Advances in Orchestrating Sustainable Smart Cities (Part 1)

Rajiv Ranjan, Prem Prakash Jayaraman, Massimo Villari, and Dimitrios Georgakopoulos

RAPID urbanization is a global megatrend with 66 percent of the world's population expected to live in urban areas by 2050. The staggering exponential increase in urbanization is leading to more people migrating to major cities in the search of better opportunities and quality of life. Cities need to increase the efficiency in which they operate and use their resources sustainability in order to meet the demands imposed by rapid urbanization. The challenge is to continue providing basic resources such as sufficient fresh water; cleaner energy; transportation alternatives to commute efficiently from one place to another; adaption to changing climatic conditions; safety and security; while also ensuring economical, social, and environment sustainability.

Advances in disruptive digital technologies ranging from cyber-physical systems, Big Data analytics, clouds, edges, internet of things etc., is widely accepted as the backbone for addressing these growing problems by impacting all city services and making cities smart. Smart cities [1] are envisioned as a wide-scale concord of internet of things [2] with sensors monitoring cyber and physical indicators; actuators that allow dynamic adaption to the changing complex urban environment; and cloud and edge computing offerings affordable easy to tap computing resources. Recently, we are witnessing an increased investment towards the development and deployment of intelligent and sustainable environments. In order to develop sustainable smart city solutions, smart cities must have an open, flexible, and secure platform on which consolidate siloes and services. Such advances opens new opportunities for improving the efficiency of services offered by smart cities while simultaneously exposing novel vulnerabilities. For example, environments that are embedded with millions and billions of dedicated things (such as sensors) and communication capabilities provide real-time access to fine-grained consumer data while exposing new privacy and security risks. Furthermore, as connected things grow, we will witness a significant increase in data that will require novel data processing pipelines, architectures, methods, and techniques. These challenges represent a huge opportunity for a paradigm shift that will require the need for data processing, analysis, and security close to the connected "things" i.e., towards the edge of the network in-order to support the growing smart city ecosystem. This paradigm shift will lead to an explosive growth of independent, owned and operated things and services including gateways, repeaters, smart infrastructure, and systems.

Such a paradigm needs to be architected in a way that is easy to operate and dramatically simplifies the management of service offerings through scalable orchestration and proper automation. It must allow management, integration, and deployment of different tenants (such as services and things independently owned) within the smart city ecosystem in a uniform way. It should also have a suitable policy framework, letting specific stakeholders have access to data produced by other tenants, and analyze and extract values from the data. In order to address these challenges, this special issue solicits high quality original research papers (including smart city experience papers) that made significant contributions to the state-of-the-art in "method and techniques to build sustainable smart city solutions" research area. The call for papers received a number of submissions. After a two-phase peer review process, we have accepted five high-quality papers related to the aforementioned areas of interest which will be published in the October-December 2017 as Part 1.

PART 1: OCTOBER-DECEMBER 2017

The first paper titled "Data Replication and Virtual Machine Migrations to Mitigate Network Overhead in Edge Computing Systems" by Nikos Tziritas, Maria Koziri, Areti Bachtsevani, Thanasis Loukopoulos, George Stamoulis, Samee U. Khan, George Stamoulis, and Cheng-Zhong Xu investigates the use of data replication in conjunction with the VM assignment problem in cloud and edge environment. In particular, they propose algorithms that decides 1) which data should be replicated and where and 2) which VM must be migrated so as to minimize the network overhead between traditional cloud and edge devices. They discuss both the un-capacitated case and the more realistic case whereby datacenters (for the traditional cloud case) and micro-datacenters (for the mobile cloud case) have limited storage and computing capacity. They propose an algorithm based on hyper-graph partitioning to solve the aforementioned problem in an optimal way regarding the unconstrained case and extend it to capture storage and computing capacity constraints. Experimental evaluation shows that the proposed algorithm yields up to

- *R. Ranjan is with Newcastle University, UK. E-mail: raj.ranjan@ncl.ac.uk.*
- *P.P. Jayaraman and D. Georgakopoulos are with RMIT University, Melbourne, Australia. E-mail: {prem.jayaraman, dimitrios.georgakopoulos}@rmit.edu.au.*
- *M. Villari is with the University of Messina, Italy. E-mail: mvillari@unime.it.*

For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org, and reference the Digital Object Identifier below.
Digital Object Identifier no. 10.1109/TSUSC.2017.2773280

53 percent network overhead reduction when compared to state-of-the-art algorithms found in the literature.

The second paper titled "Predicting Transportation Carbon Emission with Urban Big Data" by Xiangyong Lu, Kaoru Ota, Mianxiong Dong, Chen Yu, and Hai Jin addresses an important topic in developing sustainable smart cities solutions namely carbon emission. In particular, the focus of this work is on predicting the real-time and fine-grained transportation carbon emission information in a smart city, based on the spatio-temporal big datasets observed in the city, such as taxi GPS data, transportation carbon emission data, road networks, points of interests (POIs), and meteorological data. They propose a three-layer perceptron neural network (3-layerPNN) to learn the characteristics of collected data and infer the transportation carbon emission. Experimental results show that the proposed methods have significant advantages over well-known machine learning methods such as Gaussian Naive Bayes, Linear Regression, Logistic Regression, and deep learning methods such as Stacked Denoising Autoencoder and Deep Belief Networks.

The third paper titled "Sociability-Driven Framework for Data Acquisition in Mobile Crowdsensing Over Fog Computing Platforms for Smart Cities" by Claudio Fiandrino, Fazel Anjomshoa, Burak Kantarci, Dzmitry Kliazovich, Pascal Bouvry, and Jeanna Neefe Matthews investigates a sociability-driven framework for data acquisition from citizens of smart cities. Data is the main pillar of future smart city solutions and the ability to collect them from the citizens will enable customized and useful smart city services and applications. To this end, the authors propose a novel framework for data acquisition in mobile crowd sensing deployed over a fog computing platform which facilitates important operations like user recruitment and task completion. The proposed user recruitment policy called DSE (Distance, Sociability, Energy) exploits three criteria: i) spatial distance between users and tasks, ii) user sociability, which is an estimate of the willingness of users to contribute to sensing tasks, and iii) remaining battery charge the devices. Results of the performance evaluation reveal that the average number of recruited users improves by nearly 20 percent if compared to policies using only spatial distance as selection criterion.

The fourth paper titled "Lightweight Mutual Authentication for IoT and Its Applications" by Nan Li, Dongxi Liu, and Surya Nepal investigates security schemes to support the development of smart city applications underpinned by Internet of Things data. The uninterrupted and accurate functioning of IoT devices are critical to smart city applications as crucial decisions will be made based on the data received. The challenge is in designing a secure mutual authentication protocol that is affordable and efficient for resource constrained devices. They propose a lightweight mutual authentication protocol based on a novel public key encryption scheme. The proposed protocol takes a balance between the efficiency and communication cost without sacrificing the security. Experimental evaluation of the proposed scheme reveal, at a similar security level, the proposed scheme's performance is significantly better than existing RSA and ECC based schemes.

The fifth paper titled "Graph Encryption for Top-K Nearest Keyword Search Queries on Cloud" by Chang Liu, Liehuang Zhu, and Jinjun Chen investigates graph encryption method for an important type graph queries in smart cities called top-k Nearest Keyword (kNK) searches. The authors claim that though keyword searches over encrypted textual data have been extensively studied, approaches for encrypting graph-structured data with support for answering graph queries are still in its infancy. They design several indexes to store necessary information for answering queries and guarantee that private information about the graph such as vertex identifiers, keywords, and edges are encrypted or excluded. Security and efficiency of the proposed graph encryption scheme are demonstrated and supported by theoretical proofs and experiments on real-world datasets, respectively.

Overall, the special issue compiles an excellent list of papers which make ground-breaking contributions to the design and development of techniques that will empower future smart cities. We hope that both scientific and academic communities will benefit from these special issues.

Rajiv Ranjan
Prem Prakash Jayaraman
Massimo Villari
Dimitrios Georgakopoulos
Guest Editors

REFERENCES

[1] M. Vögler, J. M. Schleicher, C. Inzinger, S. Dustdar, and R. Ranjan, "Migrating smart city applications to the cloud," *IEEE Cloud Comput.*, vol. 3, no. 2, pp. 72–79, Mar.-Apr. 2016. doi: 10.1109/MCC.2016.44

[2] L. Wang and R. Ranjan, "Processing distributed internet of things data in clouds," *IEEE Cloud Computing*, vol. 2, no. 1, pp. 76–80, Jan./Feb. 2015. doi: 10.1109/MCC.2015.14

Rajiv Ranjan is a reader (associate professor) in computing science with Newcastle. He is a chair (visiting) professor with the Chinese University of Geosciences, Wuhan, and the inventor of CloudSim, the world's most adopted and cited (2150+ Google citations since 2011) Cloud and Big Data application simulation and benchmarking toolkit. He was awarded the 2016 IEEE TCSC Award for Excellence in Scalable Computing (Middle Career Researcher), and secured £5M+ in funding. For more than 10 years, he has conducted seminal research around the development of generic resource management models and methods for efficient scheduling and resource allocation for all types of parallel and distributed computing systems such as cloud computing, big data analytics, and Internet of Things (IoT). He has established an outstanding understanding of the fundamentals of computer science, which has enabled him to easily apply his research for the benefit of the society. For example, he has easily extended his research to develop novel IoT applications (e.g., smart buildings, urban hazard management) for supporting and fostering sustainable cities and societies. He has served on the technical program committees of 50+ international IEEE/ACM/EAI workshops and conferences. He has a proven record of international leading research reflected by 200+ scientific publications (full publication list available at: https://rajivranjan.net) that include 131 journal articles, 51 conference papers, 12 book chapters, and 11 edited research books. His research has been highly cited internationally (Google Scholar h-index 41, 9300+ citations; Scopus h-index 22, 4300+ citations, Thomson Reuters h-index 20, 2300+ citations). His papers have appeared in highly selective (as per Australian Ranking:- core.edu.au) journals and conferences including, the *ACM Computing Surveys* (ERA A*), the *IEEE Transactions on Parallel and Distributed Systems* (ERA A*), the *IEEE Transactions on Computers* (ERA A*), WWW (CORE A*), PVLDB (CORE A*), JCCS (ERA A*), FGCS (ERA A), SPE (ERA A), and the ACM TECS (ERA A). He serves on editorial boards of the *IEEE Transaction on Cloud Computing*, the *IEEE Transaction on Computers* (2015-2017), the *IEEE Cloud Computing*, the *IEEE IT Professional*, and *Springer Computing*.

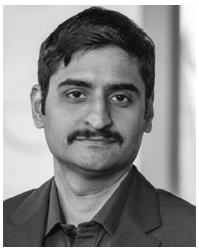

Prem Prakash Jayaraman is currently a research fellow with the Swinburne University of Technology, Melbourne. His research areas of interest include, Internet of Things, cloud computing, mobile computing, sensor network middleware, and semantic internet of things. He has authored/co-authored more than 50 research papers in international journals and conferences such as the *IEEE Transaction on Cloud Computing*, the *IEEE Selected Areas in Communication*, the *Journal of Computational Science*, the *IEEE Transactions on Emerging Topics in Computing*, *Future Generation Computing Systems*, *Springer Computing*, the *ACM Ubiquity Magazine*, and *IEEE Magazine*. He is one of the key contributors of the Open Source Internet of Things project OpenIoT that won the prestigious Black Duck Rookie of the Year Award in 2013. He has been the recipient of several awards including hackathon challenges at the 4th International Conference on IoT (2014) at MIT Media Lab, Cambridge, MA and IoT Week 2014 in London and the best paper awards at HICSS 2016/2017 and IEA/AIE-2010. Previously, he was a postdoctoral research fellow at CSIRO Digital Productivity Flagship, Australia, from 2012 to 2015.

Massimo Villari is associate professor of computer engineering with the University of Messina (Italy). He is actively working as IT security and distributed systems analyst in cloud computing, virtualization, and storage. For the EU Projects "RESERVOIR," he led the IT security activities of the whole project. For the EU Project "VISION-CLOUD," he covered the role of architectural designer for UniME. He is currently Scientific ICT Responsible in the EU Project frontierCities, the Accelerator of FIWARE on Smart Cities–Smart Mobility. He is strongly involved in EU Future Internet initiatives, specifically cloud computing and security in distributed systems. He is co-author of more of 130 scientific publications and patents in cloud computing (cloud federation), distributed systems, wireless network, network security, cloud security, and cloud and IoTs. He was general chair of ESOCC 2015 and IEEE-ISCC 2016. Since 2011, he has been a fellow of IARIA, recognized as a Cloud Computing Expert, and involved in the activities of the FIArch, the EU Working Group on Future Internet Architecture. In 2014, he was recognized by an independent assessment (the *IEEE Transactions on Cloud Computing*, vol. 2, no. 3, Jul-Sep. 2014) as one of the world-wide active scientific researchers, top 27 classification, in cloud computing area. He is general chair of EAI-CN4IoT. He is editor-in-chief of *EAI Endorsed Transactions on Smart Cities*. Currently, he is Scientific Responsible for the UniME-IRCCSME Cloud initiative in eHealth.

Dimitrios Georgakopoulos joined the Swinburne University of Technology in July 2016 to start a new Key Laboratory on IoT. Before that, he served as a research director (2008-2014) at CSIRO's ICT Centre and the executive director of the Information Engineering Laboratory, which was the largest computer science research program in CSIRO and Australia. Prior to joining CSIRO, he held research and management positions in several industrial laboratories in the USA, including Telcordia Technologies (where he helped found two of Telcordia's Research Centers in Austin, Texas, and Poznan, Poland); Microelectronics and Computer Corporation (MCC) in Austin, Texas; GTE (currently Verizon) Laboratories in Boston, Massachusetts; and Bell Communications Research (Bellcore) in Piscataway, New Jersey. He was also a full professor at RMIT University (2014-2016), and is currently an adjunct professor with the Australian National University and a CSIRO adjunct fellow. He has authored/co-authored more than 160 journal and conference publications in computer science, which include three seminal papers in the areas of service computing, workflow management, and context management for the Internet of Things (IoT). According to Google Scholar, he has more than 11,000 lifetime citations. He has served as the general or program chair of more than 20 major international conferences as well as many other smaller conferences or workshops. He has received two outstanding paper awards from the IEEE Computer Society (CS), as well as several IEEE CS service awards. He was the recipient of GTE's (Verizon's) Excellence Award. He has attracted significant external research funding (AUD$35M+) from various industry and government research funding agencies, ranging from DARPA and ARDA in the USA, to the Framework Program in the EU, to the Department of Human Services in Australia.