



Advances in big data programming, system software and HPC convergence

Ching-Hsien Hsu¹ · Geoffrey Fox² · Geyong Min³ · Sugam Sharma⁴

© Springer Science+Business Media, LLC, part of Springer Nature 2018

1 Introduction

In the era of big data, developing modern computing systems and system software that can scale to massive amounts of data becomes a key challenge to both researchers and practitioners. Scalability in distributed system usually means that the performance of a system should increase proportionally with the increase in resources. However, this is not sufficient in the big data era. The system should be designed in a way so that all the five Vs of big data can be tackled. Driven by this insight, this special issue aims at presenting the current state-of-the-art research and future trends on various aspects of big data programming and system software techniques for big data processing and attempts towards building highly adaptive big data systems that can automatically adapt their behaviours to the amount of available resources. The major subjects cover methodologies, modelling, analysis and newly introduced applications. Besides the latest research achievements, this special issue also covers innovative commercial data management systems, innovative commercial applications of big data technology and experience in applying recent research advances to real-world problems.

Papers selected for this special issue represent recent progresses in the field, including works on big data systems and optimizations, novel techniques for big data

✉ Ching-Hsien Hsu
chh@cs.ccu.edu.tw

Geoffrey Fox
gcfexchange@gmail.com

Geyong Min
G.Min@exeter.ac.uk

Sugam Sharma
sugam.k.sharma@gmail.com

¹ National Chung Cheng University, Minxiong, Taiwan

² Indiana University, Bloomington, USA

³ Exeter University, Exeter, UK

⁴ Iowa State University, Ames, USA

analytics, sustainable architectures and applications. All of these papers not only provide novel ideas and state-of-the-art techniques in the field, but also stimulate future research in the sustainable computing towards big data and HPC convergence.

2 Systems and optimizations

Heterogeneous multiprocessor platforms have been constantly adapted to the varying computational requirements of different kinds of tasks. In real-time systems, a multiprocessor platform can duplicate a task and execute copies on different processors. The paper by Weizhe Zhang, Yao Hu, Hui He, Yawei Liu and Allen Chen entitled “Linear and Dynamic Programming Algorithms for Real-Time Task Scheduling with Task Duplication” analyses the expandability of two periodic real-time task scheduling algorithms, dynamic programming algorithm and linear programming algorithm, with particular focus on their operating efficiency and stability. According to the two characteristics of the problem, i.e. the periodic real-time task with duplication and heterogeneous platform, the authors designed linear and dynamic programming algorithms to simulate the task scheduling process. The findings reveal that the linear programming algorithm is superior to the dynamic programming algorithm in terms of expandability.

In order to improve the computational efficiency and expand the scale of calculation, the many-task computing (MTC) is introduced into the distributed hydrological modelling. In such model, the computing should be reasonably split into many tasks as small as possible. The paper by Shasha Yue, Yan Ma, Lajiao Chen, Yuzhu Wang and Weijing Song entitled “Dynamic DAG scheduling for many-task computing of distributed eco-hydrological model” proposes a parallel scheme for the computing of distributed multi-process hydrological model. The computing of the distributed hydrological model is split into tremendous amount of small tasks that are directly dispatched into the multiple nodes of cluster through traditional local resource manager. The proposed task model can explicitly and completely express the various properties and independencies of tasks. The parallel scheme is tested with an optimality-based distributed eco-hydrological model. It is demonstrated that this approach provides efficient computing performance.

Metadata management is an essential part in Apache Hadoop. Performing optimization of metadata accesses enhances big data storing, processing and analysing, especially in multi-tenant environments. The paper by Minh Chau Nguyen, Heesun Won, Siwoon Son, Myeong-Seon Gil and Yang-Sae Moon entitled “Prefetching-based Metadata Management in Advanced Multi-tenant Hadoop” proposes a novel approach to improve the performance of metadata management for Hadoop in the multi-tenant environment. They present a formal algorithm to apply the prefetching mechanism into the Hadoop system and perform the actual implementation on a recent Hadoop system. Experimental results show that the proposed approach can enable the high performance for metadata management as well as maintain advanced multi-tenancy features.

The sustained improvement in the computing capability of high-performance computing (HPC) systems has enabled scientific applications to perform more complex

analysis and process larger datasets. These data-intensive applications require substantial I/O bandwidth to access the exponentially growing amount of data. The paper by Jie Yu, Guangming Liu, Wenrui Dong and Xiaoyong Li entitled “WatCache: A Workload-aware Temporary Cache on the Compute Side of HPC Systems” presents a workload-aware node allocation method to allocate fast storage devices to jobs according to their I/O requirements and merged the devices of the jobs into separate temporary cache spaces. They implemented a metadata caching strategy that reduces the metadata overhead of I/O requests to improve the performance of small I/O. WatCache is tested through a series of I/O benchmarks and applications on hundreds of compute nodes, which clearly indicates that it brings significant performance improvements, especially on small I/O, as well as rank 0 and all-rank I/O.

The increasingly rapid growth of the large, complex, heterogeneous and multidimensional big data in recent years has raised concerns about capacity of the existing data management systems. The paper by Sugam Sharma and Shashi Gadia entitled “Expanding ParaSQL for Spatio-Temporal (Big) Data” attempts to expand the bandwidth of the querying system of a robust dimensional data model. The model is one of the few efficient data models that have the capability of handling heterogeneous dimensions, ordinary, temporal and spatio-temporal, of the big data in a uniform way. The proposed data model is well equipped with a syntactically rich and powerful SQL-like querying system, called ParaSQL. The ParaSQL is syntactically rich and very easy to further expand its horizon and syntactically answer more complex data-driven questions.

3 Big data analytics

While software gives us convenience, it also causes many problems. Various research efforts are needed to defend against software plagiarism, attacks using malware/software and so on. Software similarity calculation can be applied to various fields such as software plagiarism detection and malware analysis. The paper by Tae Guen Kim, Yeo Reum Lee, Boo Joong Kang and Eul Gyu Im entitled “Binary Executable File Similarity Calculation Using Function Matching” proposes a similarity calculation method for binary executable files, based on function matching techniques. Experimental results show that similarity accuracy of the proposed binary-based similarity calculation method is similar to those of a well-known source code-based method. In addition, the authors performed an experiment of clustering malware families.

Automatically linking records from different data sources become more and more important in many big data applications. A crucial step of integrating data from multiple sources is detecting and eliminating duplicate records. The paper by Chenxiao Dou, Yi Cui, Daniel Sun, Raymond Wong, Muhammad Atif, Guoqiang Li and Rajiv Ranjan entitled “Unsupervised Blocking and Probabilistic Parallelisation for Record Matching of Distributed Big Data” studies the problem of blocking records and parallelised a density-based blocking algorithm. Compared to the traditional blocking techniques, the proposed method has a small loss of recall and reduction ratio but

good efficiency gain. Performance comparisons demonstrated the superiority of this approach in terms of reduction ratio, recall and running time.

Data mining is to extract the previously unknown and potentially useful information from a large database. It is the core process of the knowledge discovery of database. The association rule mining is one of the most important techniques in data mining. The paper by Rui Zhang, Wenguang Chen, Tse-Chuan Hsu, Hongji Yang, and Yeh-Ching Chung entitled “ANG - A Combination of Apriori and Graph Computing Techniques for Frequent Itemsets Mining” presents how to use the graph computing method to perform frequent itemset mining. The authors discussed the advantages and disadvantages of Apriori and the graph computing method when applied them to frequent itemset mining. Based on the discussions, we have proposed a hybrid method, by taking the advantages of Apriori and the graph computing method for frequent itemset mining. The experimental results show that the proposed technique can predict the switching point accurately and outperforms Apriori, DHP and the graph computing method for all test cases.

Under the circumstance of big data, traditional storage systems face the big challenge of energy consumption. Switching some storage nodes, which do not experience workloads, to a low-power state is a typical approach to reduce the consumption of energy. The frequently accessed data are stored into the active group which maintains the nodes in an active state to offer service, and the cold data accessed infrequently are stored into the low-power group. The paper by Cheng Hu and Yuhui Deng entitled “Aggregating Correlated Cold Data to Minimize the Performance Degradation and Power Consumption of Cold Storage Nodes” proposes to aggregate and store the correlated cold data in the same cold node within the low-power group. In contrast with concentrating the frequently accessed files of the hot nodes, the correlated cold files are aggregated to the relevant cold node. Experimental results demonstrate that this method effectively reduces the energy consumption, while maintaining system performance at an acceptable level in contrast to some state-of-the-art methods.

4 Big data applications

In the era of big data, an increasing amount of image data is being generated on the Internet through daily social communication. Most real data are high dimensional, such as image data. High-dimensional data significantly increases the time and space requirements and also brings the curse of dimensionality problem. The paper by Shang-guang Wang and Chuntao Ding entitled “Appropriate Points Choosing for Subspace Learning over Image Classification” proposes an appropriate points choosing method which can build a margin between samples if they belong to the different classes. It chooses points to construct inter-class and intra-class graphs in different ways, and as a result, it could establish relationship between two remote samples and find a good projection matrix for them. The experimental results show that the proposed method can be very effective for data classification.

Medical images such as MRI, CT and ultrasound are subject to various types and degrees of noise, which might be happening during transmission and acquisition. The availability of noise in the medical images has direct or indirect influence that compli-

cates the diagnosis, analysis and treatment process timely. The paper by Worku Jifara, Feng Jiang, Seungmin Rho, Maowei Cheng and Shaohui Liu entitled “Medical Image Denoising Using Convolutional Neural Network: A Residual Learning Approach” proposes an efficient deep convolutional neural network model for denoising of medical images for small training dataset. The proposed medical image denoising model was trained with small dataset and achieves competitive results. Experimental results demonstrate that, the performance of the proposed medical image-denoising model is better than the existing medical image-denoising techniques both in peak to signal-to-noise ratio (PSNR) and structural similarity (SSIM).

5 Conclusions

All of the above papers address either technical issues in big data platforms or information modelling or propose novel application models in the various big data and HPC fields. They also trigger further related research and technology improvements in application of sustainable computing. Honourably, this special issue serves as a landmark source for education, information and reference to professors, researchers and graduate students interested in updating their knowledge about or active in big data architecture and sustainability, and novel application models for big data and HPC convergence and systems.

This special issue of The Journal of Supercomputing covers different aspects of the problem, both from the theoretical to practical side. After a large open call, an international editorial committee selected six research papers. Each paper was reviewed by at least three reviewers.

The guest editor would like to express sincere gratitude to the EiC of JoS, for giving the opportunity to prepare this special issue. In addition, we are deeply indebted to numerous reviewers for their professional effort, insight and hard work put into commenting on the selected articles that reflect the essence of this special issue. Last but not least, we are grateful to all authors for their contributions and for undertaking two-cycle revisions of their manuscripts, without which this special issue could not have been produced.