

GPU Accelerated Discontinuous Galerkin Methods for Shallow Water Equations

Rajesh Gandham^{1,*}, David Medina¹ and Timothy Warburton¹

¹ *Department of Computational and Applied Mathematics, Rice University, 6100 Main Street, MS-134, Houston, TX-77005, USA.*

Received 7 January 2014; Accepted (in revised version) 27 November 2014

Abstract. We discuss the development, verification, and performance of a GPU accelerated discontinuous Galerkin method for the solutions of two dimensional nonlinear shallow water equations. The shallow water equations are hyperbolic partial differential equations and are widely used in the simulation of tsunami wave propagations. Our algorithms are tailored to take advantage of the single instruction multiple data (SIMD) architecture of graphic processing units. The time integration is accelerated by local time stepping based on a multi-rate Adams-Bashforth scheme. A total variational bounded limiter is adopted for nonlinear stability of the numerical scheme. This limiter is coupled with a mass and momentum conserving positivity preserving limiter for the special treatment of a dry or partially wet element in the triangulation. Accuracy, robustness and performance are demonstrated with the aid of test cases. Furthermore, we developed a unified multi-threading model OCCA. The kernels expressed in OCCA model can be cross-compiled with multi-threading models OpenCL, CUDA, and OpenMP. We compare the performance of the OCCA kernels when cross-compiled with these models.

AMS subject classifications: 35Q35, 35L60, 35L65, 65Z05, 65M99

Key words: Shallow water equations, discontinuous Galerkin, positivity preserving, slope limiting, GPUs, accelerators, multi-threading.

1 Introduction

The shallow water equations (SWE) are of great interest in the modeling of tsunamis, storm surges and tidal waves. They are the simplest nonlinear models for water wave propagation. The shallow water assumptions simplify three-dimensional wave propagation to two-dimensional hyperbolic partial differential equations, reducing the complexity of the model. The reduced complexity makes the shallow water equations attractive

*Corresponding author. *Email addresses:* Rajesh.Gandham@rice.edu (R. Gandham), David.S.Medina@rice.edu (D. Medina), timwar@rice.edu (T. Warburton)

for tsunami modeling. These equations are valid for long waves but may represent the wave propagation of short waves or dispersive waves poorly. However, these simplified equations provide satisfactory solutions of tsunami wave propagation [18] over long distances.

The shallow water equations are two-dimensional hyperbolic PDEs with velocity and fluid height as unknown quantities. These equations are complicated by the presence of largely varying length scales, varying bathymetry, and nonlinear effects near the shore. Stable, accurate and efficient algorithms are of great interest for these applications.

There is an extensive literature for finite difference [4, 21], finite volume [2, 18] and finite element [22] methods for shallow water equations. Recently, there has been growing interest in using discontinuous Galerkin methods (DG) for solutions of the shallow water equations [1, 8, 11, 17]. DG methods are locally mass conservative like finite volume methods and can achieve high order accuracy on unstructured meshes like finite element methods. This allows flexibility in handling irregular boundaries without compromising accuracy for problems with sufficiently smooth solutions. DG methods can achieve $\mathcal{O}(H^{N+1/2})$ accuracy with a piecewise degree N polynomial approximation [14]. Where, H is the largest length scale in the mesh.

In DG formulations, elements are coupled using weak penalty terms, resulting in localized memory access. Furthermore, a high order polynomial representation of the solution in each element results in high arithmetic intensity per degree of freedom. Both of these features are well suited for the GPU hardware architecture [12, 15, 16]. This motivated us to adopt GPUs along with a nodal DG discretization for large-scale tsunami simulations. Furthermore, to alleviate the need to write kernels for thread models like OpenMP, CUDA, and OpenCL separately, we developed OCCA: A unified approach to multi-threading languages. Kernels written in OCCA are cross compiled with any of these thread models at runtime. This gives us the flexibility in choosing the most efficient multi-threading model for a given hardware architecture, without writing new codes.

This paper is organized as follows: In Sections 2 and 3, we outline the governing equations and nodal discontinuous Galerkin discretization. In Section 4, we describe local time stepping using multi-rate Adams-Bashforth time integration. In Section 5, we explain the stabilization of numerical scheme using positivity preservation, wetting drying treatment and modified total variational bounded (TVB) limiter. Several tests for verification of accuracy and robustness are presented in Section 6. We discuss GPU kernels and their performance in Section 7. In Section 8, we will describe the fundamentals and features of the OCCA multi-threading model and compare the performance of the kernels written in OCCA, when they are cross compiled with OpenCL, CUDA, and OpenMP.

2 Governing equations

The shallow water equations are depth averaged incompressible Navier-Stokes equations, and are given in conservative form by [18],