

Articles

ID	Title	Authors	Paper Sessions
2	Benchmarking performance: influence of task location on cluster throughput	Manuel Rodríguez-Pascual, José A. Moríñigo, Rafael Mayo-García	High Performance Applications and Tools
4	A Deep Learning Mapper (DLM) for Scheduling on Heterogeneous Systems	Daniel Nemirovsky, Tugberk Arkose, Nikola Markovic, Mario Nemirovsky, Osman Unsal, Adrián Cristal, Mateo Valero	HPC Infrastructure and Datacenters
5	Task scheduling for processing big graphs in heterogeneous commodity clusters	Alejandro Corbellini, Alejandro Zunino, Cristian Mateos Diaz, Silvia Schiaffino, Daniela Godoy	Parallel and Distributed Algorithms
6	Eny Meeny Miny Moe: choosing the fault tolerance technique for my cloud workflow	Leonardo Jesus, Lucia Drummond, Daniel Oliveira	Grid, Cloud Computing, and Federations
7	PRIMULA: A framework based on finite elements to address multi scale and multi physics problems	Alejandro Soba	High Performance Applications and Tools
8	Energy aware multiobjective scheduling in a federation of heterogeneous datacenters	Santiago Iturriaga, Sergio Nesmachnow	Grid, Cloud Computing, and Federations
13	AccaSim: an HPC Simulator for Workload Management	Cristian Galleguillos, Zeynep Kiziltan, Alessio Netti	High Performance Applications and Tools
14	Markov Decision Process to Dynamically Adapt Spots Instances Ratio on the Autoscaling of Scientific Workflows in the Cloud	Yisel Garí, David Monge, Cristian Mateos Diaz, Carlos García Garino	Grid, Cloud Computing, and Federations
18	IoT Workload Distribution Impact between Edge and Cloud Computing in a Smart Grid Application	Otávio Carvalho, Manuel García, Eduardo Roloff, Philippe Navaux	Big Data, Data Management, and Visualization
19	Analysis and Characterization of GPU Benchmarks for Kernel Concurrency Efficiency	Pablo Carvalho, Lucia Drummond, Cristiana Bentes, Esteban Clua, Edson Cataldo, Leandro Marzulo	GPU, Multicore Architectures, and Accelerators
21	romeoLAB: a High Performance Training Platform for HPC, GPU and DeepLearning	Arnaud RENARD, Jean-Matthieu Etancelin, Michael Krajecki	HPC in Industry and Education
22	Towards Efficient Lattice-Boltzmann Method Implementation on Single-box Computers	Ezequiel Malamud	Ongoing work
23	Exploring Application-Level Message-Logging in Scalable HPC Programs	Esteban Meneses	Parallel and Distributed Algorithms
24	Experimental Analysis of Secret Sharing Schemes for Cloud Storage based on RNS	Vanessa Miranda, Andrei Tchernykh, Jorge M. Cortés-Mendoza, Mikhail Babenko, Gleb Radchenko, Sergio Nesmachnow, Zihui Du	Grid, Cloud Computing, and Federations
26	Bi-Objective Heterogeneous Consolidation in Cloud Computing	Fermin Armenta, Luis-Angel Galaviz-Alejos, Andrei Tchernykh, Gleb Radchenko, Alexander Drozdov, Oleg Sergiyenko, Ramin Yahyapour	Grid, Cloud Computing, and Federations
27	Accelerated Numerical Optimization with Explicit Consideration of Model Constrains	Lucia Damiani, Ariel I. Diaz, Javier Iparraguirre, Anibal Blanco	Parallel and Distributed Algorithms
30	Model-R: A Framework for Scalable and Reproducible Ecological Niche Modeling	Andrea Sánchez-Tapia, Marín de Siqueira, Rafael Lima, Felipe Barros, Guilherme Gall, Luiz Gadelha, Luís Alexandre da Silva, Carla Osthoff	Big Data, Data Management, and Visualization
34	Scaling the deployment of Virtual Machines in UnaCloud	Jaime Chavarriaga, César Forero-González, Jesse Padilla-Agudelo, Andrés Muñoz, Rodolfo Cáliz-Ospino, Harold Castro	Grid, Cloud Computing, and Federations
35	Parallel Batch Self-Organizing Map on Graphics Processing Unit using CUDA	Habib Daneshpajouh, Pierre Delisle, Jean-Charles Boisson, Michael Krajecki, Nordin Zakaria	Ongoing work
38	SherlockFog: finding Opportunities for MPI Applications in Fog and Edge Computing	Maximiliano Geier, Esteban Mocskos	High Performance Applications and Tools
40	Performance Prediction of Acoustic Wave Numerical Kernel on Intel Xeon Phi Processor	Víctor Martínez, Matheus Serpa, Fabrice DUPROS, Edson Luiz Padoin, Philippe Navaux	GPU, Multicore Architectures, and Accelerators
41	Evaluating the NVIDIA Tegra processor as a low-power alternative for sparse GPU computations	José Aliaga, Ernesto Dufrechou, Pablo Ezzatti, Enrique Quintana-Ortí	GPU, Multicore Architectures, and Accelerators
42	Power Consumption Analysis for Energy Characterization of Scientific Workloads in Multicores	Jonathan Muraña, Sergio Nesmachnow, Santiago Iturriaga, Andrei Tchernykh	HPC Infrastructure and Datacenters
43	Initial Experiences from TUPAC Supercomputer	David Vinazza, Alejandro Otero, Alejandro Soba, Esteban Mocskos	HPC Infrastructure and Datacenters
44	Parallel processing of intra-cranial electroencephalogram readings on distributed memory systems	Leonardo Piñeyro, Sergio Nesmachnow	Special Students Track
45	Distributed cosmic ray detection using cloud computing	Germán Schnyder, Sergio Nesmachnow, Gonzalo Tancredi	Special Students Track
47	Support Vector Machine Acceleration for Intel Xeon Phi Manycore Processors	Renzo Massobrio, Sergio Nesmachnow, Bernabe Dorronsoro	Special Students Track
48	Performance improvements of a parallel multithreading self-gravity algorithm	Nestor Rocchetti, Daniel Frascarelli, Sergio Nesmachnow, Gonzalo Tancredi	Parallel and Distributed Algorithms
50	A Fast GPU Dose Calculation Algorithm Based on Convolution/Superposition for Radiotherapy	Diego Carrasco, Pablo Cappagli, Flavio Colavecchia	Parallel and Distributed Algorithms