Huismann, Immo; **Stiller, Jörg**; **Fröhlich, Jochen**

Two-level parallelization of a fluid mechanics algorithm exploiting hardware heterogeneity.

(English) Zbl 1390.76008
Comput. Fluids 117, 114-124 (2015).

Summary: The prospect of wildly heterogeneous computer systems has led to a renewed discussion of programming approaches in high-performance computing, of which computational fluid dynamics is a major field. The challenge consists in harvesting the performance of all available hardware components while retaining good programmability. In particular the use of graphic cards is an important trend. This is addressed in the present paper by devising a hybrid programming model to create a heterogeneous data-parallel computation with a single source code. The concept is demonstrated for a one-dimensional spectral-element discretization of a fluid dynamics problem. To exploit the additional hardware available when coupling GPGPU-accelerated processes with excess CPU cores, a straight-forward load balancing model for such heterogeneous environments is developed. The paper presents a large number of run time measurements and demonstrates that the achieved performance gains are close to optimal. This provides valuable information for the implementation of fluid dynamics codes on modern heterogeneous hardware.

MSC:

76-04 Software, source code, etc. for problems pertaining to fluid mechanics
65M70 Spectral, collocation and related methods for initial value and initial-boundary value problems involving PDEs
65Y05 Parallel numerical computation
76M22 Spectral methods applied to problems in fluid mechanics

Keywords:

parallelization; heterogeneous computing; MPI; GPGPU; OpenMP; OpenACC

Software:

SkePU; MPI; PETSc; HOSTA; OpenACC; GPGPU; CUDA; StarPU; OmpSs

Full Text: DOI

References:

[1] Nvidia's Next Generation CUDA Compute Architecture: Kepler GK110. Technical report; 2012. <http://www.nvidia.de/content/PDF/kepler/NVI Kepler-GK110-Architecture-Whitepaper.pdf>.

[2] The OpenACC Application Programming Interface. Technical report; 2013. v. 2.0a. <http://www.openacc.org/>.

[3] Augonnet, C.; Thibault, S.; Namyst, R.; Wacrenier, P.-A., Starpu: a unified platform for task scheduling on heterogeneous multicore architectures, Concurr Comput Pract Exper, 23, 2, 187-198, (2011)

[4] Balay S, Abhyankar S, Adams MF, Brown J, Brune P, Buschelman K, et al. PETSc Web page; 2014. <http://www.mcs.anl.gov/petsc>.

[5] OpenMP Architecture Review Board. OpenMP Application Program Interface; 2011.

[6] Bolis, A.; Cantwell, C. D.; Kirby, R. M.; Sherwin, S. J., From h to p efficiently: optimal implementation strategies for explicit time-dependent problems using the spectral/hp element method, Int J Numer Methods Fluids, (2014)

[7] Bueno J, Planas J, Duran A, Badia RM, Martorell X, Ayguade E, et al. Productive programming of GPU clusters with OmpSs. In: IEEE 26th international parallel distributed processing symposium (IPDPS), 2012; May 2012. p. 557-68.

[8] Deville, M. O.; Fischer, P. F.; Mund, E. H., High-order methods for incompressible fluid flow, (2002), Cambridge University Press · Zbl 1007.76001

[9] Dong T, Dobrevand V, Kolev T, Rieben R, Tomov S, Dongarra J. Hydrodynamic computation with hybrid programming on CPU-GPU clusters. Technical report. <http://icl.cs.utk.edu/news_pub/submissions/ut-cs-13-714.pdf>.

[10] Enmyren, J.; Kessler, C. W., Skepu: a multi-backend skeleton programming library for multi-GPU systems, (Proceedings of the fourth international workshop on high-level parallel programming and applications, HLPP '10, (2010), ACM New York (NY, USA)), 5-14

[11] Dongarra, J., The international exascale software project roadmap, Int J High Perform Comput Appl, 25, 3-60, (2011)

[12] The MPI Forum. MPI: a message passing interface version 3.0; 2012.

[13] Fröhlich, J., Large eddy simulation turbulenter strömungen, (2006), B.G. Teubner Verlag Wiesbaden

[14] Göddeke, D.; Strzodka, R.; Mohd-Yusof, J.; McCormick, P. S.; Buijssen, S. H.M.; Grajewski, M., Exploring weak scalability for FEM calculations on a GPU-enhanced cluster, Parall Comput, 33, 10-11, 685-699, (2007)

[15] Gregg, C.; Hazelwood, K., Where is the data? why you cannot debate CPU vs. GPU performance without the answer, (Proceedings of the IEEE international symposium on performance analysis of systems and software, ISPASS '11, (2011), IEEE Computer Society Washington (DC, USA)), 134-144

[16] Hoshino T, Maruyama N, Matsuoka S, Takaki R. CUDA vs OpenACC: performance case studies with kernel benchmarks and a memory-bound CFD application. In: Proceedings of the 13th IEEE/ACM international symposium on cluster, cloud and grid computing; 2013. p. 136-43.

[17] Hundsdorfer, W., Partially implicit BDF2 blends for convection dominated flows, SIAM J Numer Anal, 38, 6, 1763-1783, (2001) · Zbl 1007.76052

[18] Jin, H.; Jespersen, D.; Mehrotra, P.; Biswas, R.; Huang, L.; Chapman, B., High performance computing using MPI and openmp on multi-core parallel systems, Parall Comput, 37, 562-575, (2011)

[19] Karniadakis, G. E.; Sherwin, S. J., Spectral/hp element methods for CFD, (1999), Oxford University Press · Zbl 0954.76001

[20] Kempe, T.; Fröhlich, J., An improved immersed boundary method with direct forcing for the simulation of particle laden flows, J Comput Phys, (2012) · Zbl 1402.76143

[21] Klöckner, A.; Warburton, T.; Bridge, J.; Hesthaven, J. S., Nodal discontinuous Galerkin methods on graphics processors, J Comput Phys, 228, 21, 7863-7882, (2009) · Zbl 1175.65111

[22] Komatitsch, D.; Erlebacher, G.; Göddeke, D.; Michéa, D., High-order finite-element seismic wave propagation modeling with MPI on a large GPU cluster, J Comput Phys, 229, 20, 7692-7714, (2010) · Zbl 1194.86019

[23] Lefebvre, M.; Guillen, P.; Le Gouez, J.-M.; Bastevand, C., Optimizing 2D and 3D structured Euler CFD solvers on graphical processing units, Comput Fluids, 70, 136-147, (2012) · Zbl 1365.76106

[24] Meuer H, Trohmaier E, Dongarra J, Simon H. Top500 list - June 2014; June 2014. <www.top500.org>.

[25] Mudalige, G. R.; Giles, M. B.; Thiyagalingam, J.; Reguly, I. Z.; Bertolli, C.; Kelly, P. H.J., Design and initial performance of a high-level unstructured mesh framework on heterogeneous parallel systems, Parall Comput, 39, 669-692, (2013)

[26] Patera, A. T., A spectral element method for fluid dynamics: laminar flow in a channel expansion, J Comput Phys, 54, 468-488, (1983) · Zbl 0535.76035

[27] Peters, N., Discussion of test problem A, (Peters, Norbert; Warnatz, Jürgen, Numerical methods in laminar flame propagation, Notes on numerical fluid mechanics, vol. 6, (1982), Vieweg+Teubner Verlag), 1-14

[28] Poinsot, T.; Veynante, D., Theoretical and numerical combustion, Edwards, (2005)

[29] Rus, P.; Štok, B.; Mole, N., Parallel computing with load balancing on heterogeneous distributed systems, Adv Eng Softw, 34, 185-201, (2003) · Zbl 1048.68021

[30] Xu, C.; Deng, X.; Zhang, L.; Fang, J.; Wang, G.; Jiang, Y., Collaborating CPU and GPU for large-scale high-order CFD simulations with complex grids on the tianhe-1A supercomputer, J Comput Phys, 278, 275-297, (2014) · Zbl 1349.76655