Scalable Storage Systems and High-Performance Applications

Jesus Carretero · J. Daniel Garcia

Published online: 1 January 2010
© Springer Science+Business Media, LLC 2009

Data management is becoming a more and more important issue in every kind of computing system. High-performance computing, data visualization, multimedia and data mining systems are examples of application areas where input/output and storage are critical issues for performance. A key challenge in those areas is how to transfer large amounts of data in and out of large-scale systems, like clusters, Grids and peer-to-peer systems. Another key point is how to transfer large amount of data in a timely manner. Keeping this in mind, there is a need to identify and solve barriers and workarounds in the storage and input/output fields.

This special issue brings together five research papers related to different topics of data acquisition and management and the associated scalability issues in high-performance distributed systems. The papers have been selected among those presented at the successful “International Workshop on Scalable Data Management Application and Systems” (SDMAS 2008) which was held in July, 2008 in Las Vegas, Nevada. The workshop was held within the well-known conference on “Parallel and Distributed Processing Techniques and Applications”.

In “Enabling the interactive display of large medical volume data sets by multi-resolution bricking”, Gordon Clapworthy et al. present an approach to interactive out-of-core volume data exploration that has been developed to augment the existing capabilities of the LhpBuilder software, a core component of the European project LHD. They have implemented a hierarchical, bricked, partition-based, out-of-core strategy to balance the usage of main and external memories. A new indexing scheme

J. Carretero (✉)
Avda. Universidad 30, 28911, Leganes, Spain
e-mail: jcarrete@inf.uc3m.es

J.D. Garcia
Avda. Universidad Carlos III 22, 28270, Colmenarejo, Spain

is introduced, which permits the use of a multiresolution bricked volume layout with minimum overhead and also supports fast data compression.

In “Parallel failure recovery techniques in cluster-based media servers”, Joahyung Lee and Inbum Jung propose a new failure recovery mechanism for large-scale video-on-demand (VOD) services using clusters based on the pipeline computing concept. The proposed method not only distributes the internal network traffic generated from the recovery operations but also utilizes the CPU time available in the back-end nodes. Even if a back-end node fails, the proposed method provides continuous streaming media services within a short MTTR value as well as more QoS streams than the existing method.

In “New techniques for simulating high performance MPI applications on large storage networks”, Alberto Nuñez et al. propose new techniques to analyze the behavior, the performance, and specially the scalability of High-Performance Computing applications on different computing architectures. The paper presents a new simulation framework, called SIMCAN, for HPC architectures. The main characteristic of the proposed simulation framework is the ability to be configured for simulating a wide range of possible architectures that involve any number of I/O components.

In “Scalable parallel word search in multicore/multiprocessor systems”, Frank Drews, Jens Lichtenberg, and Lonnie Welch present a parallel algorithm for fast word search to determine the set of biological words of an input DNA sequence. The algorithm is designed to scale well on state-of-the-art multiprocessor/multicore systems for large inputs and large maximum word sizes. The proposed algorithm achieves three major goals: to impose a high degree of cache locality, to largely reduce the need for data access locking, and to enable an even distribution of the overall processing load among multiple threads.

Finally in “SENFIS: A Sensor Node File System for increasing the scalability and reliability of Wireless Sensor Networks applications”, Soledad Escolar et al. propose a novel file system for sensor nodes, which addresses both scalability and reliability concerns. SENFIS can be mainly used in two broad scenarios. First, it can transparently be employed as a permanent storage for distributed TinyDB queries, in order to increase the reliability and scalability. Second, it can be directly used by a WSN application for permanent storage of data on the WSN nodes.

We hope that the reader will find the papers in this issue both informative and inspiring. The papers of this special issue are extended versions of a selection of papers from the SDMAS 2008 workshop. A new review process was performed using different reviewers from those who participated in the workshop. As a result of this high quality process, some initially invited papers were rejected. We would like to thank the authors for their contributions. We would also like to thank the reviewers from the original SDMAS workshop and from the special issue for their dedication and effort in providing timely and thorough reviews. This special issue would not have been possible without their hard work.