Available online at www.sciencedirect.com
ScienceDirect
Journal of Hydrodynamics
2016,28(6):977-985
DOI: 10.1016/S1001-6058(16)60699-6
www.sciencedirect.com/
science/journal/l0016058

New prospects for computational hydraulics by leveraging high-performance heterogeneous computing techniques*

Qiuhua LIANG[1,2], Luke SMITH[2], Xilin XIA[2]
1. Hebei University of Engineering, Handan 056038, China
2. School of Civil Engineering and Geosciences, Newcastle University, Newcastle upon Tyne, UK,
E-mail: Qiuhua.Liang@ncl.ac.uk

(Received June 11, 2016, Revised October 16, 2016)

Abstract: In the last two decades, computational hydraulics has undergone a rapid development following the advancement of data acquisition and computing technologies. Using a finite-volume Godunov-type hydrodynamic model, this work demonstrates the promise of modern high-performance computing technology to achieve real-time flood modeling at a regional scale. The software is implemented for high-performance heterogeneous computing using the OpenCL programming framework, and developed to support simulations across multiple GPUs using a domain decomposition technique and across multiple systems through an efficient implementation of the Message Passing Interface (MPI) standard. The software is applied for a convective storm induced flood event in Newcastle upon Tyne, demonstrating high computational performance across a GPU cluster, and good agreement against crowd-sourced observations. Issues relating to data availability, complex urban topography and differences in drainage capacity affect results for a small number of areas.

Key words: computational hydraulics, high-performance computing, flood modeling, shallow water equations, shock-capturing hydrodynamic model

Introduction

Computational hydraulics is the field of developing and applying numerical models to solve hydraulic problems. It is a synthesis of multiple disciplines including but not restricted to applied mathematics, fluid mechanics, numerical analysis and computer science. The field has undergone rapid development in the last three decades, particularly following the advances in remote sensing technology, facilitating a rich source of topographic and hydrological data to support various modeling applications. Full 2-D and even 3-D numerical models have been developed to predict complex flow and transport processes and applied to simulate different aspects of hydrosystems, particularly flood inundation in extended floodplains. However, due to the restrictions in computational power, the application of these sophisticated models has long been restricted to performing simulations in relatively localized domains of a limited size.

Taking 2-D flood modeling as an example, considerable research effort has been devoted to improving the computational efficiency of flood models, in order to allow simulations at higher spatial resolutions and over greater extents. The common approaches that have been attempted include simplifying the governing equations, improving the numerical methods and developing parallel computing algorithms. Simplified 2-D hydraulic models with kinematic- or diffusive-wave approximations for flood inundation modeling had dominated the literature in the first decade of the 21st century[1-3]. These published works have shown that in certain cases these simplified models can reproduce reasonably well the flood extent and depth with high computational efficiency. However, accurate prediction of the evolution of flood waves involving complex processes is impossible without accurate representation of hydrodynamic effects, and is beyond the capabilities of these simplified models. Their reduced physical complexity may also cause increased

---

* Project supported by the UK NERC SINATRA Project (Grant No. NE/K008781/1).

Biography: Qiuhua LIANG (1974-), Male, Ph. D., Professor sensitivity to, and dependence on parameterization[4,5]. Furthermore, the reduction of computational time by these simplified approaches is not consistent across simulations, but highly dependent on the simulation resolution and flow hydrodynamics of the application[6,7].

Computationally more efficient numerical methods, including dynamically adaptive grids and sub-grid parameterization techniques, have also been widely developed to improve computational efficiency. Adaptive grids can adapt to the moving wet-dry interface and other flow and topographic features, thus facilitate accurate prediction of the flood front and routing processes. By creating a refined mesh only in areas of interest, dynamic grid adaption provides an effective means to relax the computational burden inherent in full dynamic inundation models[8,9]. However, since the time step of a simulation is controlled by the cells with highest level of refinement, which is concentrated on the most complex flow dynamics and highest velocities or free-surface gradients, the speedup achieved through adaptive grid simulation is generally limited, typically up to ~3 times for practical applications.

Rather than creating high-resolution mesh to directly capture small-scale topographic or flow features, as used in the adaptive mesh methods, techniques known as sub-grid parameterization have also been developed to integrate small aspects of topographic features into flood models, to enable more accurate and efficient but still coarse-resolution simulations[10,11]. For example, Soares-Frazão et al.[10] introduced a new shallow flow model with porosity to account for the reduction in storage due to sub-grid topographic features. The performance of the porosity model was compared with that of a refined mesh model explicitly reflecting sub-grid scale urban structures, and a more classical approach of raising local bed roughness. While being able to reproduce the mean characteristics of urban flood waves with less computational burden than refined mesh simulations, the porosity model was unable to accurately predict the formulation and propagation of certain localized wave features, e.g. reflected bores.

Parallel programming approaches have also been adopted to facilitate more efficient hydraulic simulations, and shown to exhibit good weak and strong scaling when software is structured appropriately[12,13]. However, none of the above three approaches has proven to be truly successful until the advent of heterogeneous computing leveraging graphics processing units (GPUs). GPUs are designed to process large volumes of data by performing the same calculation numerous times, typically on vectors and matrices. Such hardware architectures are well-suited to the field of computational fluid dynamics. New programming languages including CUDA and OpenCL have exposed this hardware for use in general-purpose applications (GPGPU). A number of attempts have been made to explore the benefits of GPU computing for highly efficient large-scale flood simulations. Early pioneers of such methods include Lamb et al.[14] who harnessed graphics APIs directly to implement a diffusion wave model (JFlow) for GPUs, Kalyanapu et al.[15] with a finite-difference implementation of the full shallow water equations, and later Brodtkorb et al.[16] with a finite-volume scheme. Such software is becoming increasingly mainstream, Néelz and Pender[17] report results from several commercial GPU hydraulics implementations while Smith and Liang[18] demonstrate the potential for generalized approaches applicable to both CPU and GPU co-processors. The most recent research also explores how domain decomposition across multiple GPUs can provide further performance benefits[19].

This work presents a hydrodynamic model, known as the High-Performance Integrated Modelling System (HiPIMS), for simulating different types of natural hazards (results are presented for flood modelling herein). HiPIMS solves the 2-D shallow water equations (SWEs) using a first-order or second-order shock-capturing finite-volume Godunov-type numerical scheme although only the first-order scheme is used in this work. To substantially improve computational efficiency, the model is implemented for high-performance heterogeneous computing using the OpenCL programming framework, and therefore can take advantage of either CPUs or GPUs with a single codebase. The model has also been developed to support simulations across multiple GPUs using a domain decomposition technique and across multiple systems through an efficient implementation of the Message Passing Interface (MPI) standard. The unprecedented capability of HiPIMS to achieve high-resolution large-scale flood inundation modeling at an affordable computational cost is demonstrated through an application to reproduce the June 2012 Newcastle flood event. Two simulations have been carried out with a 2 m resolution, one covering 36 km^2 of Newcastle central area and another covering 400 km^2 of Tyne and Wear, which respectively involve 8×10^6 and 10^8 computational cells.

1. HiPIMS–a high-performance shallow flow model

HiPIMS solves the matrix form of the 2-D SWEs with source terms, given as follows

$$\frac{\partial \boldsymbol{q}}{\partial t}+\frac{\partial \boldsymbol{f}}{\partial x}+\frac{\partial \boldsymbol{g}}{\partial y}=\boldsymbol{R}+\boldsymbol{S}_b+\boldsymbol{S}_f \tag{1}$$

where t is time, x and y are the Cartesian directions, $\boldsymbol{q}$ is the vector containing the conserved flow