Editorial

Andrzej M. Goscinski*

School of Information Technology,
Deakin University,
Deakin Waurn Ponds Campus,
Geelong, Melbourne, VIC, 3217, Australia
Email: ang@deakin.edu.au
*Corresponding author

Dennis Gannon

School of Informatics, Computing and Engineering,
Indiana University,
9951 Lake Washington BVD NE #25, Belview, WA, 98004, USA
Email: dennis.gannon@outlook.com

Hong Zhu

School of Engineering, Computing and Mathematics,
Oxford Brookes University,
Wheatley Campus, Oxford, OX33 1HX, UK
Email: hzhu@brookes.ac.uk

Shangguang Wang

State Key Laboratory of Networking and Switching Technology,
Beijing University of Posts and Telecommunications,
187#, 10 Xi Tu Cheng Rd., Beijing, 100876, China
Email: sgwang@bupt.eud.cn

Biographical notes: Andrzej M. Goscinski is a Professor Emeritus of Computing from the School of Computing at the Deakin University where he was on the faculty from 1992. He also currently serves as an Adjunct Professor at the RMIT University. Before joining Deakin, he was with the Australian Defence Academy – University of New South Wales. He is recognised as one of the leading researchers in distributed systems, in particular cloud, fog, and edge computing supporting IoT environments. The results of his research have been published in international refereed journals and conference proceedings and presented at specialised conferences.

Dennis Gannon is a Professor Emeritus from the School of Informatics, Computing and Engineering at the Indiana University where he was on the faculty from 1985 to 2008. From 2008 until he retired in 2015, he was with the Microsoft Research in Redmond, Washington. Prior to 1985, he was with the Faculty of Computer Science at the Purdue University. He received his PhD in Computer Science from the University of Illinois, Urbana in 1980 and PhD in Mathematics from the University of California, Davis in 1976.

Hong Zhu is a Professor of Computer Science from the Oxford Brookes University, UK, where he also chairs the Applied Formal Methods Research Group. He obtained his BSc, MSc and PhD in Computer Science from the Nanjing University, China, in 1982, 1984 and 1987, respectively. He was a faculty member of Nanjing University from 1987 to 1998. He joined Oxford Brookes University in November 1998 as a Senior Lecturer in Computing and became a Professor in October 2004. His research interests are in the area of software engineering for cloud computing, including software design, modelling and testing.

Shangguang Wang is a Professor of Computer Science from the Beijing University of Posts and Telecommunications, China, where he is also the Vice Director of the State Key Laboratory of Networking and Switching Technology. He received his PhD degree from the Beijing University of Posts and Telecommunications in 2011. He was a Postdoctoral Fellow from the Beijing Jiaotong University before he joined Beijing University of Posts and Telecommunications.

In recent years, researchers and practitioners have been making progress within the area of edge, fog, and cloud computing, and separately in internet of things (IoT), addressing open problems in individual areas and with the cloud computing as engines that connect them. This special issue provides look at the state of the art for these areas of research.

The concept of the IoT has been in development and use for nearly four decades, and its name is known since the end of the 1990s. At the beginning, they were formed of simple sensors, regulators, and actuators to collect data and generate control signals, and communicated using simple P-2-P networks. In the last years, data are collected from many devices, some only powered by very small batteries and difficult to access, other, requiring a lot of energy to operate and generating a lot of data. These data are sometimes stored and processed by small pad-size computers or smart phones, other by large computers, high-performance computers, and HPC clouds. Data are collected in remote locations, in many cases not easily accessible, sometimes in areas far away from where they could be pre-processed, stored and delivered to other places that could provide means to process them to extract useful knowledge. Examples of environments that collect, pre-process, store, and process data are parking spaces, driverless cars, security cameras, hospitals, metallurgical plants, gas and oil discovery areas, people, company divisions, cities, multi-continental companies. Data are transmitted using wireless and wire-based media and moved between interested parties through routers. All these people, computers, storage, networks, and network devices form the IoT. These IoT are supported by clouds, fogs, and edges. Garner estimates that 25 billion things will be connected to the internet by 2020.

The problems are: where collected data should be stored, and processed, what thing should be a sender of data, and what thing should be a receiver to minimise communication delay and processing costs, how to increase reliability and security, provide data protection and confidentiality, and how to increase user satisfaction. Thus, new solutions are needed urgently to reach this number and achieve the IoT of this size.

Some answers to these questions could be found in this special issue. All papers printed in this special issue are based on selected articles presented at the SERVICES 2018 conference sessions, in particular of the EDGE 2018 Conference. These conference papers have been substantially revised, expanded, and rewritten so that they form new, original work. Although this is a matter of judgment, it was carried out based on a comparison of each submitted article with the original conference paper. All papers have been put through a double-blind peer review process.

The selected for this special issue papers concentrate on future cloud and fog computing infrastructures, fog-based brokers, brokers for heterogeneous IoT applications, transparent IoT for heterogeneous applications, and edge computing architectures.

The paper entitled ‘Docker-pi: Docker container deployment in fog computing infrastructures’ by Ahmed and Pierre enhances the statement that there is a trend to exploit single-board devices in the development of future fog computing infrastructures and use them for handling cloud-like types of workloads. The main bottleneck in fog computing nodes is created by their hardware limitations. Furthermore, there are trials to employ Docker containers in fog computing infrastructures for IoT applications. A good example of these efforts is use of limited devices such as Raspberry PIs to build IoT clouds. The authors, rather than assuming that the Docker container images were already cached in the local nodes, focused on downloading and installation of the Docker container deployment and proposed a technique to speed up these processes. They proposed and evaluated three optimisation techniques to speed up these processes: sequential image layer downloading, multi-threaded layer decompression, and I/O pipelining. Each optimisation addressed a different issue in the standard Docker container deployment. Finally, they combined them together, which led to significant performance improvement. This approach eliminated the delays that take place during container deployment.

Bolettieri and Bruno in their paper entitled ‘Edge-centric resource allocation for heterogenous IoT applications using a CoAP-based broker’ enhance the claim that IoT devices grow in numbers. There is a need for connecting these devices; therefore, one must remember that they are heterogeneous, cloud-based IoT architecture network bandwidth and communication latency can be a severe bottleneck and depend on small amounts of energy. In many cases they are in far away and difficult to access places. The problem is how to allocate IoT devices transparently and effectively among heterogenous applications. To solve this problem the author(s) designed a fog-based broker, specifically, an optimisation framework to determine the notification periods that maximise the applications’ QoS satisfaction under network-related constraints. The authors also proposed algorithms that leverage measurements of the degree of reliability of application transmissions to infer the congestion level of the IoT resources and adapt the notification periods accordingly. Since the authors carried out experimental computer science research, they have developed a software prototype of the broker using the standard features of the CoAP protocol. They rightly validated the achieved solution using simulations and real experiments in an IoT testbed. The authors demonstrated that as the application demands increase, the proposed approach guarantees better QoS satisfaction, higher throughput and improved energy efficiency than a conventional CoAP proxy.

Brasilino, Marri, Shroyer, Pilachowski, Kissel and Swany in the paper entitled ‘In-network processing for edge computing with InLocus’ shows a design of a fog-based broker for heterogenous IoT applications. According to the author the design is to be able to regulate the access to IoT resources transparently and effectively. The author presented literature study widely and deeply. As a follow up, he formulated the problem precisely.

The proposed model was defined adequately. Applications target IoT devices. The author stated a general optimisation to determine polling rates that should be allocated by the broker to each IoT device to maximise the overall QoS satisfaction of the active applications. This approach was substantiated precisely. The author used this as a basis of the development of a general optimisation model for resource brokering, which he used to build a practical algorithm for IoT devices brokering. The whole building process was carried out correctly; including design rationale, determination of transaction loss ratio, congestion detection, algorithm construction, and tuning of the polling rates. The author addressed enabling technologies proposed to be used and be used for the implementation stage of their research. Two types of experimentation were employed to demonstrate feasibility and quality of the solution: simulation and testbed experiments. The author carried them out properly. Conclusions were written well, and future work presented.

The goal of the ‘SMIoT: a software architecture for maintainable internet-of-things applications’ paper by Bohé, Willocx and Naessens is the presentation of a new approach that proposes a paradigm shift from sensor-centric towards application-centric IoT ecosystem design. For this purpose, the authors proposed new architecture, the so called SMIoT architecture and demonstrated that it is possible to build applications independently of sensors and actuators, IoT devices. The architecture allows developers of IoT systems to focus on the business logic by abstracting low level IoT protocols, and communication and security mechanisms. As the outcome, the proposed architecture supports the development of complex and maintainable IoT applications; this is the main contribution of the paper. The paper also contains two more contributions; first, the proposed principles were used in the implementation of an IoT system in an Android framework. The authors used the implementation to demonstrate the flexibility of the proposed architecture. Second, the introduced SMIoT architecture was validated through the design and development of a care home ecosystem. These three major components were prepared based on research carried out using methodology of experimental computer science, which were executed correctly. The reader will notice that the presented problem was described well, and the following it goal was defined precisely. Literature study is wide and deep and reflects the current state of research in the area of interest.

The goal of the paper by Huang entitled ‘The edge architecture for semi-autonomous industrial robotic inspection systems’ by Filho, Yu, Huang, Venkataramana, El-Messidi, Sharber, Westerheide and Alkadi is to present the edge computing architecture of a model-driven robotic UAV system used for industrial inspections, showing its use in different scenarios. The area of interest of the author of this submission is semi-automated inspection of industrial assets; semi-automated inspection means that both robots and human supervisors cooperate when performing time consuming, dangerous, difficult to execute jobs. Since human supervisors exploit measuring, sensing and photographing, there are IoTs that form an interface between them and industrial assets. However, communication, in practice network connectivity, and device power for the purpose of robotic inspection in industrial space are limited. To demonstrate that the goal was achieved, and the outcome of research and development is satisfactory, the author presents the requirements of the system in the context of industrial inspection scenarios. He also presented how to address these requirements using a system that combines aerial unmanned autonomous vehicles (UAVs), base station, cloud services, and human operators. This was followed by a description of the basic interfaces between major system components, including design trade-offs and implementation. The author validated the approach in two industrial applications: the detection of methane leaks in industrial oil and gas facilities, and the inspection of flare stacks in oil refineries.