, and do not capture the individual detail for reliability, such as which devices are responsible for failures, or which parts of the network are responsible for traffic problems.

Behera et al. (2015) proposed a method of modelling reliability in a service oriented IoT. Specifically, algorithms were proposed to evaluate reliability in a Centralised Heterogeneous IoT Service System (CHISS). The authors proposed that reliability could be measured by modelling the availability of the program to run the service, the availability of input required for the service to run and the service reliability of subsystems associated with the system. The algorithms were tested on a case study of a fire alarm system, which was running under normal operation at the time. The algorithms were able to determine if the program and file was available for each component in the IoT system. This methodology did not, however, consider the notion that the IoT components could fail at any moment and begin sending anomalous data, or that the network could fall victim to a spreading threat or virus. In order to present a true reflection of reliability, it is necessary to have a mechanism which can alert the user to failures in the system before critical actuations are made.

Kharchenko et al. (2017) proposed the use of a Markov model to predict the reliability requirements of an IoT system. The Markov model considered that the application could be in a range of 15 states, from normal condition to complete failure. The probabilistic nature of the Markov model facilitates prediction that the system will move from one state to the next and can establish the probability of a failure at a given point in time. This model only considers the states specified in the design of the model and is not capable of reacting to new situations that were not catered for in the design of the model.

5.4 Anomaly detection

With the vulnerable state of IoT networks, given their constrained devices and highly mobile nature, it is essential that any framework which intends to quantify the reliability of an IoT infrastructure must have knowledge of the potential presence of anomalous data in its applications. This anomalous data could have severe consequences if left undiagnosed to be sent to the application layer and used in critical actuation situations. This section presents the current research on IoT anomaly detection. IoT-specific anomaly detection is a challenging area, because the solutions must be lightweight and capable of handling the heterogeneous range of IoT devices.

Spanos et al. (2019) proposed a smart-home anomaly detection method which combines statistical and machine learning techniques according to the network behaviour of the device. During training, features are extracted from the network packet data, these features are then standardised and passed into a clustering algorithm. These clustered labels are then passed into ensemble classification methods, which determine the final result from soft-voting. The authors were able to detect mechanical exhaustion and physical damage to the devices. Nevertheless, more data and performance metrics are required here to determine if the model works at scale and with a wider set of devices.

Gonzalez-Vidal et al. (2019) examined methods to detect anomalies in IoT time-series data. Their process consisted of two steps; extract outliers and abnormal patterns using the individual time-series properties of the data, and then use the features extracted from these models to classify them from the annotated classes. For the time series anomaly detection model, the ARIMA and HOT-SAX frameworks were used, while Random Forest and Association Rule Mining methods were used in the classification component. The authors saw accuracies of up to 90% using their methods. This work is a valuable contribution in the area of sensor data-level anomaly detection, however, it is limited in that it requires time-series data to operate.

Stiawan et al. (2017) proposed a technique for early anomaly detection using network traffic analysis. This technique used the SNMP (Simple Network Mapping Protocol) to collect traffic from a heterogeneous range of IoT devices. This traffic was then visualised in graphs for further analyses. Thresholds could then be set based upon CPU and memory usage which can determine the presence of an anomalous communication in the network. This approach is lightweight and suited to the IoT, however, the solution does not include a method to automatically or statistically determine a threshold for failures, which could generate a high volume of false alarms.

Sedjelmaci et al. (2016) proposed an energy-efficient anomaly detection technique which caters for low-resource IoT devices. The technique uses a game theoretic methodology in order to reach the optimal energy efficiency by combining two known techniques for intrusion detection in IoT; signature-based detection and anomaly detection. The anomaly detection component learns activity and builds a classification rule, which is then passed to the signature detection component so that the next time the anomaly occurs it can be recognised by its signature rather than having to rerun the classifier to detect it. Game theory was then applied to this hybrid technique to create further energy savings, which opposes two “players” against each other, one being the attacker launching the new attack signatures and the other running the algorithm to detect anomalous new signatures. When the game finishes the historical data can

be examined to determine the probability of a new signature and thus can state a time at which anomaly detection should be run to build new rules. The study compared the proposed lightweight game-theoretic technique to other known hybrid techniques in the research literature. The study found that accuracy was reduced in the game-theoretic technique, which was to be expected given the predictive nature of the technique. When comparing energy consumption, however, the study found that it was possible to save up to 6000 mJ of energy when running the lightweight technique, which represents a worthwhile energy saving given the low-resource nature of IoT.

Desnitsky et al. (2015) proposed a method for detecting anomalies in IoT applications using domain-specific knowledge to create a list of constraints for the application. For example, the temperature in a home should not exceed 30 degrees Celsius, or the constraints could be drawn from the history of the data, for example a motion sensor in an office stops providing data. If one of these constraints is exceeded this indicates the presence of an anomalous situation. This model is useful for detecting simple anomalous scenarios, however, it is entirely dependent upon the rule base which is designed by the domain-expert. This limitation means that if an anomaly is not accounted for in the constraints then it will not be detected.

Abeshu and Chilamkurti (2018) proposed a deep learning approach for detecting attacks based upon the fog-computing paradigm in IoT. Using the fog-computing paradigm can significantly reduce delays versus the traditional cloud centric paradigm, which is useful in mission-critical IoT scenarios. The study compared the performance of a deep learning model which used a pre-trained stacked autoencoder for feature engineering and SoftMax for classification against a shallow learning model. The study found that the deep model was consistently more accurate than the shallow model, on average this accuracy gap was 4% which is a large gap in a mission-critical application. Furthermore, the study revealed that the deep model coped with a scaling number of nodes much more comfortably than the shallow model, as when the shallow model was exposed to more than 80 fog nodes the accuracy fell by 2%.

Thanigaivelan et al. (2016) proposed an anomaly detection system for IoT where each node monitors the behaviour of its one-hop neighbours. The proposed system has three main components; the MGSS (Metrics and Grading Subsystem), the RSS (Reporting Subsystem) and the ISS (Isolation Subsystem). The MGSS is the component responsible for grading the neighbouring nodes, these nodes are graded based upon packet size and data rate. The RSS is responsible for reporting any nodes which are confirmed to be anomalous, which the ISS component will then isolate to remove the threat from the network. Further research is required within this solution in order to derive a more

comprehensive list of network parameters to monitor, and a statistical method is needed to determine if a node is anomalous or not.

Nomm et al. (2019) proposed a method of detecting botnet attacks in IoT deployments. The method evaluated feature selection techniques to reduce the dimensionality of the data before passing it into a classifier. The dataset used in the experiment was a genuine dataset from a Mirai botnet attack, containing 115 discrete numerical features generated by 9 IoT devices. The features described various network characteristics, such as source and destination IP, jitter and socket information. The author used three different techniques to reduce the dimensionality of the data; entropy, variance and Hopkins statistics. Three classifiers were then used to classify the data; LOF (local outlier factor), one-class SVM (support vector machine) and an IF (Isolation Forest). The study found that feature reduction by entropy combined with the IF classifier was able to achieve accuracy results of 90% by using 5 features. This feature reduction is well suited to the IoT given that it is a much greener approach to machine learning, as opposed to a classifier having to train and test on 115 features. This anomaly detection technique is successfully able to detect anomalies at the network level but does not consider the anomalies that may occur in the payload of the packet being sent by the IoT devices themselves.

The papers reviewed here with regard to IoT anomaly detection represent a clear drive in the research community to create a more reliable IoT ecosystem. With this in mind, it should be stated that anomaly detection is an extremely large field, with application in IoT, network security and a vast array of other computing disciplines. Within the scope of this work, it is not possible to review all available anomaly detection methods, and as such, only the pertinent IoT examples are reviewed here in detail. A more detailed review of anomaly detection methods can be found within the literature (Zarpelão et al. 2017; Moustafa et al. 2019; Cook et al. 2020; da Costa et al. 2019).

Many methods were discussed in this section which provide accurate and varied mechanisms for detecting anomalies in IoT systems. Nevertheless, further research is required to determine how anomalies actually affect the reliability of an IoT system, given that the presence of an anomaly does not necessarily need to hinder or prevent IoT services from operating. This being said, the presence of anomalies is a clear indicator that the IoT system is not performing optimally.

5.5 Discussion of surveyed work

The range of research presented in this section demonstrates a growing demand for quantifying reliability in IoT networks. This is not a straightforward task, given that we must be able to assess reliability at both a device and at a network

Table 1 Summarised literature review and contributions of each work towards reliability assessment

Contribution	Work	Metrics/techniques used	Findings
Device layer reliability	Zin et al. (2016)	R(t), failure rate, expected time to failure, availability, mean time of repair	Several metrics are proposed throughout these works, some of which are standard reliability metrics and others which are non-standard. These metrics are then used as a method of quantifying the state of reliability in these IoT deployments. In general, these approaches lack an awareness of reliability at the network level
	Saini (2016) Mavrogorgou et al. (2018) Li et al. (2012) Sicari et al. (2016), Sicari et al. (2014) Kim (2016)	Reliability over a period of time (ROPT), Trust factor MTTR, MTTF, MTBF, availability Currency, availability, validity Accuracy, precision, timeliness, completeness Reliability, maturity, failure rate, recoverability	
Network layer reliability	Maalel et al. (2013)	Link quality indicator (LQI)	These approaches do not consider the possibility that hardware failures can occur at the device layer in the IoT paradigm. If a failure were to occur at the device level, then the information would not reach the intended target regardless of the network availability. Furthermore, no paradigm is presented, such as HTTL, which would allow a person to intervene in critical situations to prevent total system failure
	Kamyod (2018)	E2E delay, throughput, retransmission attempts	
	Al-Masri (2018)	Response time, throughput, availability, reliability, cost	
	Li and Huang (2017)	Response time, MTTR, MTTF, failure rate	
	Sinche et al. (2018) Brogi and Forti (2017) Alam (2018)	Return trip time (RTT) Latency, bandwidth Delay	
System reliability	Behera et al. (2015)	Centralised heterogeneous IoT service system (CHISS)	Only considers the availability of programs required for execution, and does not consider the possibility of device or network failure
	Kharchenko et al. (2017)	Markov modelling	This approach does not consider reliability for individual components within the IoT paradigm, and is limited to the states specified at the outset. This means that information cannot be gained with regards to where in the system an error may have occurred
Anomaly detection	Stiawan et al. (2017)	Network parameters measured against empirical thresholds	These methods represent the current research towards IoT anomaly detection. While they perform well in terms of accuracy and detection, there is currently a lack of research available which translates this anomaly information into how it may affect the reliability of the IoT deployment
	Sedjelmaci et al. (2016)	Game-theoretic methodology	
	Su (2011)	Rule-based	
	Abeshu and Chilamkurti (2018)	Deep learning	
	Thanigaivelan et al. (2016)	Metric grading subsystem	
	Nomm et al. (2019)	Classification of Botnet attacks	
	Gonzalez-Vidal et al. (2019)	ARIMA, HOT-SAX, RF, ARLA	
	Spanos et al. (2019)	Ensemble classification	

level whilst also being able to detect anomalies as they occur in the system. The research studies presented in this paper all only tackle one facet of the problem, as is evidenced in Table 1 which summarises the contributions of these works. A complete solution would need to be able to integrate all of this valuable IoT reliability information into one reliability framework. The research presented in this paper presents a clear gap in the knowledge and understanding of IoT: there is currently not a solution available capable of, in an end-to-end sense, assessing the reliability of IoT infrastructure.

From the works aimed at quantifying device reliability, there are several different contributions made. Some works, such as Mavrogiorgou et al. (2018), Zin et al. (2016) and Kim (2016) use standard reliability metrics to quantify the state of reliability in IoT devices. These standard metrics include MTTF, MTTR, Availability, Maintainability and Failure Rate (Fries 2006). When given enough device data, these metrics can be used to mathematically reason about the reliability of IoT devices. Some works, such as Saini (2016) and Li et al. (2012) proposed non-standard metrics, like ROPT, Trust Factor and Maturity. Again, these metrics can provide some view of how reliable an IoT device or set of devices is.

The device reliability metrics, regardless of being standard or non-standard, offer up several opportunities for expansion and further research. Firstly, perhaps these metrics could also be extended to include network infrastructure and communications protocols. Doing so would enable the solution to be a more holistic one and bring it closer to managing reliability for the full end-to-end stack. Secondly, these metrics are able to attest to reliability of IoT devices at a certain point in time—could these metrics then be extended to allow the systems to predict and preempt failure? Doing this would be a valuable step towards a more reliable IoT, especially in scenarios where the IoT is supporting mission critical applications. This leads on to the third area for expansion here—while these metrics are valuable at solving reliability for a given set of sensors in a given environment, there is research required to understand how this generalises into other applications. Importantly, do different thresholds need to be applied when considering one IoT vertical over another? Some research is also required to understand how these reliability metrics might react as new and previously unseen devices are added to the applications. One would expect that new devices may carry a significantly different failure profile, and thus may influence the reliability metrics in different ways. The research on IoT device reliability, therefore, should be extended where possible to include the scenario in which the IoT is capable of handling new and unseen devices, operating over a wide range of communication protocols. Lastly, there is an interplay between IoT device reliability and anomaly detection which was

not fully exploited in the works surveyed. Given that we know IoT devices are prone to both spontaneous failure and attack from malicious users, this notion will have a strong influence on the reliability of IoT devices. Therefore, research is required to understand the impact of anomalies on IoT device reliability. For example, some applications may be highly sensitive to noise and anomalies, while other applications may fail completely with the presence of a single anomaly. As such, anomaly detection methods provide a valuable insight into the current state of reliability for IoT devices. A potential research question exists here in trying to understand if reliability information can be synthesised from anomaly detection models.

With regard to the works researching network reliability, again we can observe that some metrics were proposed, both standard (Al-Masri 2018; Li and Huang 2017; Alam 2018) and non-standard (Sinche et al. 2018). We can also observe that some new communication protocols were proposed for enabling a more reliable IoT. Some research was also conducted to help address the need for IoT solutions to be considerate of the various vertical markets, for example emergency IoT applications (Maalel et al. 2013). Methods were also introduced to profile devices before they joined the IoT deployment, using reliability data as the decision factor (Brogi and Forti 2017).

The research conducted on network reliability opens up several areas for future research to enable a more reliable IoT. Firstly, while some research has been conducted to understand the sensitivity of different IoT verticals, there is still a growing need for research in this area to help in understanding the impact that these vertical markets have on reliability engineering in the IoT. Given the large predictions for growth in IoT services, we can only expect demand to increase and diversify in terms of the applications being offered. Therefore, in order to be fully reliable, the IoT must be cognisant of these vertical markets, and measure reliability in a tailored fashion. For example, do faults need to be reported in real-time, such as with emergency applications? Or perhaps we may be able to tolerate faults being reported in larger time windows, such as a day, as with smart home applications.

One of the main issues with the studies aimed at assessing network reliability is that they do not have an awareness of the reliability of the devices themselves. Therefore, it is pertinent that some research is conducted to help tie these two facets together in order to enable reliability across the full IoT stack.

As with device reliability, we can also speculate about the importance of anomalies and intrusions in network traffic. It is important that we understand the impact that these anomalies have on the reliability of a particular application. Moreover, if we are able to leverage intrusion detection methods and anomaly detection methods for networks and use them

to ascertain reliability information then this represents a step towards a more reliable IoT. Also similar to the device case, some research would be pertinent to understand if it were possible to predict faults before they occur at the network level. The ability to perform this prediction would enable IoT architects to preemptively manage failure, resulting in a more reliable IoT—especially in the case of mission-critical IoT applications.

The system reliability modelling works reviewed in this paper were not specific to either the device or network component of the IoT architecture. Nevertheless, the methods in these works are at an early stage of development and lack the complexity required to deal with a complex IoT environment.

Referring to the works reviewed for anomaly detection, it is clear from these works that anomaly detection is a growing field within the IoT and computing in general. While the anomaly detection methods included were capable of detecting anomalies, there is still a lack of research and knowledge on how we might leverage this anomaly information to quantify the reliability of an IoT deployment. A key area of future research here will be to take these anomaly detection methods and to try to synthesise reliability information from them.

6 The five research directions for IoT reliability

Having catalogued and analysed the combined efforts made by the IoT reliability research community, some assessments can be drawn as to what the ideal reliability solution should look like. While none of the works surveyed in this paper fully satisfy end-to-end reliability in IoT, they each add a piece of the puzzle towards this goal. As such, we can derive from these works five crucial elements that an end-to-end reliability management system for the IoT must adhere to.

6.1 Direction 1: Vertical and real-time measurement

If the IoT is set to manage critical infrastructure, such as security and critical traffic systems, then we must be able to attest to the reliability of the system in real-time, or as close to real-time as possible. As shown in the study by Maalel et al. (2013), it is necessary that we pay particular attention to those applications which operate emergency services and require a rapid and reliable response. Moreover, there is a need to define reliability requirements in each individual domain. For example, a smart-building solution may have a delay tolerance of up to a few seconds. An industrial process, on the other hand, will likely only be able to tolerate delays of microseconds. As such, research is required to categorise

these requirements and design effective solutions to handle reliability in each of these vertical domains.

6.2 Direction 2: All devices, all protocols

This survey has demonstrated the very wide array of protocols and devices which are set to connect to and consume services from the IoT. Standards for communication protocols are continuing to evolve daily with efforts from many research groups aiming to design more lightweight and efficient communication protocols. Moreover, new IoT devices and hardware continue to emerge in the consumer market daily. Therefore, the ideal reliability solution must be both hardware, software and communication protocol agnostic.

6.3 Direction 3: Full stack awareness

One of the conclusions drawn from the literature review was that, while many researchers had successfully solved a particular problem, or subset of problems, in IoT reliability research, no study has been undertaken which had full awareness of end-to-end reliability. Given the scale and complexity of emerging IoT deployments, this is no easy task. This is not to say, however, that researchers should aim to design a “one size fits all” reliability approach, as this would contradict the first research direction outlined in this work. Rather, individual reliability solutions should be proposed for each IoT vertical that encompass the full IoT architecture. Nevertheless, designing an end-to-end reliability solution for the IoT would be a significant and novel research finding with the potential to greatly enhance IoT end-user experience.

6.4 Direction 4: Synthesising reliability information from anomalies

Much work has gone into detecting and reporting anomalies when they appear in IoT services. While this work is both useful and necessary, it does not necessarily aid reliability without an extra step. Knowledge of an anomaly does not necessarily tell the user if the IoT system has become less reliable. Therefore, there is a need to research how we can synthesise information about emergent anomalies in IoT systems into information on how the reliability has been affected. For example, if a sensor breaks in a smart home which is monitoring an assisted living scenario, there may not necessarily be an immediate risk to life. Whereas, if a thermal sensor begins sending erroneous readings in a smart factory, there is potential for dangerous machinery to malfunction.

6.5 Direction 5: Predict and preemptively manage failure

Measuring reliability is the task discussed at length in this work. If the research is to move a step beyond this goal, then the task of predictive maintenance can be considered. If we are able to reason about the quantified reliability of a system, can we then extrapolate this into an accurate maintenance date? Moreover, can this be further classified at a component level and be a dynamic process which determines results based on real-time reliability data, rather than using a history of past failures to estimate a future failure date? Solving this research question would represent a valuable step in the research of IoT reliability.

7 Conclusion

To the best of our knowledge, this study represents the first review or survey studying the topic of IoT reliability. A detailed history of the evolution of reliability was given, starting from the fundamentals of reliability engineering, moving into reliability in computing and then finally a detailed discussion on the arena of IoT-specific reliability. IoT reliability was defined and discussed across the four main layers of the architecture. A detailed literature review was presented, which looked at research in device, network and system reliability, while also reviewing the current state of the art anomaly detection methods for the IoT. Lastly, the findings and outputs of this detailed survey have been used to formulate five key research directions for the area of reliability in the Internet of Things. This finding now presents a need for the IoT research community to design and implement solutions according to the directions identified in this paper. These solutions will serve to strengthen and support the reliability of our IoT infrastructure, resulting in a safer and more stable paradigm for its users.

Acknowledgements This research is supported by the BTIC (British Telecom Ireland Innovation Centre) project, funded by BT and Invest Northern Ireland.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Abeshu, A., Chilamkurti, N.: Deep learning: the Frontier for distributed attack detection in fog-to-things computing. *IEEE Commun. Mag.* **56**(2), 169–175 (2018)
- Ahmed, I., Saleel, A.P., Beheshti, B., Khan, Z.A., Ahmad, I.: Security in the Internet of Things (IoT). In: 2017 Fourth HCT Information Technology Trends (ITT), pp. 84–90. IEEE, New York (2017)
- Alaba, F.A., Othman, M., Hashem, I.A.T., Alotaibi, F.: Internet of Things security: a survey. *J. Netw. Comput. Appl.* **88**(March), 10–28 (2017)
- Alam, T.: A reliable communication framework and its use in Internet of Things (IoT). *Int. J. Sci. Res. Comput. Sci. Eng. Inf. Technol.* **5**(10), 450–456 (2018)
- Al-Fuqaha, A., Guizani, M., Mohammadi, M., Aledhari, M., Ayyash, M.: Internet of Things: a survey on enabling technologies, protocols, and applications. *IEEE Commun. Surv. Tutor.* **17**(4), 2347–2376 (2015)
- Allhoff, F., Henschke, A.: The Internet of Things: foundational ethical issues. *Internet Things* **1–2**, 55–66 (2018)
- Al-Masri, E.: QoS-aware IIoT microservices architecture. In: 2018 IEEE International Conference on Industrial Internet (ICII), pp. 171–172 (2018)
- Atzori, L., Iera, A., Morabito, G.: Understanding the Internet of Things: definition, potentials, and societal role of a fast evolving paradigm. *Ad Hoc Netw.* **56**, 122–140 (2017)
- Behera, R.K., Reddy, K.H.K., Roy, D.S.: Reliability modelling of service oriented Internet of Things. In: 2015 4th International Conference on Reliability, Infocom Technologies and Optimization: Trends and Future Directions, ICRITO 2015, pp. 1–6 (2015)
- Bradley, E.: Reliability Engineering: A Life Cycle Approach, 1st edn. CRC Press, Boca Raton (2016)
- Broggi, A., Forti, S.: QoS-aware deployment of IoT applications through the fog. *IEEE Internet Things J.* **4**(5), 1–8 (2017)
- Chiang, M., Zhang, T.: Fog and IoT: an overview of research opportunities. *IEEE Internet Things J.* **3**(6), 854–864 (2016)
- Cook, A., Misirli, G., Fan, Z.: Anomaly detection for IoT time-series data: a survey. *IEEE Internet Things J.* **1**, 7(7), 6481–6494 (2020)
- da Costa, K.A.P., Papa, J.P., Lisboa, C.O., Munoz, R., de Albuquerque, V.H.C.: Internet of Things: a survey on machine learning-based intrusion detection approaches. *Comput. Netw.* **151**, 147–157 (2019)
- Da Li, X., He, W., Li, S.: Internet of Things in industries: a survey. *IEEE Trans. Ind. Inform.* **10**(4), 2233–2243 (2014)
- DDCMS: Code of practice for consumer IoT security (2018). https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/773867/Code_of_Practice_for_Consumer_IoT_Security_October_2018.pdf
- Desnitsky, V.A., Kotenko, I.V., Nogin, S.B.: Detection of anomalies in data for monitoring of security components in the Internet of Things. In: Proceedings of International Conference on Soft Computing and Measurements, SCM 2015, pp. 189–192 (2015)
- Evans, D.: The Internet of Things—how the next evolution of the internet is changing everything. CISCO white paper, pp 1–11 (2011)
- Fekade, B., Maksymyuk, T., Kyryk, M., Jo, M.: Probabilistic recovery of incomplete sensed data in IoT. *IEEE Internet Things J.* **5**(4), 2282–2292 (2017)
- Fries, R.C.: Reliable Design of Medical Devices, 2nd edn, 3rd edn, Number 2. Taylor and Francis Group/CRC Press, London (2006)
- Ghorbani, H.R., Ahmadzadegan, M.H.: Security challenges in Internet of Things: survey. In: 2017 IEEE Conference on Wireless Sensors (ICWiSe), pp. 1–6. IEEE, New York (2017)
- Gonzalez-Vidal, A., Cuenca-Jara, J., Skarmeta, A.F.: IoT for water management: towards intelligent anomaly detection. In: IEEE 5th

- World Forum on Internet of Things, WF-IoT 2019—Conference Proceedings, pp. 858–863 (2019)
- ISO, P., 9000: 2015 Quality management systems-Fundamentals and vocabulary. International Organization for Standardization (ISO), Geneva: ISO (2015)
- Kamyod, C.: End-to-end reliability analysis of an IoT based smart agriculture. In: 3rd International Conference on Digital Arts, Media and Technology, ICDAMT 2018, pp. 258–261 (2018)
- Karkouch, A., Mousannif, H., Al Moatassime, H., Noel, T.: Data quality in Internet of Things: a state-of-the-art survey. *J. Netw. Comput. Appl.* **73**, 57–81 (2016)
- Karkouch, A., Mousannif, H., Al Moatassime, H., Noel, T.: A model-driven architecture-based data quality management framework for the Internet of Things. In: Proceedings of 2016 International Conference on Cloud Computing Technologies and Applications, CloudTech 2016, pp. 252–259 (2017)
- Kharchenko, V., Kolisnyk, M., Piskachova, I., Bardis, N.: Reliability and security issues for IoT-based smart business center: architecture and Markov model. In: Proceedings—2016 3rd International Conference on Mathematics and Computers in Sciences and in Industry, MCSI 2016, pp. 313–318 (2017)
- Kim, M.: A quality model for evaluating IoT applications. *Int. J. Comput. Electr. Eng.* **8**(1), 66–76 (2016)
- Kouicem, D.E., Bouabdallah, A., Lakhlef, H.: Internet of Things security: a top-down survey. *Comput. Netw.* **141**, 199–221 (2018)
- Kushalnagar N., Gabriel M., Christian S.: "IPv6 over low-power wireless personal area networks (6LoWPANs): overview, assumptions, problem statement, and goals." 1–11 (2007)
- Li, S., Huang, J.: GSPN-based reliability-aware performance evaluation of IoT services. In: Proceedings—2017 IEEE 14th International Conference on Services Computing, SCC 2017, pp. 483–486 (2017)
- Li, F., Nastic, S., Dustdar, S.: Data quality observation in pervasive environments. In: Proceedings—15th IEEE International Conference on Computational Science and Engineering, CSE 2012 and 10th IEEE/IFIP International Conference on Embedded and Ubiquitous Computing, EUC 2012, pp. 602–609 (2012)
- Maalel, N., Natalizio, E., Bouabdallah, A., Roux, P., Kellil, M.: Reliability for emergency applications in Internet of Things. In: Proceedings—IEEE International Conference on Distributed Computing in Sensor Systems, DCoSS 2013, pp. 361–366 (2013)
- Mavrogiorgou, A., Kiourtis, A., Symvoulidis, C., Kyriazis, D.: Capturing the reliability of unknown devices in the IoT world. In: 2018 5th International Conference on Internet of Things: Systems, Management and Security, IoTSMS 2018, pp. 62–69 (2018)
- Moore, S.J., Nugent, C.D., Cleland, I., Zhang, S.: Impact analysis of erroneous data on IoT reliability. In: Proceedings of the 2019 IEEE SmartWorld Smart City Innovation Conference, pp. 1908–1915 (2019)
- Moustafa, N., Hu, J., Slay, J.: A holistic review of Network Anomaly Detection Systems: a comprehensive survey. *J. Netw. Comput. Appl.* **128**(October 2018), 33–55 (2019)
- Nomm, S., Bahsi, H.: Unsupervised anomaly based botnet detection in IoT networks. In: Proceedings—17th IEEE International Conference on Machine Learning and Applications, ICMLA 2018, pp. 1048–1053 (2019)
- Rafferty, J., Synnott, J., Nugent, C.D., Ennis, A., Catherwood, P.A., Mcchesney, I., Cleland, I., Mcclean, S.: A scalable, research oriented, generic, sensor data platform. *IEEE Access* **6**, 45473–45484 (2018)
- Rayes, A., Salam, S.: Internet of Things-from Hype to Reality: The Road to Digitization, 1st edn. Springer, Cham (2016)
- Rizzardi, A., Miorandi, D., Sicari, S., Cappiello, C., Coen-Porisini, A.: Networked smart objects: moving data processing closer to the source. *Lect. Notes Inst. Comput. Sci. Soc. Inform. Telecommun. Eng. LNICST* **170**, 28–35 (2016)
- Saini, N.K.: Trust factor and reliability-over-a-period-of-time as key differentiators in IoT enabled services. In: 2016 International Conference on Internet of Things and Applications, IOTA 2016, pp. 411–414 (2016)
- Sato, H., Kanai, A., Tanimoto, S., Kobayashi, T.: Establishing trust in the emerging era of IoT. In: Proceedings—2016 IEEE Symposium on Service-Oriented System Engineering, SOSE 2016, pp. 398–406 (2016)
- Sedjelmaci, H., Senouci, S.M., Al-Bahri, M.: A lightweight anomaly detection technique for low-resource IoT devices: a game-theoretic methodology. In: 2016 IEEE International Conference on Communications, ICC 2016, pp. 1–6 (2016)
- Shi, W., Cao, J., Zhang, Q., Li, Y., Lanyu, X.: Edge computing: vision and challenges. *IEEE Internet Things J.* **3**(5), 637–646 (2016)
- Sicari, S., Cappiello, C., De Pellegrini, F., Miorandi, D., Coen-Porisini, A.: A security-and quality-aware system architecture for Internet of Things. *Inf Syst Front.* **18**(4) 665–677 (2014)
- Sicari, S., Rizzardi, A., Miorandi, D., Cappiello, C., Coen-Porisini, A.: A secure and quality-aware prototypical architecture for the Internet of Things. *Inf. Syst.* **58**, 43–55 (2016)
- Sinche, S., Polo, O., Raposo, D., Fernandes, M., Boavida, F., Rodrigues, A., Pereira, V., Sa Silva, J.: Assessing redundancy models for IoT reliability. In: 19th IEEE International Symposium on a World of Wireless, Mobile and Multimedia Networks, WoWMoM 2018, pp. 14–15 (2018)
- Singh, D., Tripathi, G., Jara, A.J.: A survey of Internet-of-Things: future vision, architecture, challenges and services. In: 2014 IEEE World Forum on Internet of Things, WF-IoT 2014, pp. 287–292 (2014)
- Spanos, G., Giannoutakis, K.M., Votis, K., Tzovaras, D.: Combining statistical and machine learning techniques in IoT anomaly detection for smart homes. In: IEEE International Workshop on Computer Aided Modeling and Design of Communication Links and Networks, CAMAD, 2019 Sept., pp. 1–6 (2019)
- Stankovic, J.A.: Research directions for the Internet of Things. *IEEE Internet Things J.* **1**(c), 3–9 (2014)
- Stiawan, D., Idris, M.Y., Malik, R.F., Nurmaini, S., Budiarto, R.: Anomaly detection and monitoring in Internet of Things communication. In: Proceedings of 2016 8th International Conference on Information Technology and Electrical Engineering: Empowering Technology for Better Future, ICITEE 2016, pp. 1–4 (2017)
- Su, M.Y.: Real-time anomaly detection systems for Denial-of-Service attacks by weighted k-nearest-neighbor classifiers. *Expert Syst. Appl.* **38**(4), 3492–3498 (2011)
- Thanigaivelan, N.K., Nigussie, E., Kanth, R.K., Virtanen, S., Isoaho, J.: Distributed internal anomaly detection system for Internet-of-Things. In: 2016 13th IEEE Annual Consumer Communications and Networking Conference, CCNC 2016, pp. 319–320 (2016)
- Tsai, C.W., Lai, C.F., Vasilakos, A.V.: Future Internet of Things: open issues and challenges. *Wirel. Netw.* **20**(8), 2201–2217 (2014)
- Wang, Y.: Trust quantification for networked cyber-physical systems. *IEEE Internet Things J.* **5**(3), 2055–2070 (2018)
- Xie, M., Dai, Y.-S., Poh, K.-L.: Computing system reliability: models and analysis. Springer Science & Business Media (2004)
- Yaqoob, I., Ahmed, E., Hashem, I.A.T., Ahmed, A.I.A., Gani, A., Imran, M., Guizani, M.: Internet of things architecture: Recent advances, taxonomy, requirements, and open challenges. *IEEE wireless commun.* **24**(3), 10–16 (2017)
- Zarpelão, B.B., Miani, R.S., Kawakani, C.T., de Alvarenga, S.C.: A survey of intrusion detection in Internet of Things. *J. Netw. Comput. Appl.* **84**(February), 25–37 (2017)
- Zin, T.T., Tin, P., Hama, H.: Reliability and availability measures for Internet of Things consumer world perspectives. In: 2016 IEEE 5th Global Conference on Consumer Electronics, GCCE 2016, pp. 1–2 (2016)

Samuel J. Moore received the B.Sc. degree in computing systems from Ulster University. From 2012 to 2018 he worked as a technology consultant and data analytics specialist in PwC and Citigroup respectively. In 2018 he began research toward a D.Phil. degree in the School of Computing in Ulster University as a part of the British Telecom Ireland Innovation Centre (BTIIC). His research areas primarily concern the IoT, cybersecurity, reliability and artificial intelligence.

Ian Cleland received the B.Sc. degree in Biomedical Engineering and the PhD degree from Ulster University, U.K. He is currently a Lecturer in Data Analytics, within the School of Computing at Ulster University. His research focuses on the development and evaluation of novel healthcare technologies that incorporate concepts from pervasive computing, biomedical engineering and behavioural science.

Chris D. Nugent received the B.Eng. degree in electronic systems and the D.Phil. degree in biomedical engineering from Ulster University. He is currently a Professor of biomedical engineering in the School of Computing, Ulster University. His research interests include data analytics for smart environments and the design and evaluation of pervasive and mobile solutions within the context of ambient-assisted living.

Shuai Zhang is a Lecturer in Computing Science in the School of Computing. Her research interests are in data analytics in the area of Connected Health applications including activity and behavioural recognition in smart environment, change point detection for sensor data annotation, and modelling user engagement with and adoption of assistive technologies for people with dementia, more recently on commercial customer and services analytics. She is a grant holder for projects

funded from H2020, the Royal Society, Wellcome Trust-Wolfson and Industrial funding.