

A note on software tools and technologies for delivering smart media-optimized big data applications in the cloud

Rajiv Ranjan¹ · Dimitrios Georgakopoulos² · Lizhe Wang³

Published online: 23 August 2015
© Springer-Verlag Wien 2015

Welcome to the special issue of Springer Computing on Software Tools and Technologies for Delivering Smart Media-Optimized Applications in the Cloud. This special issue deals with the intersection of big media data and its needs for exploiting elastic cloud computing services for efficiently processing and analysing such big data to support emerging media-optimised applications.

Big data and cloud computing are big news around the world. For example, studies by CISCO [1] and IBM [2] found we now generate 2.5 quintillion bytes of data per day. According to IDC [3], the data universe is set to explode to 40 yottabytes by 2020—5,200 gigabytes for every person on earth. Together with the rapidly increasing volume of data being generated, opportunities for analysing such data to distil knowledge that can help solve problems that benefit individuals, the industry and government, and our entire society presents a monumental technical challenge. This special issue focuses on a major and less exploited segment of big data: multimedia content. Major industry players, such as CISCO [4], predict that by 2016, 90 % of internet traffic will be

✉ Lizhe Wang
lizhe.wang@computer.org

Rajiv Ranjan
raj.ranjan@csiro.au
<http://www.rajivranjan.net>

Dimitrios Georgakopoulos
dimitrios.georgakopoulos@rmit.edu.au

- ¹ Digital Productivity and Services Flagship, CSIRO, GPO Box 664, ACT 2601 Woodstock, Australia
- ² School of Computer Science and Information Technology, Royal Melbourne Institute of Technology, 414-418, Swanston Street, Melbourne VIC 3000, Australia
- ³ School of Computer Science, China University of Geosciences (Wuhan), No. 388 Lumo Road, Wuhan 430074, P.R. China

multimedia content (3D images, high resolution video, and audio), we refer to as big media. In addition to established markets in media on demand entertainment, advertising and social media content sharing that are currently generate and process such big media, emerging demand for management, analysis and distribution of big media is also reported in the domains of healthcare, aged-care, security, and education (just to name a few). For example, in healthcare [5] when patients return home after a surgery or a treatment will receive follow up advice and education via live as well as archived video. This may include video instructions about how to change the dressing on a healing wound, how to brush their teeth after having braces installed, or how to exercise to improve mobility. Providing services around the management and distribution of content produced by advanced medical imaging is another example form the medical domain. In the aged-care domain, carers can utilize real-time or recorded video feeds from their patients' homes to detect clinical signs and indicators such as skin colour, mood, affect, and to determine whether a patient is utilizing provided devices and medications appropriately. In the physical security area, thousands or millions of video and audio streams from increasing numbers of security cameras, and at higher resolutions, can be analysed in real time to detect threats and generate alerts as need to protect citizens, homes and critical infrastructure. Finally, a new wave of innovation in the education sector involves satisfying the demand for big media management and that will provide students with the opportunity to access lectures form around the world and perform lab work in virtual laboratories without leaving their home even in remote areas 100s or 1000s of kilometres away and from specialised teachers and lab facilities that are not available locally.

Over the past few years, cloud computing [6, 7] has emerged as the latest and most dominant solution for processing big data by offering both hardware and data analysis software resources as virtualization-enabled services. Cloud computing providers such as Amazon Web Services and Microsoft Azure currently provide application owners the option of deploying their applications over a network of a virtually infinite resource pool with practically no up-front capital investment and with operating cost proportional to the actual use (i.e., implementing a pay-as-you-go model). An ever increasing number of cloud vendors now offer resources such as hardware (CPUs, GPUs, storage, and networks); software infrastructure for data management and analysis (e.g., databases, web servers, stream-processing systems, and data-mining packages); and applications (e.g., email, sales applications, and social networks) as infrastructure as a service (IAAS), platform as a service (PAAS), and software as a service (SAAS), respectively. While Cloud computing can provide virtually infinite computing resources with an option to pay-per-use and as needed, and allow users to access media application from anywhere at any time, and by using devices ranging from smart phones to desktop computers, it does not (yet) provide an complete solution for hosting applications that exploit big media.

Developing a software tools and technologies that can support end-to-end lifecycle operations of different multimedia content-driven applications on cloud infrastructures (while ensuring the QoS targets) remains a challenging research problem. This challenge mainly arises from the uncertainty posed by application dynamic and poorly understood workloads (streaming vs. static content), resources capacity demands (e.g., bandwidth, memory, storage, and processors), failures (e.g., failure of a network link),

access patterns (e.g., number of end-users and location), different device types (e.g., mobile phone, laptop, and smart TV), indexing needs (e.g., text and content based), and different network types (e.g. wired and wireless). For example, in the health, age care, physical security and education application scenarios discussed earlier it is clear that the following requirements must be satisfied by a media-optimised cloud: (i) hundreds of petabytes of multimedia content will be generated in real-time as well as offline which will be required to be efficiently processed (stored, distributed, and indexed with a schema and semantics) in a way that does not compromise end-users' Quality of Service (QoS) in terms of content availability, content search delay, content analysis delay, and the like and (ii) the information retrieved from content will be consumed and processed in many places concurrently and collaboratively. Existing solutions and in particular media-related cloud infrastructures service that store, distribute, and index hundreds of petabytes of multimedia content fall short of this challenge or do not exist.

To address these challenges, this special issue compiles a number of excellent technical contributions that significantly advance the state-of-the-art in the areas of media-optimised cloud services and related applications. The first paper [8] in this special issue by Yan Ma, et al., titled "Parallel Programming Templates for Remote Sensing Image Processing on GPU Architectures: Design and Implementation" explores the use of specialized hardware, more specifically Graphic Process Units (GPUs), for image analysis in remote sensing image processing applications. In particular, this paper explores the design of a reusable GPU-based remote sensing image parallel processing model and a set of parallel programming templates, which provide programmers with a more simple and effective way for programming parallel remote sensing image analysis algorithms.

A related application domain is video surveillance for providing physical security and the paper [9] by Zheng Xu, et al., titled "Semantic Enhanced Cloud Environment for Surveillance Data Management using Video Structural Description" deal directly with the problem of increasing number of surveillance video in smart cities, which creates enormous difficulties for searching and analysing such content, which is typically stored in the cloud, to detect intrusion detection and aid criminal investigation. The paper proposes a semantic-based cloud environment to facilitate search and analysis of surveillance video data.

MapReduce is one of the most important techniques, and currently the preferred choice of cloud providers, for providing cloud-based data analysis services. The paper [10] titled, "Enabling Collaborative MapReduce on the Cloud with A Single-Sign-On Mechanism" by Jiaqi Zhao, Jie Tao, and Achim Streit deal with the problem of eliminating the need of a cloud cluster manager to support the deployment of MapReduce for content analysis in the Cloud. This was needed until now because the individual virtual machines that are created on the Cloud need to be organized into clusters in order to be capable of running a MapReduce application. To solve this problem the paper introduces a software framework for individual virtual machines to execute a MapReduce application in a parallel/collaborative way without the necessity of installing a software package for cluster management. The proposed framework includes a Single-Sign-On (SSON) mechanism that enables the remote access to the individual machines.

The next paper [11], titled "A Hybrid Cloud-P2P Architecture for Multimedia Information Retrieval on VoD Services", by Vladimir Rocha et al., proposes a novel Cloud

Computing and Peer-to-Peer (P2P) hybrid architecture for multimedia information retrieval for providing VoD services that support random seeking while providing scalability and efficiency. They propose a novel architecture comprises Cloud and P2P layers. The Cloud layer provides for video segment metadata retrieval, using ontologies to improve the relevance of the retrieved information and for distributing the metadata structures among cloud nodes. The P2P layer is responsible for finding peers that store the physical location of a segment. The paper uses simulations to show that this architecture can provide the desired system scalability and performance, while avoiding the use of powerful centralized servers

The issue of improving cloud data storage for managing media content is explored by the paper [12] titled “On Construction of a Distributed Data Storage System in Cloud”, by Chao-Tung Yanga, et al. This paper focuses on extending and enhancing the functionality of the Hadoop distributed system (HDFS). The resulting distributed cloud storage system implements a node and network resource monitor that is used to optimize the virtualization system and data storage system.

The next paper [13] titled, “A Two-Stage Approach for Task and Resource Management in Multimedia Cloud Environment”, by Biao Song, et al. focuses on cloud resource management for multimedia services, and in particular how to provide the quality of service (QoS) needed by multimedia applications. The paper introduces a queuing based approach for task management and heuristic algorithms for resource management are proposed. The paper uses simulations to show that the proposed solutions provide better performance than other existing state-of-the-art approaches.

The following paper focuses on media delivery optimization in software defined networks. In particular, the paper [14] “Adaptive Scheduling Algorithm for Media-Optimized Traffic Management in Software Defined Networks” by Florin Pop, et al., proposes novel adaptive scheduling strategies that improve the resource utilization. The paper also includes extensive simulation results that show that the proposed scheduling strategies can achieve the described goals.

The paper [15] titled “WebRTC-based Video Conferencing Service for Tele-health”, by Julian Jang-Jaccard, et al., investigates the area of tele-health and describes the design of Real-Time Communication (WebRTC)-based video conferencing systems by enabling web browsers with Real-Time Communications (RTC) capabilities via simple JavaScript APIs. The paper also presents the authors’ experience in the use of this system as a part of the largest tele-home monitoring project that has been carried out in Australia.

The issues of multimedia service selection based on the service-provided QoS are the topic of the paper [16] titled “Weighted Principal Component Analysis-based Service Selection Method for Multimedia Services in Cloud”, by Lianyong Qi at al. The paper proposes a multimedia service selection method based on weighted Principal Component Analysis, called Weighted PCA-based Multimedia Service Selection Method (W_PCA_MSSM). The paper shows that W_PCA_MSSM can reduce the number of QoS criteria for service evaluation, and can eliminate the correlations between different QoS criteria. The W_PCA_MSSM effectiveness is validated with experiments using a real-world service QoS data set.

The domain of e-learning is covered in the last paper [17] in this special issue that describes a study involving “A Hybrid Intelligence-Aided Approach to

Affect-Sensitive E-learning”, by Jingying Chen, et al. This paper proposes a hybrid intelligence-aided approach to affect-sensitive e-learning. The development of a multi-modal system that incorporates affect recognition and intervention to improve the learner’s learning experience is also presented.

Our hope is that this special issue will shed light to this relatively unexplored space at intersection between big media and cloud computing that you will find the articles of this special issue to be informative and useful to your research.

References

1. Pepper R, Garrity J. The internet of everything: how the network unleashes the benefits of big data. CISCO. <http://blogs.cisco.com/wp-content/uploads/GITR-2014-Cisco-Chapter.pdf>
2. Bringing Big Data to the Enterprise, IBM, <http://www-01.ibm.com/software/in/data/bigdata/>
3. Gantz J et al. The Digital Universe in 2020: Big Data, Bigger Digital Shadows, and Biggest Growth in the Far East, IDC IVIEW, Sponsored by EMC Corporation
4. Tomorrow’s Internet Today, CISCO, Available at: http://www.cisco.com/c/dam/en/us/products/collateral/routers/asr-9000-series-aggregation-services-routers/brochure_tomorrows_internet_today.pdf
5. Shah T, Rabhi F, Ray P (2015) Investigating an ontology-based approach for big data analysis of inter-dependent medical and oral health conditions. *Cluster Comp* 18(1):351–367. doi:10.1007/s10586-014-0406-8
6. Ranjan R, Benatallah B, Dustdar S, Papazoglou M (2014) Cloud Resource Orchestration programming: overview, issues and directions. *IEEE Internet Computing*, IEEE Computer Society. (Accepted September 2, 2014)
7. Armbrust M et al (2010) A view of Cloud Computing, *Communications of the ACM Magazine*, Volume 53, Issue 4. ACM Press, New York, USA, 2010, pp 50–58
8. Ma Y, Chen L, Liu P, Lu K (2014) Parallel programming templates for remote sensing image processing on GPU architectures: design and implementation. *Spring J Comp*. doi:10.1007/s00607-014-0392-y
9. Xu Z, Mei L, Liu Y, Hu C, Chen L (2014) Semantic enhanced cloud environment for surveillance data management using video structural description. *Spring J Comp*. doi:10.1007/s00607-014-0408-7
10. Zhao J, Tao J, Streit A (2014) Enabling collaborative mapreduce on the cloud with a single-sign-on mechanism. *Springer Journal of Computing*, pp. 1-18. doi:10.1007/s00607-014-0390-0
11. Rocha V, Kon F, Cobe R, Wassermann R (2014) A Hybrid Cloud P2P Architecture for Multimedia Information Retrieval on VoD Services. *Springer Journal of Computing*, Accepted September 2014, doi:10.1007/s00607-014-0428-3
12. Yang CT, Shih W, Huang C, Cheng F, Chu W (2014) On construction of a distributed data storage systems in cloud. *Spring J Comp*. doi:10.1007/s00607-014-0399-4
13. Song B, Hassan M, Alamri A, Alelaiwi A, Tian Y, Pathan M, Almogren A (2014) A two stage approach for task and resource management in multimedia cloud environment. *Spring J Comp*. doi:10.1007/s00607-014-0411-z
14. Pop F, Dobre C, Comaneci D, Kolodziej J (2014) Adaptive scheduling algorithm for media-optimized traffic management in software defined networks. *Spring J Compu*. doi:10.1007/s00607-014-0406-9
15. Jang-Jaccard J, Nepal S, Celler B, Yan B (2014) WebRTC-based video conferencing service for tele-health. *Spring J Comp*. doi:10.1007/s00607-014-0429-2
16. Qi L, Dou W, Chen J (2014) Weighted principal component analysis-based service selection method for multimedia services in cloud computing. *Spring J Comp*. doi:10.1007/s00607-014-0413-x
17. Chen J, Luo N, Liu Y, Liu L, Zhang K, Kolodziej J (2014) A hybrid Intelligence-aided approach to affect-sensitive e-learning. *Spring J Comp*. doi:10.1007/s00607-014-0430-9