Guest Editorial: Special Section on Advances in Big Data Analytics for Management

Giuliano Casale, *Member, IEEE*, Yixin Diao, *Fellow, IEEE*, Marco Mellia, *Senior Member, IEEE*, Rajiv Ranjan, *Senior Member, IEEE*, and Nur Zincir-Heywood, *Member, IEEE*

I. Introduction

CLOUD and network analytics can harness the immense stream of operational data from clouds and networks, and can perform analytics processing to improve reliability, automated configuration, performance, and optimized network management in general. In this area, we have witnessed a growing trend towards using statistical analysis and machine learning techniques to improve operations and management of IT systems and networks.

Research is therefore needed to understand and improve the potential and suitability of Big Data analytics in the context of systems and network management. This will not only provide deeper understanding and better decision making based on largely collected and available operational data, but present opportunities for improving data analytics algorithms and methods on aspects such as accuracy and scalability. Moreover, there is an opportunity to define novel platforms that can harness the vast operational data and advanced data analysis algorithms to drive management decisions in networks, data centers, and clouds.

This special section of IEEE TRANSACTIONS ON NETWORK AND SERVICE MANAGEMENT deals with *Advances in Big Data Analytics for Management*. It is a second special section, after the success of last year's edition [1]. It presents recent, emerging approaches and technical solutions that can exploit Big Data and analytics in management solutions, as well as platform improvements and specific designs to speed up the processing of large amounts of data that a network typically generates.

We have accepted 10 papers from 36 papers submitted to the open call in this special section that address the underlying challenges of *Big Data Analytics for Management* and present novel theoretical and experimentation results.

II. Special Section Overview

The ten accepted papers in this special section cover three important areas of *Big Data Analytics for Management*: analytics for network management, anomaly detection and security, and analytics platforms and applications.

G. Casale is with Imperial College London, London SW7 2AZ, U.K.
Y. Diao is with IBM T. J. Watson Research Center, Yorktown Heights, NY 10598 USA (e-mail: diao@us.ibm.com).
M. Mellia is with Politecnico di Torino, 10129 Turin, Italy.
R. Ranjan is with Newcastle University, Newcastle upon Tyne NE1 7RU, U.K.
N. Zincir-Heywood is with Dalhousie University, Halifax, NS B3H 4R2, Canada.
Digital Object Identifier 10.1109/TNSM.2018.2806897

A. Analytics for Network Management

Four papers in this special section focus on Analytics for Network management, considering data centers and SDN networks.

In "AWESoME: Big Data for Automatic Web Service Management in SDN," Trevisan *et al.* [item 1) in the Appendix] present a comprehensive Web traffic management approach based on a "per service" management concept to identify and prioritize all traffic of important Web services. It leverages big data algorithms to automatically build models describing the traffic of Web services, and uses the models to install rules in SDN switches to steer all flows related to the originating services. Extensive experimental results are shown using various available traffic traces.

In "Big Data Analysis-Based Secure Cluster Management for Optimized Control Plane in Software-Defined Networks," Wu *et al.* [item 2) in the Appendix] explore a big data analysis-based security cluster management architecture for an optimized control plane. To achieve this, the authors propose an ant colony based optimization scheme and a security authentication scheme. Simulations and evaluations show the feasibility and the efficiency of the proposed cluster management architecture.

In "Spatial-Temporal Prediction Models for Active Ticket Managing in Data Centers," Xue *et al.* [item 3) in the Appendix] present a methodology for active ticket managing to achieve efficient time series prediction and capacity planning. In doing so, their goal is to reduce virtual machine and box usage tickets that are issued in production data centers. They evaluate their methodology on a large number of data center production traces as well as on a cluster running MediaWiki.

In "Mining Causality of Network Events in Log Data," Kobayashi *et al.* [item 4) in the Appendix] propose a method to mine causality of network events from large size of heterogeneous network log messages. The proposal leverages a causal inference algorithm that reconstructs causal structures from a set of time series of events. The authors evaluate their method on 15 months of network syslog data obtained in a nation-wide academic network.

B. Anomaly Detection and Security

Three papers in this special section focus on Anomaly Detection and Security, using Big Data approaches.

In "Detecting Botclouds at Large Scale: A Decentralized and Robust Detection Method for Multi-Tenant Virtualized

Environments," Cogranne *et al.* [item 5) in the Appendix] propose a two-step approach based on Principle Component Analysis and statistical hypothesis theory for the detection of malicious activities perpetrated by virtual hosts infected by botnets. They validate their approach on a large scale dataset providing real container traces as well as simulated ones. This enables them to reach to public cloud context in their evaluations.

In "Anomaly Detection in Complex Real World Application Systems," Gow *et al.* [item 6) in the Appendix] present a Whole of Service Anomaly Detection (WoSAD) methodology that is based on a mix of data driven and model based methods for service profiling and measurement. The authors examine anomaly detection efficiency in two case studies using six detection models in a large Australian Financial Services Organisation. The proposed WoSAD methodology is demonstrated to be more efficient than alternative individual transactions and service models.

In "Social Plane for Recommenders in Network Performance Expectation Management," Zhang *et al.* [item 7) in the Appendix] propose a "social plane" that relies on recommended measurements based on "content-based filtering" to run similarity analysis and subscribe to useful measurements, and "collaborative filtering" to share knowledge on anomaly symptoms. The authors show the effectiveness of the social plane approach within a SoyKB Big Data application case study using social network creation and mingling of experts.

C. Analytics Platforms and Applications

Three papers in this special section focus on Analytics Platforms and Applications, leveraging network properties to improve performance.

In "Efficient Deep Neural Network Serving: Fast and Furious," Yan *et al.* [item 8) in the Appendix] present a dynamic scheduling framework powered by an interference-aware queueing-based analytical model. It identifies and switches to the optimal parallel configuration of the serving system to minimize response latency for deep neural network serving. The benchmark based evaluation demonstrates its good latency prediction accuracy and efficiency as well as its abilities to identify optimal parallel configurations and to adapt to changing load conditions.

In "Group Mobility Detection and User Connectivity Models for Evaluation of Mobile Network Functions," Suzuki *et al.* [item 9) in the Appendix] present a group mobility detection method based solely on signaling data. Then, using this method, the authors build connected/idle duration models for users to characterize network utilization. The simulation results show that the group mobility detection and the connection/idle duration models based on control plane data analytics are useful for the development of mobility-aware functions in base stations.

In "Automatic Generation of Workload Profiles Using Unsupervised Learning Pipelines," Prats *et al.* [item 10) in the Appendix] introduce a methodology for discovery of resource consumption phases in data center applications. The authors propose to combine Conditional Restricted Boltzmann Machines and Hidden Markov Models to deliver an unsupervised phase detection method. The technique is shown to be effective in identifying workload phases in Apache Hadoop and Spark workloads, among others.

ACKNOWLEDGMENT

The authors sincerely thank the authors for their contributions. This special section would not have been possible without their support and sharing of their expertise to benefit the broader audience of this special section. They are very grateful to all the reviewers, who have helped them improve the quality and presentation of each paper. The authors are especially thankful to their Editor-in-Chief, Rolf Stadler, for providing enthusiastic support throughout the process.

APPENDIX
RELATED WORK

1) M. Trevisan, I. Drago, M. Mellia, H. H. Song, and M. Baldi, "AWESoME: Big data for automatic Web service management in SDN," *IEEE Trans. Netw. Service Manag.*, vol. 15, no. 1, pp. 13–26, Mar. 2018.
2) J. Wu, M. Dong, K. Ota, J. Li, and Z. Guan, "Big data analysis-based secure cluster management for optimized control plane in software-defined networks," *IEEE Trans. Netw. Service Manag.*, vol. 15, no. 1, pp. 27–38, Mar. 2018.
3) J. Xue, R. Birke, L. Y. Chen, and E. Smirni, "Spatial-temporal prediction models for active ticket managing in data centers," *IEEE Trans. Netw. Service Manag.*, vol. 15, no. 1, pp. 39–52, Mar. 2018.
4) S. Kobayashi, K. Otomo, K. Fukuda, and H. Esaki, "Mining causality of network events in log data," *IEEE Trans. Netw. Service Manag.*, vol. 15, no. 1, pp. 53–67, Mar. 2018.
5) R. Cogranne, G. Doyen, N. Ghadban, and B. Hammi, "Detecting botclouds at large scale: A decentralized and robust detection method for multi-tenant virtualized environments," *IEEE Trans. Netw. Service Manag.*, vol. 15, no. 1, pp. 68–82, Mar. 2018.
6) R. Gow, F. A. Rabhi, and S. Venugopal, "Anomaly detection in complex real world application systems," *IEEE Trans. Netw. Service Manag.*, vol. 15, no. 1, pp. 83–96, Mar. 2018.
7) Y. Zhang, P. Calyam, S. Debroy, and S. S. Nuguri, "Social plane for recommenders in network performance expectation management," *IEEE Trans. Netw. Service Manag.*, vol. 15, no. 1, pp. 97–111, Mar. 2018.
8) F. Yan, Y. He, O. Ruwase, and E. Smirni, "Efficient deep neural network serving: Fast and furious," *IEEE Trans. Netw. Service Manag.*, vol. 15, no. 1, pp. 112–126, Mar. 2018.
9) M. Suzuki, T. Kitahara, S. Ano, and M. Tsuru, "Group mobility detection and user connectivity models for evaluation of mobile network functions," *IEEE Trans. Netw. Service Manag.*, vol. 15, no. 1, pp. 127–141, Mar. 2018.
10) D. B. Prats, J. L. Berral, and D. Carrera, "Automatic generation of workload profiles using unsupervised learning pipelines," *IEEE Trans. Netw. Service Manag.*, vol. 15, no. 1, pp. 142–155, Mar. 2018.

REFERENCE

[1] F. De Turck, P. Chemouil, R. Boutaba, M. Yu, C. E. Rothenberg, and K. Shiomoto, "Guest editors' introduction: Special issue on big data analytics for management," *IEEE Trans. Netw. Service Manag.*, vol. 13, no. 3, pp. 578–580, Sep. 2016.

Giuliano Casale received the Ph.D. degree in computer engineering from Politecnico di Milano, Italy, in 2006. He joined the Department of Computing, Imperial College London, U.K., in 2010, where he is currently a Senior Lecturer in modeling and simulation. He was a Scientist with SAP Research, U.K. and as a Consultant in the capacity planning industry. He teaches and does research in performance engineering, cloud computing, and Big data, topics on which he has published over 120 refereed papers. He has served on the technical program committee of over 80 conferences and workshops and as the co-chair for conferences in the area of performance engineering such as ACM SIGMETRICS/Performance. He is a member of the IFIP WG 7.3 group on Computer Performance Analysis and since 2015 he has been serving in the ACM SIGMETRICS Board of Directors.

Yixin Diao received the Ph.D. degree in electrical engineering from Ohio State University, Columbus, OH, USA. He is currently a Research Staff Member with IBM T. J. Watson Research Center, Yorktown Heights, NY, USA. He has authored over 80 papers in systems and services management and has coauthored the book entitled *Feedback Control of Computing Systems*. He was a recipient of several Best Paper Awards from the IEEE/IFIP Network Operations and Management Symposium, the IFAC Engineering Applications of Artificial Intelligence, and the IEEE International Conference on Services Computing. He is an Associate Editor for the IEEE TRANSACTIONS ON NETWORK AND SERVICE MANAGEMENT and the *Journal of Network and Systems Management*.

Marco Mellia received the Ph.D. degree in electrical engineering from Politecnico di Torino, Italy. He holds a position as Associate Professor with Politecnico di Torino, Italy. His current research interests are in the area of traffic monitoring and analysis, in cyber monitoring, and Big Data analytics. He has co-authored over 250 papers published in international journals and presented in leading international conferences. He was a recipient of the IRTF ANR Prize at IETF-88 and the Best Paper Award at IEEE P2P'12, ACM CoNEXT'13, and IEEE ICDCS'15. He is part of the editorial board of the ACM/IEEE TRANSACTIONS ON NETWORKING, the IEEE TRANSACTIONS ON NETWORK AND SERVICE MANAGEMENT, and the *ACM Computer Communication Review*.

Rajiv Ranjan received the Ph.D. degree in computer science and software engineering from the University of Melbourne, Australia, in 2009. He has been a Reader (Associate Professor) of computing science with Newcastle University since 2015. He is an Internationally Renowned Researcher in the areas of cloud computing, Internet of Things (IoT), and big data. He has authorship of about 225 peer-reviewed scientific publications, including publications in the IEEE TRANSACTIONS ON PARALLEL AND DISTRIBUTED SYSTEMS, the IEEE TRANSACTIONS ON COMPUTERS, the *Journal of Computer and System Sciences*, *ACM Computing Surveys*, and IEEE/ACM World Wide Web conference, PVLDB. He is the Inventor of CloudSim which is the world's most adopted and cited (over 2800 Google scholar citations since 2011) distributed systems simulation and benchmarking framework. Along with researchers from Chinese Academy of Sciences, he invented G-Hadoop—the first-ever framework to support big data processing across distributed datacentres. G-Hadoop successfully integrated, managed, and processed 2.36 PB of remote sensing big data across eight satellite datacentres, and holds the current world record for remote sensing big data management and processing. He is one of the world's most cited computer scientists (top 0.09% of 2 million researchers) with over 9950 citations, an H-index of 42, and G-index of 94 (ref: Google Scholar). His Web of Science (H-index of 21 and over 2550 citations) and Scopus (H-index of 25 and over 4600 citations) indices are also internationally leading.

Nur Zincir-Heywood received the Ph.D. degree in computer science and engineering from Ege University, Turkey, in 1998. She is a Full Professor of computer science with Dalhousie University, Canada. Her research interests include data driven techniques for cyber security and network management, topics on which she has published over 200 fully reviewed papers. She was a recipient of several best paper awards as well as the supervisor for recipient of the IFIP/IEEE IM 2013 Best Ph.D. Dissertation Award in Network Management and the 2017 Women Leaders in the Digital Economy Award and a member of the ACM. She is on the editorial board of the IEEE TRANSACTIONS ON NETWORK AND SERVICE MANAGEMENT and is the Technical Program Co-Chair of IFIP/IEEE Traffic Measurement and Analysis Conference 2018. She has been a Co-Organizer for the IEEE/IFIP International Workshop on Analytics for Network and Service Management since 2016.