Editorial

Guest Editors' Introduction: Special Issue on Supercomputing: Applications and Technologies

Andrei Tchernykh [1], René Luna-García [2], and Juan Manuel Ramírez-Alcaraz [3]

[1] Computer Science Department, CICESE Research Center, Ensenada, BC, México.
chernykh@cicese.mx
[2] Centro de Investigación en Computación - IPN, México City, México. René Luna-García
rene.luna.garcia@gmail.com
[3] Colima University, C.P. 28040. Colima, Col., México, jmramir@ucol.mx

Abstract. This special issue contains a selection of papers that had been presented at the International Supercomputing Conference in Mexico (ISUM). The papers cover several important topics and present some of the directions in rapidly expanding area of research and development related with applications and technologies of supercomputing. Authors show the use of supercomputers to solve complex modeling and simulation problems in a spectrum of disciplines such as: Climate modeling, Cloud computing, Dynamic analysis, Finite element method, Gaussian transformation, GPU, Image security, Load balancing, LU decomposition, Molecular biology, Numerical optimization, Parallel numerical computing, Performance analysis, Rendering and animation, Stochastic optimization, Thermal and stress analysis. The objective of this special issue is to publish and overview recent trends in the interdisciplinary area of parallel and distributed computing, applications and technologies. We hope that the set of selected papers provides the community with a better understanding of the current research areas, introducing new research, development, and deployment efforts in running applications on high performance computing infrastructures. The ISUM succeeded in highlighting many of the significant topics and trends within the area of concern. Some of them are reflected in the articles selected for this issue. 12 papers from Argentina, Germany, Mexico, and Spain were selected from 78 papers presented at the conference, which had an overall acceptance rate of 16%..

Keywords. ISUM, CyS.

Introducción de Editores Invitados a Número Especial Sobre Supercomputación: Aplicaciones y Tecnologías

Resumen. Esta edición especial contiene una selección de los trabajos presentados en el International Supercomputing Conference in México (ISUM). Los artículos cubren varios de los temas más importantes que presentan una rápida expansión en áreas de investigación y desarrollo, relacionadas con aplicaciones y tecnologías de supercomputación. Los autores muestran el uso de supercomputadoras para resolver problemas complejos de modelado y simulación en disciplinas tales como: Modelación del clima, Cómputo en la nube, Análisis dinámico, Método de elemento finito, Transformación de Gauss, GPU, Imágenes de seguridad, Balanceo de carga, Descomposición LU, Biología molecular, Optimización numérica, Cómputo Paralelo numérico, Análisis de rendimiento, Visualización y animación, Optimización estocástica, Análisis térmico y esfuerzos. El objetivo de esta edición especial es publicar, así como, dar una visión general de las tendencias en áreas interdisciplinarias de la informática, aplicaciones y tecnologías paralelas y distribuidas. Esperamos que el conjunto de los trabajos seleccionados proporcione a la comunidad una mejor comprensión de las actuales áreas de investigación, desarrollos y esfuerzos de implementación de aplicaciones para su ejecución en infraestructuras informáticas de alto rendimiento. El ISUM destacó exitosamente muchos temas y tendencias importantes en diferentes áreas de interés, algunas de ellas están reflejadas en los artículos seleccionados para esta edición especial. Se seleccionaron 12 artículos procedentes de Argentina, Alemania, México y España de un total de 78. El porcentaje de trabajos aceptados para esta edición especial fue del 16%.

Palabras clave. ISUM, CyS.

1 Introduction

High performance computing (HPC) has been under many years of development, coupling with different research and application trends such as grid computing, cloud computing, green computing, etc. It has become a key technology in research and development activities in many academic and industrial areas, especially when the solution of large and complex problems copes with timing constraints. Nowadays, the design, analysis, and evaluation of scientific applications for parallel and distributed computing systems are almost a commonplace. However, to reach high performance, many obstructions that have their origin in a broad spectrum of problem areas of parallel and distributed processing, resource optimization, and algorithms for high-performance computing have to be overcome.

2 ISUM

The ISUM provides first class open forums for scientists in academia, engineers in industry and government to present and discuss issues, trends and results, which shape the high performance computing.

The scope of the conference includes, but not limited to the following topics: Cloud/Grid computing systems, Cyber-Infrastructure, Database applications and data mining, Distributed systems security, E-Science, Emerging Technologies, High-performance scientific computing, HPC applications and emerging technologies, Intelligent computing and neural networks, Interconnection networks, Modeling and simulation, Parallel/distributed algorithms, Reliability and fault-tolerance, Resource allocation and management, Scientific visualization, Software tools and environments, Task scheduling, Wireless networks and mobile computing, among others.

At this forum, the scientific community presents the technological tendencies and scientific applications serving as support to the body of knowledge in the diverse disciplines. It serves as a space to give continuity to the inter-institutional collaborations.

With a track record of four consecutive years, ISUM has served as a successful forum to promote the exchange of experiences, knowledge and research related to the use of high performance computing. It comprises plenary talks, lectures, workshops, technology expositions, poster presentations, and panel discussions. More than 300-500 participants attend the conference each year.

A number of recognized researchers from Europe, Latin America, Mexico, and United States attended the event, and shared their work through keynote talks, conference presentations, and workshops. We mention academic keynote speakers in alphabetic order: Alexandre Tkatchenko (Fritz Haber Institute, Germany), Amit Majumdar (San Diego Supercomputer Center, USA), Andrei Tchernykh (CICESE Research Center, México), Behrooz Parhami (University of California, Santa Barbara, USA), Carl Kesselman (University of Southern California, USA), Carlos Jaime Barrios (Universidad Industrial de Santander in Bucaramanga, Colombia), Glenn Bresnahan (Boston University, USA), Ian Foster (Argonne National Laboratory, USA), Jack Dongarra (University of Tennessee, USA), Jaime Klapp (Instituto Nacional de Investigaciones Nucleares, México), Klaus Ecker (Ohio University, USA), Marc Snir (University of Illinois at Urbana-Champaign, USA), Mateo Valero (Barcelona Supercomputing Center, Spain), Michael Norman (San Diego Supercomputing Center, USA), Moisés Torres Martínez (University of Guadalajara, Mexico), Nicholas Cardo (Lawrence Berkeley National Laboratory, USA), Pablo Emilio Guillén Rondón (Univesidad de los Andes, Venezuela), Rajkumar Buyya (University of Melbourne, Australia), Ramin Yahyapour (GWDG – University of Göttingen, Germany), Richard Jorgensen (Cinvestav, México), Rick Stevens (Argonne National Laboratory, USA), Thomas Sterling (Louisiana State University, USA), Uwe Schwiegelshohn (Dortmund University, Germany), Vassil Alexandrov (University of Reading, UK), William Thigpen (NASA Advanced Supercomputing Division, USA).

Some of technology keynote speakers are: Padmanabhan Iyer (Intel), Bob Anderson (Cray), Claude Paquette (Cray), Boris Cortes Silva (NetApp/EMTEC), Joshua Mora (AMD), Bill

Nitzberg (PBS), Frederick Reid (SGI), Ivan Tovar (SSC), and Rajesh Suckramani (IBM).

This event is held by coordination of the ISUM National Organizing Committee formed by various institutions: Centro de Investigación Científica y de Educación Superior de Ensenada (CICESE), Centro de Investigación en Matemáticas (CIMAT), Centro de Investigación y de Estudios Avanzados del IPN (CINVESTAV), Centro Nacional de Supercómputo (CNS), Corporación Universitaria para el Desarrollo de Internet (CUDI), Instituto Politécnico Nacional (IPN), Instituto Potosino de Investigación Científica y Tecnológica (IPICYT), Instituto Tecnológico de Monterrey (ITESM), Universidad Autónoma Metropolitana (UAM), Universidad de Colima (UCOL), Universidad de Guadalajara (UDG), Universidad de Guanajuato (UGTO), Universidad de Sonora (USON), Universidad Nacional Autónoma de México (UNAM), with support from CONACYT.

Furthermore, a set of leading technology companies such as: AMD, CAPA4, CGG, Cisco Systems, Dell, EMC, Estratel, Global Computing, Hewlett Packard (HP), IBM, Intel Corporation, IQ Tech, Lufac, Microsoft, NVIDIA, PBS Works, Red Supercomputo, RMMC, SGI, SGrupo, CNS, SUN Microsystems, Telmex, WOLFRAM, etc. provided their support as sponsors to carry out this event.

ISUM 2014 is the next event in a series of highly successful International Supercomputing Conference in Mexico, previously held as ISUM-2010 (Guadalajara, México, March 2010), ISUM 2011 (San Luis Potosí, México, March 2011), ISUM 2012 (Guanajuato, México, March 2012), ISUM 2013 (Manzanillo, México, March 2013). ISUM 2014 will be held in Ensenada, Baja California, México, March 2014.

3 In this issue

Papers in the first group deal with a broad spectrum of implementation issues, mostly those related to the using the Finite Element Method (FEM).

The paper by Cardoso-Nungaray et al., "Parallel processing strategy to solve the thermal-mechanical coupled problem applied to a 4D system using the Finite Element Method" proposes a strategy to solve the thermal-mechanical coupled problem using FEM. The authors simulate the deformation of a solid body due to internal forces provoked by temperature changes. They test the strategy on a car braking system, where a rotating disk velocity is decreased by the friction. A quantitative analysis of the stress, strain and temperatures in some points of the geometry, and qualitative analysis with visualizations of the simulation results are conducted.

The paper by González-García et al., "Load Balancing for Parallel Computations with the Finite Element Method" presents an overview of efforts to improve current techniques of load-balancing and efficiency of FEM computations on large-scale parallel machines. The authors introduce a multilevel load balancer to improve the local load imbalance. They present the current state of the art in the area, and show resource aware parallelization approaches, which lead to scalable applications on more than thousands of heterogeneous processors.

The article by García-Blanquel et al., "Parallel Adaptive Method for Selecting Points of Interest on Structures: Cranial Deformation" presents a cranial deformation simulation. The authors applied stress model using Finite Element Method, propose the boundary conditions according to the position of forces in stereotactic frame fixed on the human head. They implement parallel image segmentation algorithm using PosixThread and MPI to build a three-dimensional cranium structure, and adaptive methods for reducing the computational cost and discretization errors. For the experiments they use 134 computed tomography images.

A discrete element method (DEM), as one of numerical methods, becomes widely accepted as an effective technique of addressing engineering problems. The article by Medel et al., "Design and Optimization of Tunnel Boring Machines by Simulating the Cutting Rock Process Using the Discrete Element Method" proposes the use the DEM to build models that simulate the rock cutting process by a cutting disk, and measure the forces-hard-rock interactions essential in the design of tunnel boring machines. The authors make a performance prediction of their design

parameters that are critical in mechanical excavation.

The second group consists of four papers that deals with various issues related to the use of graphics processing units (GPUs), which highly parallel structure makes them effective for algorithms, where processing of data can be done in parallel.

The article by Lopresti et al., "Solving Multiple Queries through the Permutation Index in GPU" describes an algorithm of building the permutation index used to approximate similarity search on databases and solve many content queries at the same time. The authors evaluate the tradeoff between the answer quality and time performance of the implementation. The metric space model is used to model the similarity search problem. However, for very large metric database, it is not enough to preprocess the dataset by building an index, it is also necessary to speed up the queries by using high performance computing.

The article by Romero-Vivas et al., "Analysis of Genetic Expression with Microarrays using GPU Implemented Algorithms" presents the implementation of algorithms using Compute Unified Device Architecture (CUDA) to determine statistical significance in the evaluation of gene expression for a microarray hybridization experiments. The results are compared with respect to traditional sequential implementations, and showed speedup by a factor of around 5-30. DNA microarrays are used to simultaneously analyze the expression level of thousands of genes under multiple conditions. Massive amount of generated data makes its analysis a challenge, and an ideal candidate for massive parallel processing.

In the article by Rudomin et al., "GPU Generation of Large Varied Animated Crowds," the authors discuss methods for simulating, generating, animating and rendering crowds and their behaviors. Such a data processing is complex and costly. However, the author show that such systems can scale up almost linearly using GPU clusters and HPC.

The paper by García-Cano et al., "A Parallel PSO Algorithm for a Watermarking Application on a GPU" presents the usability, advantages and disadvantages of using CUDA for implementing a particle swarm optimization algorithm. To test

proposed solutions, a hide watermark image applications are used. The authors use the insertion/extraction algorithms, and consider two objectives: fidelity and robustness using Mean Squared Error (MSE) and Normalized Correlation (NC). These optimization criteria are evaluated by using Pareto dominance.

The paper by Gil-Costa et al., "Suffix Array Performance Analysis for Multi-core Platforms," evaluates the performance achieved by the suffix array over a 32-core platform. Suffix arrays are efficient data structures for solving complex queries in a number of applications related to text databases, for instance, biological databases. The authors propose an optimization technique to improve the use of the cache memory, by reducing the number of cache memory replacement performed every time a new query is processed.

The cloud computing is a type of Internet-based computing, where different services, such as computer infrastructure, software, applications and storage are accessed, shared as virtual resources, and delivered to low-cost consumer PC through the Internet. Dynamic resource allocation based on virtualization technology is analyzed in the paper "Performance Evaluation of Infrastructure as a Service Clouds with SLA Constraints" by Lezama-Barquet et al. The use of Service Level Agreements (SLAs) is a fundamentally new approach for job scheduling. With this approach, schedulers are based on satisfying QoS constraints. The main idea is to provide different levels of service, each addressing different set of customers. The authors developed a set of heuristics to allocate virtual resources to physical infrastructures in an efficient way. Algorithms are studied in the context of executing real workload traces available for HPC community. In order to provide performance comparison, a joint analysis of several metrics is performed. The authors show that the system can be dynamically adjusted in response to the changes in the configuration and/or the workload. To this end, the past workload within a given time interval can be analyzed to determine an appropriate parameters. The time interval for this adaptation depends on the dynamics of the workload characteristics and IaaS configuration.

In the article by González et al., "A New LU Decomposition On Hybrid GPU-Accelerated Multicore Systems", the authors introduce a theorem on the decomposition of a matrix into determinants, and the new linear transformations. They show that linear equation systems can be simultaneously solved with these new linear transformations. The authors have also proposed a modified Doolittle-Gauss-LU-Decomposition in two versions: the first one applied to the matrix and the second to the augmented matrix. The first one is a new algorithm to compute determinants in exact form, and the second is a new LU elimination process to solve linear system of equations in parallel using multicore systems with GPU.

The last paper by Heras-Sánchez et al., "Tridimensional-temporal-thematic hydro-climate modeling of distributed parameters for the San Miguel river basin" proposes a numerical analysis of diverse physiographic features and hydro-climate phenomena such as rainfall, temperature, soil-evaporation, and topography, among others The authors design the geographic database (GDB) that integrates data from multiple sources, and develop strategies for structuring, identifying, handling and improving the GDB information. An iterative digitizing and geo-referencing process is formulated to match remotely collected hydro-climate data with digital elevation model and thematic images. Digital maps, images, thematic information and spatially-referenced vector data are used together. Continuous data are represented by discrete data structures that fit to the mathematical models used to represent the physical phenomena. An increasing database size and model complexity causes difficulties to support problem-solving. High performance computing is used to overcome the computational intractability of large and complex spatial analysis.

These 12 papers cover a diverse range of technologies, modeling, simulation and computing paradigms that exemplifies the unique nature of the ISUM event.

Acknowledgements

References

1. ISUM International Supercomputing Conference in Mexico (Congreso Internacional de Supercómputo en México), http://www.isum.mx/

Andrei Tchernykh is a researcher in the Computer Science Department, CICESE Research Center, Ensenada, Baja California, Mexico. From 1975 to 1990 he was with the Institute of Precise Mechanics and Computer Technology of the Russian Academy of Sciences (Moscow, Russia). He received the Ph.D. in Computer Science in 1986. In CICESE, he is a coordinator of the Parallel Computing Laboratory. He is member of the National System of Researchers of Mexico (SNI), Level II. He leads number of national and international research projects. He served as a program committee member of professional international conferences, and as a general chair for International Conferences on Parallel Computing Systems. His main interests include scheduling, load balancing, adaptive resource allocation, scalable energy-aware algorithms, green grid and cloud computing, eco-friendly P2P scheduling, multi-objective optimization, scheduling in real time systems, computational intelligence, heuristics, meta-heuristics, and incomplete information processing**.**

Rene Luna-García is Professor at the Artificial Intelligence Laboratory of Centro de Investigación en Computación-IPN México City, México. He is member of High Altitude Water Cherenkov experiment for Cosmic Rays. He has Master degree in nuclear physics and Ph. D. in Cosmic Rays. He worked as posdoc in the DZERO-TEVATRON particle accelerator at Fermi Lab as part of diffraction group. He has more than 200 papers with the DZERO collaboration. His research interests include High Energy Physics, Cosmic Rays, modeling and simulate, scientific computing.

Juan Manuel Ramírez-Alcaraz is an associate professor at University of Colima. He received the Ph.D. in Computer Science from CICESE Research Center in 2011. His work was focused on the area of job scheduling in hierarchical computational Grids. He received his B.S. degree in Computational Systems Engineer in 1994 and his Master degree in Telematics in 2001, both from University of Colima, México. He has worked during 25 years as professor in Computer Science at University of Colima. His principal interest areas include parallel programming, job scheduling, Grid and Cloud computing.