

Contents

HPC Infrastructures and Datacenters

A Deep Learning Mapper (DLM) for Scheduling on Heterogeneous Systems	3
<i>Daniel Nemirovsky, Tugberk Arkose, Nikola Markovic, Mario Nemirovsky, Osman Unsal, Adrian Cristal, and Mateo Valero</i>	
Power Consumption Characterization of Synthetic Benchmarks in Multicores	21
<i>Jonathan Muraña, Sergio Nesmachnow, Santiago Iturriaga, and Andrei Tchernykh</i>	
Initial Experiences from TUPAC Supercomputer.	38
<i>David Vinazza, Alejandro Otero, Alejandro Soba, and Esteban Mocskos</i>	

HPC Industry and Education

romeoLAB: A High Performance Training Platform for HPC, GPU and DeepLearning	55
<i>Arnaud Renard, Jean-Mathieu Etancelin, and Michael Krajecki</i>	

GPU, Multicores, Accelerators

Analysis and Characterization of GPU Benchmarks for Kernel Concurrency Efficiency	71
<i>Pablo Carvalho, Lúcia M. A. Drummond, Cristiana Bentes, Esteban Clua, Edson Cataldo, and Leandro A. J. Marzulo</i>	
Parallel Batch Self-Organizing Map on Graphics Processing Unit Using CUDA	87
<i>Habib Daneshpajouh, Pierre Delisle, Jean-Charles Boisson, Michael Krajecki, and Nordin Zakaria</i>	
Performance Prediction of Acoustic Wave Numerical Kernel on Intel Xeon Phi Processor	101
<i>Víctor Martínez, Matheus Serpa, Fabrice Dupros, Edson L. Padoin, and Philippe Navaux</i>	

Evaluating the NVIDIA Tegra Processor as a Low-Power Alternative for Sparse GPU Computations	111
<i>José I. Aliaga, Ernesto Dufrechou, Pablo Ezzatti, and Enrique S. Quintana-Ortí</i>	

HPC Applications and Tools

Benchmarking Performance: Influence of Task Location on Cluster Throughput.	125
<i>Manuel Rodríguez-Pascual, José Antonio Moríñigo, and Rafael Mayo-García</i>	
PRIMULA: A Framework Based on Finite Elements to Address Multi Scale and Multi Physics Problems	139
<i>Alejandro Soba</i>	
FaaSter, Better, Cheaper: The Prospect of Serverless Scientific Computing and HPC	154
<i>Josef Spillner, Cristian Mateos, and David A. Monge</i>	
AccaSim: An HPC Simulator for Workload Management.	169
<i>Cristian Galleguillos, Zeynep Kiziltan, and Alessio Netti</i>	
SherlockFog: Finding Opportunities for MPI Applications in Fog and Edge Computing.	185
<i>Maximiliano Geier and Esteban Mocsos</i>	

Big Data and Data Management

IoT Workload Distribution Impact Between Edge and Cloud Computing in a Smart Grid Application	203
<i>Otávio Carvalho, Manuel Garcia, Eduardo Roloff, Emmanuell Diaz Carreño, and Philippe O. A. Navaux</i>	
Model-R: A Framework for Scalable and Reproducible Ecological Niche Modeling	218
<i>Andrea Sánchez-Tapia, Marínez Ferreira de Siqueira, Rafael Oliveira Lima, Felipe Sodré M. Barros, Guilherme M. Gall, Luiz M. R. Gadelha Jr., Luís Alexandre E. da Silva, and Carla Osthoff</i>	

Parallel and Distributed Algorithms

Task Scheduling for Processing Big Graphs in Heterogeneous Commodity Clusters	235
<i>Alejandro Corbellini, Daniela Godoy, Cristian Mateos, Silvia Schiaffino, and Alejandro Zunino</i>	

Exploring Application-Level Message-Logging in Scalable HPC Programs. . .	250
<i>Esteban Meneses</i>	
Accelerated Numerical Optimization with Explicit Consideration of Model Constraints.	255
<i>Lucia Damiani, Ariel Ivan Diaz, Javier Iparraguirre, and Anibal M. Blanco</i>	
Parallel Processing of Intra-cranial Electroencephalogram Readings on Distributed Memory Systems	262
<i>Leonardo Piñeyro and Sergio Nesmachnow</i>	
Support Vector Machine Acceleration for Intel Xeon Phi Manycore Processors	277
<i>Renzo Massobrio, Sergio Nesmachnow, and Bernabé Dorronsoro</i>	
Performance Improvements of a Parallel Multithreading Self-gravity Algorithm.	291
<i>Nestor Rocchetti, Daniel Frascarelli, Sergio Nesmachnow, and Gonzalo Tancredi</i>	
A Fast GPU Convolution/Superposition Method for Radiotherapy Dose Calculation.	307
<i>Diego Carrasco, Pablo Cappagli, and Flavio D. Colavecchia</i>	
Grid, Cloud and Federations	
Eeny Meeny Miny Moe: Choosing the Fault Tolerance Technique for my Cloud Workflow	321
<i>Leonardo Araújo de Jesus, Lúcia M. A. Drummond, and Daniel de Oliveira</i>	
Energy Aware Multiobjective Scheduling in a Federation of Heterogeneous Datacenters	337
<i>Santiago Iturriaga and Sergio Nesmachnow</i>	
Markov Decision Process to Dynamically Adapt Spots Instances Ratio on the Autoscaling of Scientific Workflows in the Cloud	353
<i>Yisel Gari, David A. Monge, Cristian Mateos, and Carlos García Garino</i>	
Experimental Analysis of Secret Sharing Schemes for Cloud Storage Based on RNS	370
<i>Vanessa Miranda-López, Andrei Tchernykh, Jorge M. Cortés-Mendoza, Mikhail Babenko, Gleb Radchenko, Sergio Nesmachnow, and Zhihui Du</i>	

Bi-objective Heterogeneous Consolidation in Cloud Computing 384
*Luis-Angel Galaviz-Alejos, Fermín Armenta-Cano, Andrei Tchernykh,
Gleb Radchenko, Alexander Yu. Drozdov, Oleg Sergiyenko,
and Ramin Yahyapour*

Scaling the Deployment of Virtual Machines in UnaCloud 399
*Jaime Chavarriaga, César Forero-González, Jesse Padilla-Agudelo,
Andrés Muñoz, Rodolfo Cáliz-Ospino, and Harold Castro*

Distributed Cosmic Ray Detection Using Cloud Computing 414
*Germán Schnyder, Sergio Nesmachnow,
and Gonzalo Tancredi*

Author Index 431

High Performance Computing

4th Latin American Conference, CARLA 2017, Buenos Aires, Argentina, and Colonia del Sacramento, Uruguay, September 20-22, 2017, Revised Selected Papers

Mocskos, E.; Nesmachnow, S. (Eds.)

2018, XIV, 432 p. 167 illus., Softcover

ISBN: 978-3-319-73352-4