

# Contents

## Performance Evaluation and Analysis

Evaluating On-Node GPU Interconnects for Deep Learning Workloads . . . . .	3
<i>Nathan R. Tallent, Nitin A. Gawande, Charles Siegel, Abhinav Vishnu, and Adolfo Hoisie</i>	
Comparison of Parallelisation Approaches, Languages, and Compilers for Unstructured Mesh Algorithms on GPUs . . . . .	22
<i>G. D. Balogh, I. Z. Reguly, and G. R. Mudalige</i>	
Periodic I/O Scheduling for Super-Computers . . . . .	44
<i>Guillaume Aupy, Ana Gainaru, and Valentin Le Fèvre</i>	
A Performance Study of Quantum ESPRESSO's PWscf Code on Multi-core and GPU Systems . . . . .	67
<i>Joshua Romero, Everett Phillips, Gregory Ruetsch, Massimiliano Fatica, Filippo Spiga, and Paolo Giannozzi</i>	

## Performance Modeling and Simulation

Modeling Large Compute Nodes with Heterogeneous Memories with Cache-Aware Roofline Model . . . . .	91
<i>Nicolas Denoyelle, Brice Goglin, Aleksandar Ilic, Emmanuel Jeannot, and Leonel Sousa</i>	
A Scalable Analytical Memory Model for CPU Performance Prediction . . . . .	114
<i>Gopinath Chennupati, Nandakishore Santhi, Robert Bird, Sunil Thulasidasan, Abdel-Hameed A. Badawy, Satyajayant Misra, and Stephan Eidenbenz</i>	
Modeling UGAL on the Dragonfly Topology . . . . .	136
<i>Md Atiqul Mollah, Peyman Faizian, Md Shafayat Rahman, Xin Yuan, Scott Pakin, and Michael Lang</i>	
Resilient N-Body Tree Computations with Algorithm-Based Focused Recovery: Model and Performance Analysis . . . . .	158
<i>Aurélien Cavelan, Aiman Fang, Andrew A. Chien, and Yves Robert</i>	
Multi-fidelity Surrogate Modeling for Application/Architecture Co-design . . .	179
<i>Yiming Zhang, Aravind Neelakantan, Nalini Kumar, Chanyoung Park, Raphael T. Haftka, Nam H. Kim, and Herman Lam</i>	

A Slurm Simulator: Implementation and Parametric Analysis . . . . . 197  
*Nikolay A. Simakov, Martins D. Innus, Matthew D. Jones,  
Robert L. DeLeon, Joseph P. White, Steven M. Gallo,  
Abani K. Patra, and Thomas R. Furlani*

**Short Papers**

Path-Synchronous Performance Monitoring in HPC  
Interconnection Networks with Source-Code Attribution . . . . . 221  
*Adarsh Yoga and Milind Chabbi*

Performance and Energy Usage of Workloads on KNL  
and Haswell Architectures . . . . . 236  
*Tyler Allen, Christopher S. Daley, Douglas Doerfler,  
Brian Austin, and Nicholas J. Wright*

A Survey of Application Memory Usage on a National Supercomputer:  
An Analysis of Memory Requirements on ARCHER. . . . . 250  
*Andy Turner and Simon McIntosh-Smith*

**Author Index** . . . . . 261

High Performance Computing Systems. Performance  
Modeling, Benchmarking, and Simulation  
8th International Workshop, PMBS 2017, Denver, CO,  
USA, November 13, 2017, Proceedings  
Jarvis, S.; Wright, S.; Hammond, S. (Eds.)  
2018, XII, 261 p. 105 illus., Softcover  
ISBN: 978-3-319-72970-1