



On a fast GPU-accelerated massively parallel method for fully nonlinear water wave computations

Engsig-Karup, Allan Peter; Madsen, M. G.; Glimberg, Stefan Lemvig

Published in:
Proceedings of ENUMATH

Publication date:
2011

Document Version
Publisher's PDF, also known as Version of record

[Link back to DTU Orbit](#)

Citation (APA):
Engsig-Karup, A. P., Madsen, M. G., & Glimberg, S. L. (2011). On a fast GPU-accelerated massively parallel method for fully nonlinear water wave computations. In *Proceedings of ENUMATH*
<http://www2.le.ac.uk/departments/mathematics/research/enumath2011>

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

On a fast GPU-accelerated massively parallel method for fully nonlinear water wave computations

A.P. Engsig-Karup, M. G. Madsen, S. L. Glimberg

Department of Informatics and Mathematical Modelling

Technical University of Denmark, 2800 Kgs. Lyngby, Denmark.

April 27, 2011

Background

Recent work that will be presented at the seminar is concerned with the ongoing development of an efficient high-throughput scalable parallel model for simulation of unsteady fully nonlinear free surface water waves over uneven depths. 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.

Recent developments significantly improves the strategy proposed by Li & Fleming (1997). A flexible-order finite difference algorithm for solving the governing equations in two (see Bingham & Zhang (2007)) and three space dimensions (Engsig-Karup et al. (2009); Engsig-Karup (2010)) enable efficient, scalable and low-storage solution of the governing equations. This combined with recent developments in modern many-core hardware and programming tools for general-purpose scientific computing suggest that these lines of development can be combined for further improving the overall performance.

We seek to demonstrate that it is now possible to significantly reduce the barriers for practical use of full potential flow theory as the modeling basis for efficient solution of coastal and offshore engineering problems. Our strategy is to do proof-of-concept by utilizing modern Graphics Processing Units (GPUs) for massively parallel computation using a heterogenous CPU-GPU hardware setup. Interestingly, such a hardware setup constitutes what can be considered an affordable standard consumer desktop environment.

Development of a massively parallel wave analysis tool

We will present recent work and progress on the development of a GPU-accelerated nonlinear free-surface model (OceanWave3D) for simulation of unsteady fully nonlinear water waves over uneven depths. In short, the flexible-order finite difference OceanWave3D model is based on a unified potential flow formulation. These model equations can account for fully nonlinear and dispersive waves within the breaking limit and under the assumption of irrotational inviscid flow. To establish the model as an efficient massively parallel tool for coastal and offshore engineering applications, we have both redesigned and reimplemented the entire algorithm to enable efficient utilization of allocated hardware resources - currently targeting modern many-core GPUs. Algorithmic efficiency is achieved by solving the bottleneck problem, a large sparse linear system of equations, iteratively by employing a defect correction method which is preconditioned by a robust multigrid method. This

strategy resulted in an achieved gain in performance of more than an order magnitude in both problem size and practical speedup (relative to optimized single-threaded CPU code) for the GPU baseline implementation. See figures 1 and 2.

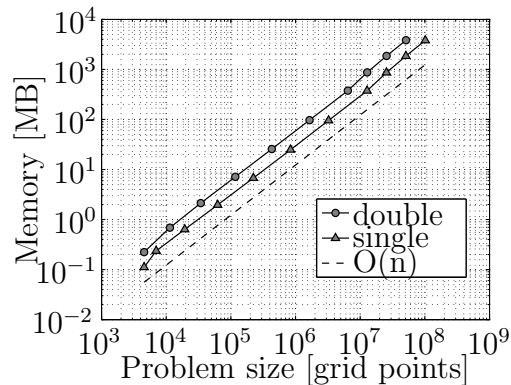


Figure 1: Scalability test of measured memory footprint.

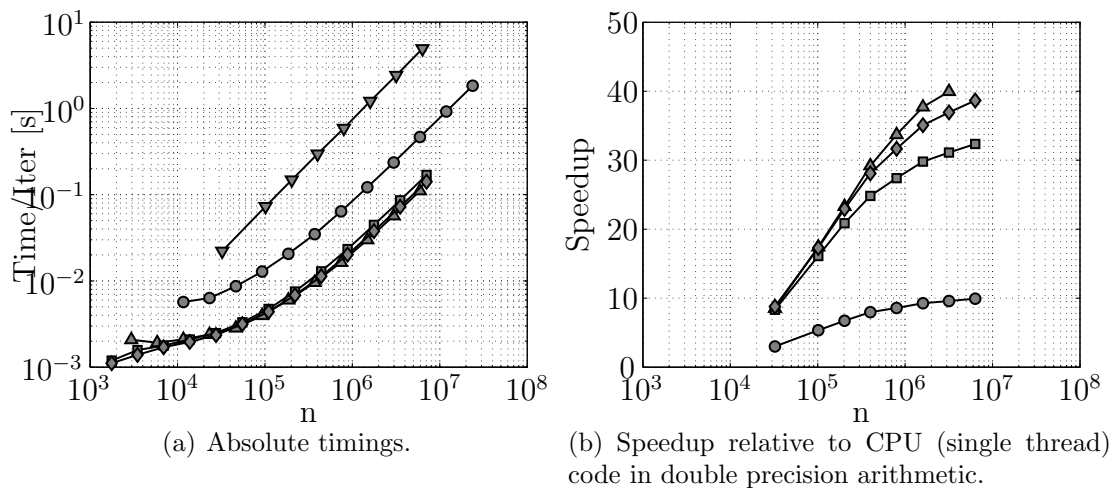


Figure 2: Scalability tests and performance comparisons in double precision arithmetic for Quadro FX 5800 (—●—), GeForce GTX 480 (—▲—), C2050 with ECC (—■—) and C2050 without ECC (—◆—) versus CPU (single thread) code (—▼—). Sixth order spatial discretization employed. The iterative Defect Correction method has been left-preconditioned with a Zebra Line Gauss-Seidel V-cycle multigrid strategy on each architecture.

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 non-linear 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).