References

[1] H.M. Aktulga, J.C. Fogarty, S.A. Pandit, and A.Y. Grama. Parallel reactive molecular dynamics: Numerical methods and algorithmic techniques. *Parallel Computing*, 38(4-5):245–259, 2012.

[2] Amit Amritkar, Danesh Tafti, Rui Liu, Rick Kufrin, and Barbara Chapman. Openmp parallelism for fluid and fluid-particulate systems. *Parallel Computing*, 38(9):501–517, 2012.

[3] F. Argüello, D.B. Heras, M. Bóo, and J. Lamas-Rodríguez. The split-and-merge method in general purpose computation on gpus. *Parallel Computing*, 38(6-7):277–288, 2012.

[4] Gerassimos Barlas. Cluster-based optimized parallel video transcoding. *Parallel Computing*, 38(4-5):226–244, 2012.

[5] Wlodzimierz Bielecki, Marek Palkowski, and Tomasz Klimek. Free scheduling for statement instances of parameterized arbitrarily nested affine loops. *Parallel Computing*, 38(9):518–532, 2012.

[6] George Bosilca, Aurelien Bouteiller, Anthony Danalis, Thomas Herault, Pierre Lemarinier, and Jack Dongarra. Dague: A generic distributed dag engine for high performance computing. *Parallel Computing*, 38(1-2):37–51, 2012.

[7] Ümit V. Çatalyürek, John Feo, Assefaw H. Gebremedhin, Mahantesh Halappanavar, and Alex Pothen. Graph coloring algorithms for multicore and massively multithreaded architectures. *Parallel Computing*, 38(10-11):576–594, 2012.

[8] Yong Chen, Huaiyu Zhu, Hui Jin, and Xian-He Sun. Algorithm-level feedback-controlled adaptive data prefetcher: Accelerating data access for high-performance processors. *Parallel Computing*, 38(10-11):533–551, 2012.

[9] Peng Di, Hui Wu, Jingling Xue, Feng Wang, and Canqun Yang. Parallelizing sor for gpgpus using alternate loop tiling. *Parallel Computing*, 38(6-7):310–328, 2012.

[10] Peng Du, Rick Weber, Piotr Luszczek, Stanimire Tomov, Gregory Peterson, and Jack Dongarra. From cuda to opencl: Towards a performance-portable solution for multi-platform gpu programming. *Parallel Computing*, 38(8):391–407, 2012.

[11] M. Etinski, J. Corbalan, J. Labarta, and M. Valero. Parallel job scheduling for power constrained hpc systems. *Parallel Computing*, 38(12):615–630, 2012.

[12] Basilio B. Fraguela, Ganesh Bikshandi, Jia Guo, María J. Garzarán, David Padua, and Christoph von Praun. Optimization techniques for efficient hta programs. *Parallel Computing*, 38(9):465–484, 2012.

[13] Timothy D.R. Hartley, Erik Saule, and Ümit V. Çatalyürek. Improving performance of adaptive component-based dataflow middleware. *Parallel Computing*, 38(6-7):289–309, 2012.

[14] Pieter Hijma, Rob V. van Nieuwpoort, Ceriel J.H. Jacobs, and Henri E. Bal. Generating synchronization statements in divide-and-conquer programs. *Parallel Computing*, 38(1-2):75–89, 2012.

[15] Joshua Hursey and Richard L. Graham. Analyzing fault aware collective performance in a process fault tolerant mpi. *Parallel Computing*, 38(1-2):15–25, 2012.

[16] Takeshi Iwashita, Yu Hirotani, Takeshi Mifune, Toshio Murayama, and Hideki Ohtani. Large-scale time-harmonic electromagnetic field analysis using a multigrid solver on a distributed memory parallel computer. *Parallel Computing*, 38(9):485–500, 2012.

[17] Minhaj Ahmad Khan. Scheduling for heterogeneous systems using constrained critical paths. *Parallel Computing*, 38(4-5):175–193, 2012.

[18] Andreas Klöckner, Nicolas Pinto, Yunsup Lee, Bryan Catanzaro, Paul Ivanov, and Ahmed Fasih. Pycuda and pyopencl: A scripting-based approach to gpu run-time code generation. *Parallel Computing*, 38(3):157–174, 2012.

[19] Phil Miller, Aaron Becker, and Laxmikant Kalé. Using shared arrays in message-driven parallel programs. *Parallel Computing*, 38(1-2):66–74, 2012.

[20] Kathryn Mohror and Karen L. Karavanic. Trace profiling: Scalable event tracing on high-end parallel systems. *Parallel Computing*, 38(4-5):194–225, 2012.

[21] A. Moreno, E. Cesar, A. Guevara, J. Sorribes, and T. Margalef. Load balancing in homogeneous pipeline based applications. *Parallel Computing*, 38(3):125–139, 2012.

[22] Rahul Nagpal and Anasua Bhowmik. Criticality guided energy aware speculation for speculative multithreaded processors. *Parallel Computing*, 38(6-7):329–341, 2012.

[23] Aleksandr Ovcharenko, Daniel Ibanez, Fabien Delalondre, Onkar Sahni, Kenneth E. Jansen, Christopher D. Carothers, and Mark S. Shephard. Neighborhood communication paradigm to increase scalability in large-scale dynamic scientific applications. *Parallel Computing*, 38(3):140–156, 2012.

[24] Frederico Pratas, Pedro Trancoso, Leonel Sousa, Alexandros Stamatakis, Guochun Shi, and Volodymyr Kindratenko. Fine-grain parallelism using multi-core, cell/be, and gpu systems. *Parallel Computing*, 38(8):365–390, 2012.

[25] Nick Rutar and Jeffrey K. Hollingsworth. Data centric techniques for mapping performance data to program variables. *Parallel Computing*, 38(1-2):2–14, 2012.

[26] Martin Sandrieser, Siegfried Benkner, and Sabri Pllana. Using explicit platform descriptions to support programming of heterogeneous many-core systems. *Parallel Computing*, 38(1-2):52–65, 2012.

[27] Madan Sathe, Olaf Schenk, and Helmar Burkhart. An auction-based weighted matching implementation on massively parallel architectures. *Parallel Computing*, 38(12):595–614, 2012.

[28] Holger Scherl, Markus Kowarschik, Hannes G. Hofmann, Benjamin Keck, and Joachim Hornegger. Evaluation of state-of-the-art hardware architectures for fast cone-beam ct reconstruction. *Parallel Computing*, 38(3):111–124, 2012.

[29] Andrew G. Schmidt, Siddhartha Datta, Ashwin A. Mendon, and Ron Sass. Investigation into scaling i/o bound streaming applications productively with an all-fpga cluster. *Parallel Computing*, 38(8):344–364, 2012.

[30] Lucas Mello Schnorr, Guillaume Huard, and Philippe Olivier Alexandre Navaux. A hierarchical aggregation model to achieve visualization scalability in the analysis of parallel applications. *Parallel Computing*, 38(3):91–110, 2012.

[31] Jesper Larsson Träff. Alternative, uniformly expressive and more scalable interfaces for collective communication in mpi. *Parallel Computing*, 38(1-2):26–36, 2012.

[32] Francisco Vázquez, José Jesús Fernández, and Ester M. Garzón. Automatic tuning of the sparse matrix vector product on gpus based on the ellr-t approach. *Parallel Computing*, 38(8):408–420, 2012.

[33] Mickeal Verschoor and Andrei C. Jalba. Analysis and performance estimation of the conjugate gradient method on multiple gpus. *Parallel Computing*, 38(10-11):552–575, 2012.

[34] John R. Wernsing and Greg Stitt. Elastic computing: A portable optimization framework for hybrid computers. *Parallel Computing*, 38(8):438–464, 2012.

[35] Roman Wyrzykowski, Krzysztof Rojek, and Lukasz Szustak. Model-driven adaptation of double-precision matrix multiplication to the cell processor architecture. *Parallel Computing*, 38(4-5):260–276, 2012.

[36] Depeng Yang, Gregory D. Peterson, and Husheng Li. Compressed sensing and cholesky decomposition on fpgas and gpus. *Parallel Computing*, 38(8):421–437, 2012.