Modelling and Simulation Challenges in Internet of Things

Gabor Kecskemeti
Liverpool John Moores University

Giuliano Casale
Imperial College London

Devki Nandan Jha
Newcastle University

Justin Lyon
Simudyne

Rajiv Ranjan
Newcastle University

With the rise of Internet of Things (IoT) technology, it is anticipated that large-scale sensor-based systems will permeate society, calling for novel methodologies to design, test, and operate these systems. IoT relies on networked, interconnected physical devices that often feature computational capabilities.[1] The sheer number of these interconnected devices plays a key role in the IoT revolution. For example, Gartner research predicts that IoT will connect up to 50 to 100 billion devices by 2020.[2] It is estimated that IoT will generate ~1.7 trillion US dollars by this time, with an approximate growth rate of 20% year over year.[3]

IoT devices are not only engaged in sensing, but they also perform some actions based on the sensed data and/or external queries. Because of this new vision, these devices are already the backbone of emerging applications, such as smart buildings, smart cities, smart vehicles, environmental sensing and forecasting, and disaster management, among others.

To better understand the potential of this revolution, recognize that most of IoT data collected today is not used, and data that is used is not yet fully exploited. For instance, less than one percent of data is utilized today. This means that 99 percent of IoT data is lost, either because it is not captured or it is captured but not analyzed or used for business analytics.[4] Most data that is used on factory floors, finds application in real-time control or anomaly detection, to send alarms when the sensor detects something out of tolerance. A great deal of additional value remains to be captured by using the data for predictive maintenance or to optimize operational processes.

Typical IoT Data Analytics Platform

Figure 1 shows the architecture of a typical IoT Data Analytics Platform (IoTDAP), which starts with raw data collection from sensing devices and ends with complex data analytics and decision making activities. An IoTDAP requires timely and intelligent coordination of data and control flows between IoT devices, IoT gateways and in-transit network devices in the Edge Datacenter (EDC), with the big data programming models and virtualized hardware resources hosted in large Cloud Datacenter (CDC)

As noted in the previous installment of “Blue Skies” and shown in Figure 1, IoT devices can be

 2325-6095/17/$33.00

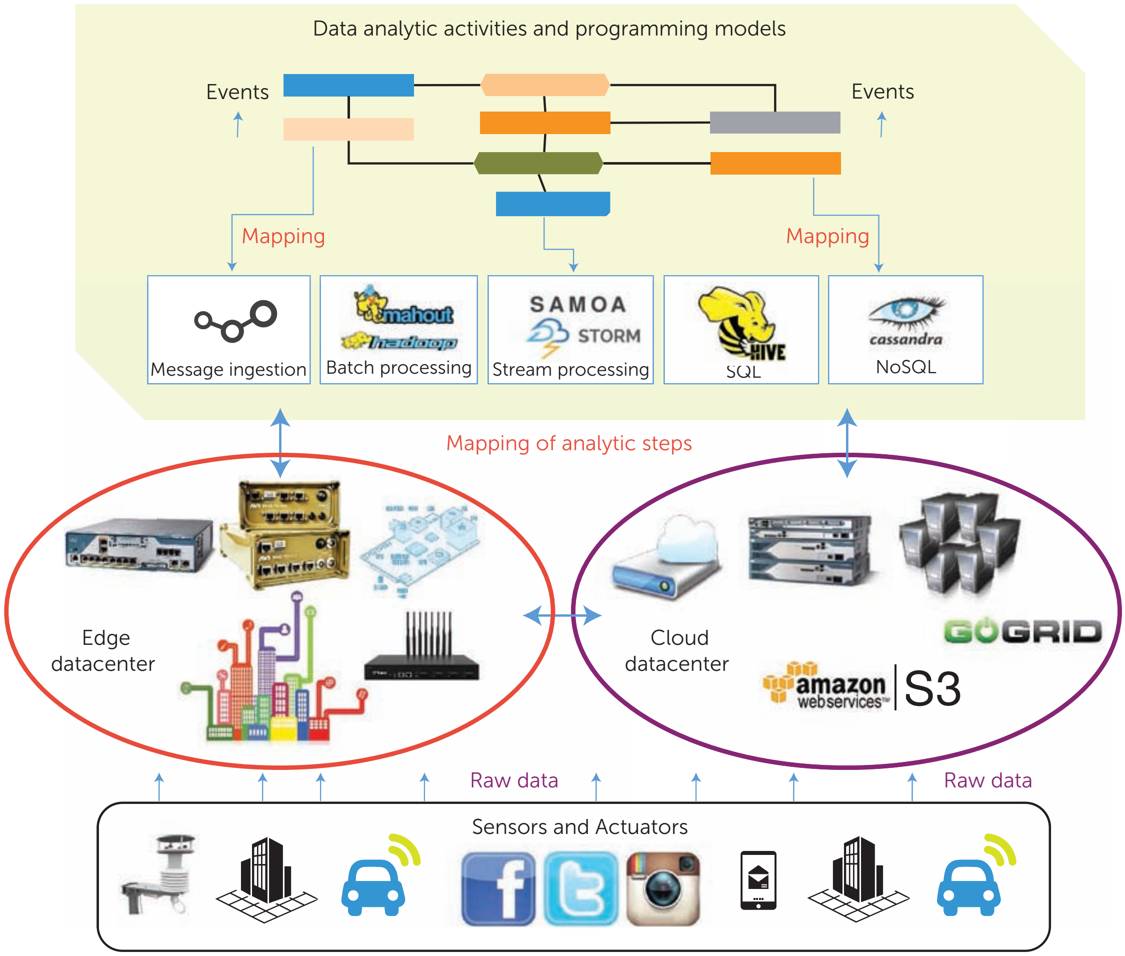

FIGURE 1. Typical Internet of Things (IOT) data analytics platform consisting of sensors, actuators, edge, and cloud datacenters.

sensors, mobile phones, radio frequency identification, actuators (such as machines/equipment fitted with sensors and deployed for mining, oil exploration, or manufacturing operations), lab instruments, and smart consumer appliances (TV, phone, and so on).[5,6] Social media, clickstreams, and business transactions are also workloads in IoT.

Next, EDC can be defined as a "collection of heterogeneous resources including smart IoT devices, IoT gateways (e.g., raspberry pi 3, UDOO board, esp8266, etc.), and Software Defined Networking (SDN) and Network Function Virtualisation (NFV) devices (e.g., Cisco IOx, Hewlett Packard (HP) OpenFlow and Middlebox Technologies) at the network edge that can offer computing and storage capabilities on a pervasive—yet much smaller—scale than CDCs. The scope and role of each resource in an EDC differs. For example, IoT gateways collect, aggregate, and process the data generated by the sensing devices. IoT gateway accepts and routes commands sent from the backend to the respective device. It is also responsible for authenticating and authorizing the devices to participate in IoTDAP. It ensures secure communication between the devices and the centralized command center. The gateway is also capable of dealing with multiple protocols (e.g., Constrained Application Protocol, MQ Telemetry Transport) and data formats. Finally, in-transit SDN and NFV devices offer useful solutions for supporting in-network/in-transit data processing (between edge and CDC) and providing network management abstraction independent of the underlying technology.

Meanwhile, the massive data storage and processing activities (data mining and big data analytics) are performed using complex big data programming models (e.g., message ingestion, batch processing, streaming, Structured Query Language (SQL), NoSQL) in CDCs. CDC exploit virtualization (both hypervisor and container-based) and platform services to elastically scale up/down storage and processing capabilities.

Compared to traditional modelling and simulation frameworks, we envision that simulation of large-scale IoTDAP will require answering complex questions.

It is widely accepted that every bit of the data generated by every IoT sensing devices need not make its way back to the CDC. For some data, it might make sense to collect, store, interpret, and respond to locally in the EDC. But IoTDAP stakeholders need a strategy about which data needs to be processed in EDC and which data reaches all the way up to the CDC. For example, rather than sending all IoT sensor data to the CDC, part of the data filtering and aggregation could happen on EDC which may then send a trigger to CDC for initialising large-scale data analytics required for decision making.

Research Challenges

Due to deep complex intertwining among different data sources, data analytic activities, EDC, CDC, the design, implementation, and testing of these new IoTDAP face many challenges that arise from the rapid increase of the number and types of sensing devices, type, and scope of resources in EDC and CDC, type and scope of big data programming models, functional complexity, data formats heterogeneity, the increasing deployment of distributed and networked architecture.

From an IoT application designer's perspective, IoTDAP are challenging to understand during the early stages of development, due to the dependencies between devices, EDC, and CDC. From an operator's perspective, deploying and testing a system in a real environment is a costly and time-consuming task, which therefore should be carried out parsimoniously, after gaining some confidence that the platform will provide appropriate performance. These challenges raise stringent requirements for developing new holistic simulation and modelling frameworks for aiding design, development, and testing of performance, safety, security, reliability, energy-efficiency, and fault tolerance related to IoTDAP.

Compared to traditional modelling and simulation frameworks, we envision that simulation of large-scale IoTDAP will require answering complex questions concerning, for example, the state of the environment surrounding the sensors and feedback-loop control of a large population of sensors. Moreover, IoTDAP simulation environments should be able to capture the interdependence between sensing, control, actuation, and computational logic in CDC and at the EDC. We identify some key challenges that need to be further investigated within the realm of IoTDAP modelling and simulation research, namely:

Scalability of the simulation fabric. In parallel with the increase in the number of IoT devices on the world, it is likely that the number of devices to be considered in a single IoTDAP will rise to unprecedented levels as well. To enable the design and rapid prototyping of these systems, it is expected from future IoT oriented simulators to model the behavior of tens of millions of sensors, actuators and their corresponding computing elements in edge and cloud computing facilities in a timely fashion. The sheer number of the simulated devices already points towards the requirement of distributed and parallel solutions behind simulation frameworks. On the other hand, to enable efficient operation of IoTDAP systems, we envision that in the foreseeable future IoT simulators could be exploited in online decision making about operation concerns. Under these circumstances, the deadline constrained simulators are expected to not only scale up to meet the deadlines but also—if needed—adjust their accuracy (i.e., to keep costs and performance at balance) according to the mission-criticalness of the decision to be taken.

Elastically modelling sensor and actuator population. Companies and operators of such large scale IoTDAP would often face unexpected turns of events like sudden unavailability of sensing devices in the ambient environment due to cascading power failures and loss of data in sensor and communication networks. Hence, the modellers need the ability to conduct various "what if" scenarios to understand how population and density of sensing devices affects resiliency of IoTDAP and the accuracy of decision making. It is therefore increasingly important that simulators offer the ability to elastically model populations of sensors, including "birth" and "death" of devices, by explicitly representing the influence of the environment on the sensors themselves.

Modelling device heterogeneity. Each IoTDAP use case (e.g., smart cities, smart vehicles, environmental sensing and forecasting, and disaster management) requires heterogeneous combination of sensing and actuation devices. Hence the modelling and simulation framework must be able to accommodate diverse sensor and actuator types (and resulting data types)—from sensing temperature and pressure to location data, and to streaming live videos for video analytics.

Modelling heterogeneous CDC and EDC resource abstractions. The simulation and modelling framework needs the ability to model performance of diverse types of computing and storage resources available within CDC and EDC. This is an extremely hard undertaking as the type and scope of resources across CDC and EDC varies considerably. The CDC and EDC resources have heterogeneous hardware and virtualisation features, for example: (i) CDCs are based on traditional hypervisors such as VMware, Xen, Hyper-V, and/or containers) and (ii) EDCs are based on lightweight containers such as Docker, OpenVZ, or Linux Containers.[7,8,9] Such diversity complicates the choice of a suitable trade-off between fidelity of the simulation model and time efficiency of the simulation runs.

Modelling heterogeneous data programming model abstractions. As IoTDAPs have to deal with a mix of workloads, the simulation framework needs to support modelling of: (i) heterogeneous data programming abstractions and workflow topologies such as message ingestion, batch processing, stream processing, transactional; (ii) heterogeneous dataflows (for example, static, real-time streams, and transactions); (iii) and heterogeneous query operators (continuous query operators in stream processing vs. transactional operators in SQL/NoSQL/batch processing). This requires a formalism to capture in the simulation the successive transformations applied to these streams and the temporal and spatial aspects as well as rates of the data intake.

Holistic performance evaluation methodology. Guaranteed performance of IoTDAP require a clear understanding of performance metrics and bottlenecks across each layer, including EDC, CDC, and data programming models. The problem is complicated as each of these layers require a distinct performance evaluation approach. For example, the benchmarks (workload model) required to stress test CDC resources is different from SDN/NFV networking system or a gateway in EDC. Moreover, the availability and failure distributions across these layers differ, which further complicated the performance evaluation. For example, performance analysis at the EDC level may require to assess network stability, throughput optimality, routing delays, fairness in resource sharing, available bandwidth, and sensor battery state. These performance metrics might be very different from the ones relevant to a CDC operator, who are interested in end-to-end response times, platform scalability, virtual machine utilizations, and the costs of moving data outside the CDC. In order to carry out performance evaluation across different layers of IoTDAP in an abstract way, it is therefore important for the simulator to offer a rich spectrum of metrics. It is also desirable to guide the user in the output analysis of simulation data, for example correlating results across different components and layers for identifying root causes of bottlenecks.

Realistic parameterization of IoT simulation models. It is a major challenge to provide holistic methodologies to guide parameterization of complex simulation models of IoTDAP. The ability to simulate endless possible combinations of resources, topologies, technologies, and workloads may be counterbalanced by the increased parameterization complexity of the simulation models. To worsen the situation, empirical traces collected across multiple devices and layers may not match the abstraction level of the simulation tool. Further research is needed on techniques to close the gap between data collection in real IoT systems and effective parameterization of corresponding simulation models. Large simulation tools should also establish user communities with the necessary expertise to validating the models across a broad spectrum of application domains and provide recommendations on default parameterizations.

Customization and Extensibility. IoTDAP simulators are expected to be easily extended with new device models without the need to know their internals. This allows their wider applicability and enables the currently more academia-focused simulators to be exploited by the industry at large scale. It is important to be able to customize the dynamic behavior of individual components, in order to capture the way actuators react to control signals, their impact on the environment, and the end-to-end control logic implemented by IoTDAP.

Support for Online Decision Making. As future IoTDAP simulators could be used in the process of online decision making, they need to be able to incorporate live data and need to be capable to process multiple scenarios even if involving millions of IoT devices. In order to ensure that the proposed simulation tools meet the expectations of the end users, it is central for the predictions to be accurate for the simulation fabric to be scalable, and for the tools to be easy to learn and use. It is also important to validate simulation accuracy against the real system in a systematic manner, across various stages of the study.

State of the Art of Current Simulators

Several tools have been developed by the research community to simulate systems and technologies that are commonly employed by IoTDAP solutions. Different simulation tools exist for different layers of IoTDAP, we here distinguish between simulators for data processing backends in CDC and solutions focussed on sensors and EDC.

Simulating IoT data processing in the CDC. Popular simulators can be classified based on the big data programming model, and the resource abstractions they are capable of simulating and modelling. Most of current work has focussed in particular on modelling MapReduce (MR) processing:

- MRPerf: In essence, this tool serves as deployment optimization tool for MR or batch processing programming model (e.g., Apache Hadoop—an implementation of MR) on CDC via reduction in the number of set-up parameters that have to be manually tuned.[10] MRPerf captures various aspects of a MR setup, and uses this information to predict expected application performance via simulation.
- Mumak: An open source MR simulator which uses data from real experiments to estimate performance metrics (e.g., completion time) for Map and Reduce tasks with different scheduling algorithms in a CDC environment.[11]
- SimMR: It was developed by HP lab.[12] This tool can replay execution traces of real workloads collected in Hadoop clusters (as well as synthetic traces based on statistical properties of workloads) for evaluating different resource allocation and scheduling ideas in CDC environments.
- MRSim: It is a discrete event based simulator for evaluating performance of Hadoop cluster.[13] It can capture the effects of different configurations of Hadoop cluster setup on data processing activity performance in terms of job completion times and hardware utilization

- MR-Cloudsim: It was developed (by extending Cloudsim) for simulating MR programming model.[14] This is an important extension, but the current tool only supports single-state Map and Reduce computation, without explicitly capturing network links and multitenancy.
- IoTSim (proposed by Zeng et al.): It supports and enables simulation of batch processing activity in IoT systems limiting themselves to the MR model.[15] They also presented a real case study that validates the effectiveness of their simulator.

All the above simulators are only suitable for simulating and modelling performance of the MR programming model, hence insufficient in the context of modelling and simulating the behavior of IoTDAP (see Figure 1), which requires multiple big data programming models and diverse resources types relevant to EDC and CDC environments.

Simulating IoT data gathering and edge processing in the EDC. Typically, IoT specific simulators are focused on describing sensors and edge computing aspects. Some recently-proposed solutions include:

- Brambilla et al. propose to improve the OSmobility extension of the Discrete Event Universal Simulator (a general-purpose discrete event simulation environment) with the ability to model and simulate performance of IoT sensor devices.[16] They evaluate the simulation's performance with up to 200,000 simulated sensors. However, they focus on models related to simulating mobility, communication, and energy consumption of IoT devices (e.g., smart car). Developing simulation models to support high heterogeneity IoT devices and relevant EDC/CDC resources, for supporting more realistic IoTDAP use cases, has been left for the future work.
- SimIoT focuses on modelling the sensor-data processing scenario relevant to remote healthcare IOTDAP.[17] This simulator introduces several techniques (e.g. a broker) that simulates the communication between an IoT sensor and the CDC (no support for EDC or big data programming models).
- Han et al. designed Devices Profile for Web Services (DPWS) DPWSim, which is a simulation toolkit that can support performance evaluation of IoT devices equipped with the Organization for the Advancement of Structured Information Standards standard called DPWS.[18] However, DPWSim is not yet capable of modelling other IoT device types as well as complex CDC/EDC computing environment—which are indeed critical to the design and development of IoTDAPs.
- Though iFogSim focuses on the edge computing aspects of IoT and offers models for some types of sensors/actuators, EDC resources and CDC resources, it requires more support for holistic performance evaluation methodology, model parameterisation, and different types of data programming models.[19]
- DISSECT-CF is a simulator which builds on cloud computing abstractions and it has recently received several extensions towards IoT sensor modelling.[20] This simulator provides simplified description and incorporation of IoT data streams into cloud computing scenarios. The sensor models are demonstrated through a weather forecasting scenario. Although the integrated sensor models are generic, they might still not be inapplicable in future IoT scenarios. Also, the network models applied by the simulator are likely to provide sufficient accuracy in network usage patterns foreseen in IoT landscape.

Amongst recent works, Silva et al. is one of the few which deals with modelling dynamic nature of IoTDAP including fault behaviors of sensing devices and network links.[21] Though authors introduce device and network link specific fault models, they failed to cover failure aspects of EDC/CDC resources as well as data programming models.[21] On the other hand, Dhoutaut et al. addresses IoT simulation from a completely different perspective: simulations are modelled by visual artefacts that allow the observation of the behavior of the simulated IoTDAP in real time and are more similar to industrial modelling tools.[22] This perspective, however, is only applicable to small scale modelling as it limits the level of evaluation to IoT systems that are visually conveyable.

We have identified a set of fundamental challenges faced by simulation and modelling research in order to support design and deployment of IoTDAPs. The high-degree of heterogeneity in IoT technologies, devices, and the diverse application domains imply that current simulation tools are alone insufficient for holistic performance analysis of IoTDAPs. The scale and richness of IoT systems also suggests that simulation scalability may become a critical bottleneck while comparing alternative implementations of IoTDAPs, posing an urgent need to develop a new generation of distributed simulators. These simulators should also be equipped with expressive formalisms to capture diversity across IoT layers. Novel methodologies should be defined to identify performance bottlenecks and parameterize large-scale simulation studies in a practical way. Research is urgently needed to identify the best trade-offs between expressiveness of IoT resource abstractions and simulator scalability. The complexity and scale of these challenges and their relevance to stakeholders across society, industry, and academia mean that community-based projects may be better equipped to deliver lasting solutions. An interdisciplinary approach should also be pursued to ensure that future IoT simulation tools will be able to accurately model the environment in which sensors collected data.

References

1. F. Xia, L.T. Yang, L. Wang, and A. Vinel, "Internet of Things," *Int'l J. Communication Systems*, vol. 25, no. 9, 2012, p. 1101.
2. C. Perera, A. Zaslavsky, P. Christen, and D. Georgakopoulos, "Context Aware Computing for the Internet of Things: A Survey," *IEEE Communications Surveys & Tutorials*, vol. 16, no. 1, 2014, pp. 414–454.
3. IDC, *Worldwide Internet of Things Forecast, 2015–2020*, May 2015, https://www.idc.com/research/viewtoc.jsp?containerId=256397
4. J. Bughin, J. Manyika, J. Woetzel, M. Chui, P. Bisson, and R. Dobbs, *The Internet of Things: Mapping the Value Beyond the Hyper*, McKinsey Global, June 2015.
5. R. Wang and R. Ranjan, "Processing Distributed Internet of Things Data in Clouds," *IEEE Cloud Computing*, vol. 2, no. 1,2015.
6. D. Puthal, S. Nepal, R. Ranjan, and J. Chen, "Threats to Networking Cloud and Edge Data Center in the IoT," *IEEE Cloud Computing*, vol. 3, no. 4, 2016.
7. W. Felter, A. Ferreira, R. Rajamony, and J. Rubio, "An Updated Performance Comparison of Virtual Machines and Linux Containers," *2015 IEEE Int'l Symposium Performance Analysis of Systems and Software* (ISPASS), March 2015, pp. 171–172.
8. A. Celesti, D. Mulfari, M. Fazio, M. Villari, and A. Puliafito,"Exploring Container Virtualization in IoT Clouds," *2016 IEEE Int'l Conf. Smart Computing* (SMARTCOMP), May 2016, pp. 1–6.
9. Y. Xu, V. Mahendran, and S. Radhakrishnan, "SDN Docker: Enabling Application Auto-Docking/Undocking in Edge Switch,"*2016 IEEE Conf. Computer Communications Workshops* (INFOCOM WKSHPS), April 2016, pp. 864–869.
10. G. Wang, A.R. Butt, P. Pandey, and K. Gupta, "A Simulation Approach to Evaluating Design Decisions in MapReduce Setups," *MASCOTS*, vol. 9, Sept. 2009, pp. 1–11.
11. H. Tang, "Mumak: Map-Reduce Simulator. Mumak: Apache," https://issues. Apache. org/jira/browse/MAPREDUCE-728, 2009.
12. A. Verma, L. Cherkasova, and R.H. Campbell, "Play It Again, SimMR!" *2011 IEEE Int'l Conf. Cluster Computing*, September 2011, pp. 253–261.
13. S. Hammoud, M. Li, Y. Liu, N.K. Alham, and Z. Liu, "MRSim: A Discrete Event Based MapReduce Simulator," *2010 Seventh Int'l Conf. Fuzzy Systems and Knowledge Discovery* (FSKD), vol. 6, August 2010, pp. 2993–2997.
14. J. Jung and H. Kim, "MR-CloudSim: Designing and Implementing MapReduce Computing Model on CloudSim," *2012 IEEE Int'l Conf. ICT Convergence* (ICTC), October 2012, pp. 504–509.
15. X. Zeng, S.K. Garg, P. Strazdins, P.. Jayaraman, D. Georgakopoulos, and R. Ranjan, "IoTSim: A Simulator for Analysing IoT Applications," *J. Systems Architecture*, 2016.
16. G. Brambilla, M. Picone, S. Cirani, M. Amoretti, and F. Zanichelli, "A Simulation Platform for Large-Scale Internet of Things Scenarios in Urban Environments," *Proc. 1st Int'l Conf. IoT in*

Urban Space, October 2014, pp. 50–55.
17. S. Sotiriadis, N. Bessis, E. Asimakopoulou, and N. Mustafee, "Towards Simulating the Internet of Things," *2014 28th Int'l Conf. Advanced Information Networking and Applications Workshops* (WAINA), May 2014, pp. 444–448.
18. S.N. Han, G.M. Lee, N. Crespi, K. Heo, N. Van Luong, M. Brut, and P. Gatellier, "DPWSim: A Simulation Toolkit for IoT Applications Using Devices Profile for Web Services," *2014 IEEE World Forum Internet of Things* (WF-IoT), March 2014, pp. 544–547.
19. H. Gupta, A.V. Dastjerdi, S.K. Ghosh, and R. Buyya, "IFogSim: A Toolkit for Modelling and Simulation of Resource Management Techniques in Internet of Things, Edge and Fog Computing Environments," arXiv preprint arXiv:1606.02007, 2016.
20. G. Kecskemeti, "DISSECT-CF: A Simulator to Foster Energy-Aware Scheduling in Infrastructure Clouds," *Simulation Modelling Practice and Theory*, vol. 58, 2015, pp.188–218.
21. I. Silva, R. Leandro, D. Macedo, and L.A. Guedes, "A Dependability Evaluation Tool for the Internet of Things," *Computers & Electrical Engineering*, vol. 39, no. 7, 2013, pp.2005–2018.
22. D. Dhoutaut, B. Piranda, and J. Bourgeois, "Efficient Simulation of Distributed Sensing and Control Environments," *IEEE Green Computing and Communications* (GreenCom) and *Internet of Things* (iThings/CPSCom), *IEEE Int'l Conf. and IEEE Cyber, Physical and Social Computing*, August 2013, pp. 452–459.

GABOR KECSKEMETI *(PhD, University of Westminster, 2011) has been a lecturer in the Department of Computer Science at Liverpool John Moores University, UK since 2016. His research interests include modeling energy efficient and autonomous distributed systems (e.g., clouds and IoT) as well as virtual machine/container image delivery optimization. He has published over 60 scientific papers. Contact him at g.kecskemeti@ljmu.ac.uk.*

GIULIANO CASALE *received the PhD degree in 2006 in Computer Engineering from Politecnico di Milano, Italy. In 2010 he joined the Department of Computing at Imperial College London, UK, where is currently a senior lecturer in modelling and simulation. He teaches and does research in performance engineering, cloud computing, and operations research. Contact him at g.casale@imperial.ac.uk.*

DEVKI NANDAN JHA *is a PhD student in the School of Computing Science at Newcastle University, UK. His research interests include cloud computing, big data analytics, and Internet of Things. Jha has an MTech in Computer Science and Technology from Jawaharlal Nehru University, India. Contact him at d.n.jha2@ncl.ac.uk.*

JUSTIN LYON *is the founder of the pioneering technology company, Simudyne, whose clients have included the Bank of England, Microsoft, BP, the US Department of Defense, Gulf Bank, the European Commission, and PWC. He is also a Director of Ordach, a cybersecurity solution at the cutting edge of cyber risk diagnostics. Contact him at justin@simudyne.com.*

RAJIV RANJAN *is a reader in the School of Computing Science at Newcastle University, UK; chair professor in the School of Computer, Chinese University of Geosciences, Wuhan, China; and a visiting scientist at Data61, CSIRO, Australia. His research interests include grid computing, peer-to-peer networks, cloud computing, Internet of Things, and big data analytics. Ranjan has a PhD in computer science and software engineering from the University of Melbourne (2009). Contact him at raj.ranjan@ncl.ac.uk or http://rajivranjan.net.*