Technical University of Denmark

DTU

Development of an efficient GPU-accelerated model for fully nonlinear water waves

Madsen, Morten Gorm; Engsig-Karup, Allan Peter; Dammann, Bernd; Frisvad, Jeppe Revall

DTU Library
Technical Information Center of Denmark

ON THE DEVELOPMENT OF A GPU-ACCELERATED NONLINEAR WATER WAVE MODEL

M. G. MADSEN, A.P. ENGSIG-KARUP, B. DAMMANN AND JEPPE R. FRISVAD

Section for Scientific Computing, DTU Informatics, Technical University of Denmark, 2800 Kgs. Lyngby, Denmark.

Recent work that will be presented at the seminar is concerned with the development of an efficient high-throughput scalable parallel model for simulation of unsteady fully nonlinear water waves over uneven depths (OceanWave3D). The model is applicable to solve and analyze large-scale wave problems in coastal and offshore engineering and is both accurate, robust and efficient for the resolution of a broad range of important wave phenomena such as diffraction, refraction and wave-wave interactions.

We pursue the goal through algorithm redesign and parallelization using CUDA with starting point based on an optimized sequential single-CPU algorithm [1, 2, 3, 4] based on a flexible-order Finite Difference Method in three space dimensions. For the compute-intensive bottleneck problem, namely, a Laplace problem, High performance is pursued by utilizing many-core processing in the model focusing on many-core Graphics Processing Units (GPUs) for acceleration of CPU code execution of a low-storage iterative method based on a defect correction method preconditioned with a multigrid method. Conceptually, this method can be conceived as a p-multigrid method which enables fast convergence for high-order discretizations.

The key components of the algorithm are explicit which is favorable for embarrassingly parallel computations and is therefore well suited to the Single-Intructions-Multiple-Data SIMD type execution on GPUs. A parallel version of the model therefore enables the solution of relatively large wave problems and requires minimal communication for execution on a heterogenous CPU-GPU system. To maximize utilization of hardware, one needs to take into account the memory bound nature of the algorithm. Therefore, important optimizations of kernels that will be invoked on the GPUs, focuses on optimizing memory access patterns and reducing unnecessary memory transfers to high-latency global and local memory regions on the GPU. Despite the algorithm is highly data-parallel and therefore can easily be split into many subroutines that can each be optimized individually, experience shows that only a few components really needs attention.

The basic discretization strategy for the OceanWave3D model considered in this work is based on methods which is widely applied in numerous other research and industrial applications. Therefore the experiences with porting the relatively complicated model considered in this work to GPUs might have broad interests.

REFERENCES

[1] Li, B. and Fleming, C. A. A three dimensional multigrid model for fully nonlinear water waves. *Coast. Engng.*, **30**, 235–258, (1997).

[2] Bingham, H. B. and Zhang, H. On the accuracy of finite-difference solutions for nonlinear water waves. *J. Engng. Math.*, **58**, 211–228, (2007).

[3] Engsig-Karup, A.P. and Bingham, H.B. and Lindberg, O. An efficient flexible-order model for 3D nonlinear water waves. *J. Comp. Phys.*, **228**, 2100–2118, (2009).

[4] Engsig-Karup, A.P. Efficient low-storage solution of unsteady fully nonlinear water waves using a defect correction method. *Submitted for J. Comp. Phys.*, (2010).