Journal of Parallel and Distributed Computing 132 (2019) 160–163

Contents lists available at ScienceDirect

J. Parallel Distrib. Comput.

journal homepage: www.elsevier.com/locate/jpdc

Editorial

A note on tools and techniques for end-to-end QoS monitoring in Internet of Things

The Internet of Things (IoT) paradigm promises to develop the technologies and knowledge for integrating physical data collecting sources (e.g. rain gauges, weather stations, pore pressure sensors, tilt meters), and on-line data collecting sources (e.g. social media sources such as Twitter, Instagram, and Facebook) with powerful and flexible data processing computing facilities like cloud computing. Such capability will benefit a wide range of application domains including (but not limited to) smart cities, smart home systems, smart agriculture, health care monitoring, and environmental monitoring (e.g. landslides, heatwave, flooding).

Traditionally, big data sets generated by IoT applications have been stored and processed by traditional cloud datacenters (Amazon Web Services, and Microsoft Azure for example). However, recently the traditional centralized model of cloud computing is undergoing a paradigm shift towards a decentralized model, which aims to take advantage of the recent evolution of the smart hardware devices at the network edge such as smart gateways (e.g. Raspberry Pi 3, UDOO board, and esp8266) and network function virtualisation solutions (e.g. Cisco IOx, HP OpenFlow and Middlebox Technologies). These devices on the network edge (sometimes referred to as Edge datacenters) can offer computing and storage capabilities on a smaller scale in order to support traditional cloud datacenters in tackling the future data processing and application management challenges that arise within IoT application ecosystems. Ultimately, the success of IoT applications will critically depend on the intelligence of tools and techniques that can monitor and verify the correct operation of such IoT ecosystems from end to end, including the sensors, big data programming models, and the hardware resources available in the edge and cloud datacenters.

In the past 20 years a large body of research has developed frameworks and techniques to monitor the performance of hardware resources and applications in distributed system environments (grids, clusters, and clouds). Monitoring tools that were popular in the grid and cluster computing era included R-GMA, Hawkeye, Network Weather Service (NWS), and Monitoring and Directory Service (MDS). These tools were concerned only with monitoring performance metrics at the hardware resource-level (CPU percentage, TCP/IP performance, available non-paged memory), and not at the application-level (e.g. event detection delay in the context of particular IoT applications). On the other hand, cluster-wide monitoring frameworks (Nagios, Ganglia, adopted by big data orchestration platforms such as YARN, Apache Hadoop, and Apache Spark) provide information about hardware resource-level metrics (cluster utilisation, CPU utilisation, memory utilisation). In the public cloud computing space, monitoring frameworks and techniques (e.g. Amazon CloudWatch used by Amazon Elastic MapReduce, Azure Fabric Controller) typically monitor an entire CPU resource as a black box, and so cannot monitor application-level performance metrics specific to IoT ecosystem, whereas techniques and frameworks such as Monitis and Nimsoft can monitor application-specific performance metrics (such as web server response time).

In summary, none of these approaches are capable of monitoring, and detecting root causes of failures and performance degradation for entire end-to-end IoT ecosystems across the edge datacenter (physical layer), the Network (communications layer), and the Big Data platforms (cloud datacenter layer). Developing formal approaches for monitoring end-to-end IoT ecosystems is what we term the grand challenge, and current platforms and techniques for monitoring IoT and Cloud computing fall short of this grand challenge [9].

This special issue solicits papers related to topics including: (1) Scalable algorithms for monitoring performance across IoT sensors, big data programming models and cloud/edge datacenters; (2) Intelligent energy aware performance monitoring algorithms for edge datacenters; (3) Techniques for end-to-end SLA and contract monitoring (between users, services, and data sources); (4) Models and machine learning techniques for automatically predicting root causes of performance degradation end-to-end; (5) Ontology models for capturing heterogeneous performance metrics for end-to-end IoT components across cloud/edge datacenters; (6) Novel middleware services for monitoring end-to-end IoT eco-systems across cloud/edge datacenters; (7) Innovative IoT performance benchmarking and performance profiling use cases. The call for papers for this special issue received a number of submissions. After a two-phase peer review process, we have accepted the following high quality papers related to the above areas of interest.

The paper titled *Ada-Things: An Adaptive Virtual Machine Monitoring and Migration Strategy for Internet of Things applications* [10], addresses the challenge of how to efficiently monitor and allocate VM resources to realize load balancing among edge clouds. In this paper, the authors propose Ada-Things, an adaptive VM monitoring and live migration Strategy for IoT applications in edge cloud architecture. The authors describe the idea behind Ada-Things that the migration method of a VM should be determined by its workload characteristics. Specifically, based on the variation of current memory dirty page rate in IoT applications, they explain that Ada-Things can adaptively select the

https://doi.org/10.1016/j.jpdc.2019.05.004

most appropriate migration method to copy memory pages, thus addressing two limitations of existing VM migration methods in edge cloud (application generality and performance imbalance). The paper presents evaluation results that show, compared with traditional methods, Ada-Things can significantly reduce the total migration time by 21%, the VM downtime by 38% and the amount of pages transferred by 29% in average.

The paper titled *QoS-aware Service Recommendation Based on Relational Topic Model and Factorization Machines for IoT Mashup Applications* [2], addresses the challenge of finding suitable Web APIs to build IoT Mashup applications for developers. The authors argue that even if the existing service recommendation methods show improvements in service discovery, the accuracy of them can be significantly improved due to overlooking the impact of sparsity and multiple-dimension information of QoS between Mashup and services on recommendation accuracy. The authors propose a QoS-aware service recommendation based on relational topic model and factorization machines for IoT Mashup applications. This method first uses relational topic model to characterize the relationships among Mashup, services, and their links, and mine the latent topics derived by the relationships. Second, it exploits factorization machines to train the latent topics for predicting the link relationship among Mashup and services to recommend adequate relevant top-k Web APIs for target IoT Mashup creation. Finally, the authors conduct a comprehensive evaluation to measure performance of their method. Compared with other existing recommendation approaches, experimental results show that this approach achieves a significant improvement in terms of precision, recall, and F-measure.

The paper titled *Quality of Experience (QoE)-aware Placement of Applications in Fog Computing Environments* [7], explains that the hierarchical, distributed and heterogeneous nature of computational instances make application placement in Fog a challenging task. Diversified user expectations and different features of IoT devices also intensify the application placement problem. Placement of applications to compatible Fog instances based on user expectations can enhance Quality of Experience (QoE) regarding the system services. The authors propose a QoE-aware application placement policy that prioritizes different application placement requests according to user expectations and calculates the capabilities of Fog instances considering their current status. In a Fog computing environment, it also facilitates placement of applications to suitable Fog instances so that user QoE is maximized in respect of utility access, resource consumption and service delivery. The proposed policy is evaluated by simulating a Fog environment using iFogSim. Experimental results indicate that the policy significantly improves data processing time, network congestion, resource affordability and service quality.

The paper titled *An Automatic Performance Model-based Scheduling Tool for Coupled Climate System Models* [5], addresses the challenge of improving the performance of the prediction ability of climate systems. The authors explain that the climate system model, is one of most challenging applications in scientific computing. It utilizes multi-physics simulation that couples multiple components, conducts decadal to millennium simulations, and has long been an important application on supercomputers. However, current climate system models suffer from inefficient task scheduling methods resulting in an intolerable simulation time. To address such challenge, this paper first constructs a lightweight and accurate performance model for effectively capturing and predicting the heterogeneous time-to-solution performance of end-to-end CESM (Community Earth System Model) components with a given simulation configuration. Then, based on the performance model, the authors further propose an efficient scheduling strategy based on rectangular packing method to determine the best process layout among different components, and the process numbers assigned to each component. Evaluations show that this work can achieve 58% average run time reductions on CESM comparing to the widely used sequential process layout for a scale of 144–480 cores on typical CPU clusters. They go on to explain that their work can save 4 million CPU hours when conducting one standard scientific experiment (a 2870-year simulation), which equals to saving $40,089 with a charge of $0.01 per CPU hour. Meanwhile, 26% extra performance improvements also could be gained in their methods comparing to the heuristic branch and bound algorithm with the guidance of the known curve-fitting performance model.

The paper titled *Optimal LEACH protocol with modified bat algorithm for big data sensing systems in Internet of Things* [4], tackles the challenge of how to reduce power consumption within Internet of Things applications. The authors explain that currently, low energy adaptive clustering hierarchy (LEACH) protocol is one well-known algorithm used in BDSS (Big Data Sensing Systems) with low energy cost. In this paper, a new variant of the bat algorithm combined with centroid strategy is introduced. Three different centroid strategies with six different designs are introduced. In addition, the velocity inertia-free update equation is also provided. The optimization performance is verified by CEC2013 benchmarks in those designs against standard BA. Simulation results prove that the bat algorithm with weighted harmonic centroid (WHCBA) strategy is superior to other algorithms. By integrating WHCBA into LEACH protocol, the authors develop a two-stage cluster-head node selection strategy and can save more energy compared to the standard LEACH protocol.

The paper titled *Person Re-identification with Multiple Similarity Probabilities Using Deep Metric Learning for Efficient Smart Security Applications* [11], addresses the research challenges of re-identification especially in terms of recognizing the different appearances of the same person in a harsh real-world environment: (1) the adaptability of the selected features to the dynamic environment cannot be guaranteed, and (2) existing methods rooted from metric learning aim to find a single metric function, and they lack the ability to measure the different appearances of the same person. To address these problems, this study proposes a multiple deep metric learning method empowered by the functionality of person similarity probability measurement. The proposed method exploits multiple stacked auto-encoder networks and classification networks to quantify pedestrian similarity relations. The stacked auto encoder networks directly recognize persons from surveillance images at the pixel level. The classification networks are equipped with the Softmax regression models and produce multiple similarity probabilities to characterize different appearances belonging to the same person. An Adaboost like model is designed to fuse the probabilities corresponding to multiple metrics, which ensures a high accuracy of recognition. Experimental results on two public datasets (VIPeR and CUHK-01) indicate that the proposed method outperforms existing algorithms by 2%–10% at rank 1. Based on the similarity probabilities learned by the proposed model, the algorithm for matching the person pair can achieve a time complexity as low as O(n), which can be deployed at a large scale on the distributed intelligent surveillance network, with each node maintaining limited computing capabilities.

The paper titled *A Provably Secure and Anonymous Message Authentication Scheme for Smart Grids* [6], explains that Smart grid (SG) is an agenda of interest to governments and researchers, as a successful cyberattack against smart grids (e.g. compromise of smart meters) can have devastating real-world consequences, ranging from financial loss to fatalities. In this paper, the authors present a novel and secure message authentication scheme, which provides mutual authentication and key establishment for smart grid. The scheme is also designed to preserve the identities of the gateways during message transmission. The paper then proves the security of the scheme, as well as verifying the security properties using Proverif and demonstrating the utility of the scheme using simulations.

The paper titled *Performance Evaluation of FIWARE: A Cloud-Based IoT Platform for Smart Cities* [1], highlights the neccessity of performance benchmarking IoT platforms. As the Internet of Things (IoT) becomes a reality, millions of devices will be connected to IoT platforms in smart cities. These devices will cater to several areas within a smart city such as healthcare, logistics, and transportation. These devices are expected to generate significant amounts of data and requests at high data rates, therefore, necessitating the performance benchmarking of IoT platforms to ascertain whether they can efficiently handle such devices. In this article, they present results gathered from extensive performance evaluation of the cloud-based IoT platform, FIWARE. In particular, to study FIWARE's performance, they developed a testbed and generated CoAP and MQTT data to emulate large-scale IoT deployments, crucial for future smart cities. They performed extensive tests and studied FIWARE's performance regarding vertical and horizontal scalability. They present bottlenecks and limitations regarding FIWARE components and their cloud deployment. Finally, they discuss cost-efficient FIWARE deployment strategies that can be extremely beneficial to stakeholders aiming to deploy FIWARE as an IoT platform for smart cities.

The paper titled *A Framework for Real Time End to End Monitoring and Big Data Oriented Management of Smart Environments* [3], contributes the following work on monitoring the correct operation of smart ecosystems: The success of Internet of Things (IoT) applications depend on the intelligence of tools and techniques that can monitor, manage, and verify the correct operations of smart ecosystems including sensors and big data analytics tools, typically deployed in Cloud and Edge computing datacenters. In their paper, the authors propose a framework for the monitoring and management of IoT systems that integrates the AllJoyn functionalities, useful to interconnect IoT devices, MongoDB, to implement Big Data storage, and Storm, to run real-time data analytics. They implemented the proposed framework and tested its main functionalities in a smart home application scenario. In their experimentation, they investigated three different data patterns, i.e, regular, event-based, and automated, in order to evaluate performance of our framework in terms of response time under different operational conditions. Experimental results show that the latency of the monitoring and service strongly depends on the type of management application running in the system, whereas it is lightly affected by the data patterns.

The paper titled *FOCAN: A Fog-supported Smart City Network Architecture for Management of Applications in the Internet of Everything Environments* [8], contributes the following: Smart city vision brings emerging heterogeneous communication technologies such as Fog Computing (FC) together to substantially reduce the latency and energy consumption of Internet of Everything (IoE) devices running various applications. The key feature that distinguishes the FC paradigm for smart cities is that it spreads communication and computing resources over the wired/wireless access network (e.g., proximate access points and base stations) to provide resource augmentation (e.g., cyberforaging) for resource- and energy-limited wired/wireless (possibly mobile) things. Motivated by these considerations, this paper presents a Fog-supported smart city network architecture called Fog Computing Architecture Network (FOCAN), a multitier structure in which the applications are running on things jointly compute, route, and communicate with one another through the smart city environment. FOCAN decreases latency and improves energy provisioning and the efficiency of services among things with different capabilities. In particular, three types of communications are defined between FOCAN devices — interprimary, primary, and secondary communication — to manage applications in a way that meets the quality of service standards for the Internet of Everything. One of the main advantages of the proposed architecture is that the devices can provide the services with low energy usage and in an efficient manner. Simulation results for a selected case study demonstrate the tremendous impact of the FOCAN energy-efficient solution on the communication performance of various types of things in smart cities.

Acknowledgments

We would like to thank the authors and all the reviewers for their hard work in helping us put together this special issue. We would also like to thank the Editor-in-Chief, and acknowledge the support of the JDPC editorial staff.

References

[1] V. Araujo, K. Mitra, S. Saguna, C. Åhlund, Performance evaluation of FIWARE: A cloud-based IoT platform for smart cities, J. Parallel Distrib. Comput. 132 (2019) 250–261.
[2] B. Cao, J. Liu, Y. Wen, H. Li, Q. Xiao, J. Chen, QoS-aware service recommendation based on relational topic model and factorization machines for IoT mashup applications, J. Parallel Distrib. Comput. 132 (2019) 177–189.
[3] A. Celesti, M. Fazio, A framework for real time end to end monitoring and big data oriented management of smart environments, J. Parallel Distrib. Comput. 132 (2019) 262–273.
[4] Z. Cui, Y. Cao, X. Cai, J. Cai, J. Chen, Optimal LEACH protocol with modified bat algorithm for big data sensing systems in Internet of Things, J. Parallel Distrib. Comput. 132 (2019) 217–229.
[5] N. Ding, W. Xue, Z. Song, H. Fu, S. Xu, W. Zheng, An automatic performance model-based scheduling tool for coupled climate system models, J. Parallel Distrib. Comput. 132 (2019) 204–216.
[6] X. Li, F. Wu, S. Kumar, L. Xu, A.K. Sangaiah, K.-K.R. Choo, A provably secure and anonymous message authentication scheme for smart grids, J. Parallel Distrib. Comput. 132 (2019) 242–249.
[7] R. Mahmud, S.N. Srirama, K. Ramamohanarao, R. Buyyaa, Quality of experience (QoE)-aware placement of applications in fog computing environments, J. Parallel Distrib. Comput. 132 (2019) 190–203.
[8] P.G.V. Naranjo, Z. Pooranian, M. Shojafar, M. Conti, R. Buyya, FOCAN: A Fog-supported smart city network architecture for management of applications in the Internet of Everything environments, J. Parallel Distrib. Comput. 132 (2019) 274–283.
[9] E. Solaiman, R. Ranjan, P. Jayaraman, K. Mitra, Monitoring internet of things application ecosystems for failure, IEEE IT Prof. (2016).
[10] Z. Wang, D. Sun, G. Xue, S. Qian, G. Li, M. Li, Ada-Things: An adaptive virtual machine monitoring and migration strategy for internet of things applications, J. Parallel Distrib. Comput. 132 (2019) 164–176.
[11] M. Xiong, D. Chen, J. Chen, J. Chen, B. Shi, C. Liang, R. Hu, Person re-identification with multiple similarity probabilities using deep metric learning for efficient smart security applications, J. Parallel Distrib. Comput. 132 (2019) 230–241.

Rajiv Ranjan is a Professor in Computing Science at Newcastle University, United Kingdom. Prior to that, he was a Senior Research and Julius Fellow at CSIRO, Canberra, where he was working on projects related to Cloud and big data computing. He has been conducting leading research in the area of Cloud and big data computing developing techniques for: (i) Quality of Service based management and processing of multimedia and big data analytics applications across multiple Cloud data centers (e.g., CSIRO Cloud, Amazon and GoGrid); and (ii) automated decision support for migrating applications to data centers. He has published about 110 papers that include

60+ journal papers. He serves on the editorial board of IEEE Transactions on Computers, IEEE Transactions on Cloud Computing, IEEE Cloud Computing, and Future Generation Computer System Journals. He is one of the highly cited authors (top 0.09%) in computer science and software engineering worldwide (h-index=43, g-index=94, and 11,000+ google scholar citations).

Ellis Solaiman is a Lecturer at the School of Computing, Newcastle University. He previously received his Ph.D. in Computing Science also from Newcastle University, where he subsequently worked as a Research Associate and Teaching Fellow. His research interests are mainly in the areas of Dependability and Trust in Distributed Systems such as the Cloud and the Internet of Things. He is also interested in the automated monitoring of these systems using technologies such as Smart Contracts. He is a Fellow of the UK Higher Education Academy (FHEA) since 2016.

Massimo Villari an associate professor of computer science at the University of Messina. His research interests include cloud computing, Internet of Things, big data analytics, and security systems. Villari has a Ph.D. in computer engineering from the University of Messina. He's a member of IEEE and IARIA boards. Contact him at mvillari@unime.it.

Paul Watson is a profess or of computer science at Newcastle University. His research interests include scalable computing systems, data-intensive problems, and cloud computing. Watson has a Ph.D. in computer science from Manchester University. Contact him at paul.watson@newcastle.ac.uk.

Rajiv Ranjan
Ellis Solaiman
School of Computing, Urban Sciences Building, Newcastle University, 1 Science Square, Newcastle Helix, Newcastle upon Tyne, NE4 5TG, United Kingdom
E-mail addresses: raj.ranjan@ncl.ac.uk (R. Ranjan), ellis.solaiman@newcastle.ac.uk (E. Solaiman).

Massimo Villari
Faculty of Engineering, University of Messina, C.Da Di Dio, N1 - 98166 - Messina, Italy
E-mail address: mvillari@unime.it.

Paul Watson
School of Computing, Urban Sciences Building, Newcastle University, 1 Science Square, Newcastle Helix, Newcastle upon Tyne, NE4 5TG, United Kingdom
E-mail address: paul.watson@newcastle.ac.uk.