Computers and Mathematics with Applications 72 (2016) 755–767

Contents lists available at ScienceDirect

Computers and Mathematics with Applications

journal homepage: www.elsevier.com/locate/camwa

A GPU accelerated level set reinitialization for an adaptive discontinuous Galerkin method

A. Karakus [a,*], T. Warburton [b], M.H. Aksel [a], C. Sert [a]

[a] Department of Mechanical Engineering, Middle East Technical University, Ankara, 06800, Turkey
[b] Department of Mathematics, Virginia Tech, Blacksburg, VA 24061-0123, United States

ARTICLE INFO

Article history:
Received 23 November 2015
Received in revised form 10 April 2016
Accepted 22 May 2016
Available online 16 June 2016

Keywords:
Discontinuous Galerkin
Level set
Local time integration
Hamilton–Jacobi
GPU computing

ABSTRACT

GPU accelerated high order reconstruction of signed distance function of the level set method is studied. The flow based reinitialization equation is discretized in space by using a nodal discontinuous Galerkin method on adaptive unstructured grids. Artificial diffusion with a modal decay rate based regularity estimator is used to damp out high frequency solution components near kinks, where mesh adaptivity is applied. A two rate Adams-Bashforth time integrator is developed to avoid time step restrictions resulting from artificial diffusion stabilization and local mesh refinement. Platform independence of the solver is achieved by using an extensible multi-threading programming API that allows runtime selection of different computing devices (GPU and CPU) and threading interfaces (CUDA, OpenCL and OpenMP). Overall, a highly scalable numerical scheme that preserves the simplicity of the original level set method is obtained. Performance and accuracy of the method to construct signed distance function on highly disturbed initial data with smooth and non-smooth interfaces are tested through distinct two- and three-dimensional problems.

1. Introduction

Level set (LS) methods [1] are commonly used to represent surface dynamics of multiphase flows. Integration of LS equations in time often distorts the scalar LS function such that it has flat or steep gradients near the interface. Reinitialization replaces the LS function with the signed distance function, which is regular and has uniform gradients.

Commonly used reinitialization methods can be categorized as fast sweeping [2] and fast marching [3] methods, which are based on the solution of static boundary value problems, and flow based methods [4], which make use of an artificial flow field to obtain the signed distance function in steady-state. Flow based methods are more flexible, accurate and easier to parallelize, and therefore preferred in the current study. The most popular flow based reinitialization equation is the following pseudo time first order PDE [4].

$$\frac{\partial \phi}{\partial t} + \operatorname{sgn}(\phi_0)(|\nabla \phi| - 1) = 0, \quad \phi(\mathbf{x}, 0) = \phi_0. \tag{1}$$

Characteristics of Eq. (1) emanate from the interface in the normal direction with unit speed, hence reinitialization starts from the interface and propagates in the normal direction. The equation requires boundary conditions when the interface

* Corresponding author.
E-mail address: akarakus@metu.edu.tr (A. Karakus).

http://dx.doi.org/10.1016/j.camwa.2016.05.025
0898-1221/

intersects the domain boundaries and creates regions without characteristic information, as in the case of contact line problems. These problems require additional information about the interface topology and are beyond the scope of this study. Except for the sign term, Eq. (1) is a Hamilton–Jacobi (HJ) equation. The general approach is to smear the sign term in a narrow band and treat Eq. (1) as a standard HJ equation with a smooth Hamiltonian. Finite difference Essentially Non-Oscillatory (ENO) and Weighted ENO (WENO) schemes have been developed to solve this equation on Cartesian grids [5,6]. Although these schemes have been adapted to unstructured grids [7], they are complicated to implement and inefficient computationally. Additionally, they result in mass loss problems in interface capturing applications [8].

Discontinuous Galerkin (DG) methods are a class of finite element methods that make use of completely discontinuous, piecewise polynomial approximations for spatial discretization and have excellent properties to overcome the problems mentioned above. High-order DG methods can accurately track interfaces with kinks having discontinuous derivatives over long time periods due to their low numerical dissipation [9]. Using DG for LS advection gives more accurate results compared to the use of standard HJ-ENO/WENO finite difference schemes [10–12]. However, it is rarely used for reinitialization, because the HJ equations cannot be written in conservative form and it is difficult to define suitable numerical fluxes for them in the DG framework.

The first DG method for the solution of HJ equations is introduced in [13]. Accuracy and stability of this method are analyzed in [14] and reinterpreted for a simplified implementation with reduced computational cost in [15], though this approach made the algorithm indirect, complicated and not optimal for reinitialization [12]. The local discontinuous Galerkin (LDG) method, which was first developed to discretize second order operators, was also used for the direct solution of HJ equations [16]. Recently, the DG method has found an application in LS interface modeling in which reinitialization is required to maintain regularity of the LS function. These studies used geometric reinitialization either based on height function on 2D Cartesian grids [12] or recursive contouring algorithm [17–19]. The reported results had noisy interfaces with considerable mass loss, which could be avoided only by using high recursion levels.

High-order DG methods, like other high-order methods, are known to produce oscillations when the approximation space is inadequate to resolve the main features of the true solution. Stabilization techniques for DG discretizations can be classified as limiting, polynomial reconstruction, spectral filtering and artificial diffusion ([20] and references therein). Among these, artificial diffusion relies on explicitly adding viscous terms to the governing equations in order to smooth the solution near the discontinuities. Although artificial diffusion offers a fast and reliable stabilization, it requires sophisticated discontinuity detectors to find the amount of diffusion and the regions over which it needs to be added [21,22]. Also it reduces the allowable time step size when a global explicit time integrator is used. Time step restriction can be relaxed with local time-stepping, which uses different time step sizes to satisfy the local CFL condition on different elements. Due to their efficiency and ease of implementation, local time integration schemes based on multi-rate Adams–Bashforth (MRAB) [23,24] and multi-rate Runge–Kutta (MRRK) methods [25,26] are used frequently. Different from the previous studies, we designed an efficient local time stepping strategy for the DG method, which does not require additional storage or computational effort, resulting in an efficient implementation on multi-threaded architectures.

Weak element connection and high-order approximation space of the DG method lead to local memory access and high arithmetic intensity. These properties make the DG method well suited for multi-threaded architectures, especially GPUs. Performance of the nodal DG methods on massively parallel architectures have been demonstrated for several applications [23,27–29]. Reinitialization formulation of the current study is accelerated using modern GPUs and many-core CPUs. Platform independence is achieved using the OCCA [30] kernel language that abstracts common multi-threading languages (OpenCL, CUDA, pThreads and OpenMP) and offers flexibility to choose the architecture and the programming language at runtime.

In this study, we introduce a GPU accelerated, explicit, multi-rate discontinuous Galerkin method for high-order solution of level set reinitialization on unstructured dynamic meshes. The rest of the paper is organized as follows: Section 2 provides the mathematical formulation including discretization of reinitialization equation, local time integration and basic properties of mesh adaptivity. Parallelization of the method on multi-threaded architectures is discussed in Section 3. Results demonstrating accuracy, performance and scalability of the method for two- and three-dimensional test cases are given in Section 4.

2. Formulation

2.1. Preliminaries

It is assumed that d dimensional domain $\Omega \subset R^d$ is well approximated by the computational domain, Ω_h, which is partitioned into non-overlapping, possibly non-conforming triangular/tetrahedral elements, $\Omega_h = \cup_{k=1}^{K} D_k$. Two elements, D_k^- and D_k^+ have a common face if $\partial D_k^- \cap \partial D_k^+ \neq \emptyset$, where ∂D_k denotes the element boundary. Also $\partial D_k = \sum_{f=1}^{N_f} \partial D_k^{f,-}$ with N_f being the total number of connections for an element which is equal to number of faces for conformal discretizations. Let $\mathbf{n}^- = -\mathbf{n}^+$ be the unit outward normal vector to ∂D_k. Q_k^- and Q_k^+ denote the traces of any scalar function, Q, when evaluated at the element boundaries of D_k^- and D_k^+, respectively.

The discontinuous approximate spaces are

$$V_N = \{v \in L^2(\Omega) | v|_{D_k} \in P_N(D_k),\ \forall D_k \in \Omega_h\} \tag{2}$$