Computers and Mathematics with Applications 72 (2016) 375–385

Contents lists available at ScienceDirect

Computers and Mathematics with Applications

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

Thermal link-wise artificial compressibility method: GPU implementation and validation of a double-population model

Christian Obrecht [a,*], Pietro Asinari [b], Frédéric Kuznik [a], Jean-Jacques Roux [a]

[a] INSA de Lyon, CETHIL UMR 5008, Université de Lyon, France
[b] SMaLL, Dipartimento Energia, Politecnico di Torino, Italy

ARTICLE INFO

Article history:
Available online 9 June 2015

Keywords:
Computational fluid dynamics
Link-wise artificial compressibility method
High-performance computing
Differentially heated cubic cavity
CUDA

ABSTRACT

The link-wise artificial compressibility method (LW-ACM) is a novel formulation of the artificial compressibility method for the incompressible Navier–Stokes equations showing strong analogies with the lattice Boltzmann method (LBM). The LW-ACM operates on regular Cartesian meshes and is therefore well-suited for massively parallel processors such as graphics processing units (GPUs). In this work, we describe the GPU implementation of a three-dimensional thermal flow solver based on a double-population LW-ACM model. Focusing on large scale simulations of the differentially heated cubic cavity, we compare the present method to hybrid approaches based on either multiple-relaxation-time LBM (MRT-LBM) or LW-ACM, where the energy equation is solved through finite differences on a compact stencil. Since thermal LW-ACM requires only the storing of fluid density and velocity in addition to temperature, both double-population thermal LW-ACM and hybrid thermal LW-ACM reduce the memory requirements by a factor of 4.4 compared to a D3Q19 hybrid thermal LBM implementation following a two-grid approach. Using a single graphics card featuring 6 GiB[1] of memory, we were able to perform single-precision computations on meshes containing up to 536^3 nodes, i.e. about 154 million nodes. We show that all three methods are comparable both in terms of accuracy and performance on recent GPUs. For Rayleigh numbers ranging from 10^4 to 10^6, the thermal fluxes as well as the flow features are in similar good agreement with reference values from the literature.

1. Introduction

Today's computational fluid dynamics (CFD) softwares are often designed to operate on unstructured meshes in order to take into account complex physical interfaces. During the last years however, there has been a resurgence of interest in numerical methods based on regular Cartesian grids. Although the representation of complex geometries with such approaches can be challenging, they usually benefit from a good asymptotic behaviour and are well-suited for high-performance implementations. A cardinal example of CFD methods operating on Cartesian grids is the well-acknowledged lattice Boltzmann method (LBM) which is nowadays used in several industry-grade softwares such as PowerFLOW, X-Flow, Fluidyna, and LaBS.

---

* Corresponding author.
E-mail address: christian.obrecht@insa-lyon.fr (C. Obrecht).

[1] Instead of the widespread but ambiguous GB and KB notations, we use the notations of the International System of Quantities, namely 1 GiB $= 2^{30}$ B, 1 KiB $= 2^{10}$ B, and 1 kB $= 10^3$ B.

http://dx.doi.org/10.1016/j.camwa.2015.05.022
0898-1221/

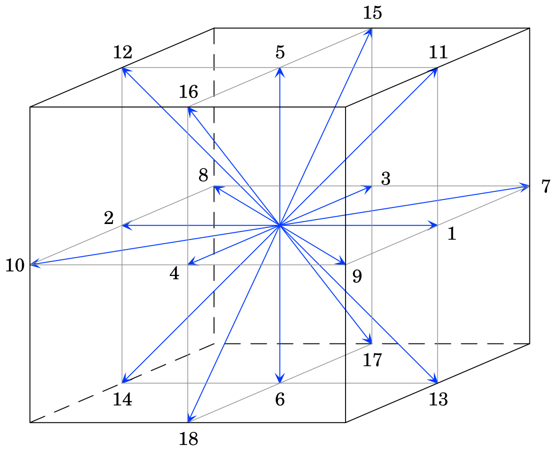

Fig. 1. The D3Q19 stencil—the arrows represent the $\boldsymbol{\xi}_\alpha$ velocities. This stencil links any bulk node to 18 of its nearest neighbours.

The link-wise artificial compressibility method (LW-ACM), introduced in [1], is a novel approach for solving the incompressible Navier–Stokes equations on Cartesian grids. Asymptotic analysis demonstrates that the LW-ACM yields the artificial compressibility equation [2], first proposed by Chorin in 1967. Although it primarily involves hydrodynamic macroscopic variables instead of discrete distribution functions, the LW-ACM exhibits strong formal similarities with LBM, making possible the use of both finite difference or LBM procedures, in particular regarding the treatment of boundary conditions. Recent work [3] has shown that the LW-ACM is well-suited for implementations on graphics processing units (GPUs), leading to a speed-up factor of 1.8 on recent hardware with respect to a state-of-the-art three-dimensional isothermal LBM solver. The large scale simulations of the lid-driven cubic cavity performed with this GPU implementation reveal that the LW-ACM is almost as accurate as multiple-relaxation-time LBM [4]. The chosen reference GPU LBM implementation uses a two-grid data layout in order to avoid read-after-write issues. With respect to this approach, the LW-ACM reduces memory consumption by a factor of $(19 \times 2 + 4)/(4 \times 2) = 5.25$. It should be mentioned that several approaches suited for GPU implementations make possible to use a single grid (see e.g. [5]) and that in this case the memory reduction factor is only $(19 + 4)/(4 \times 2) = 2.875$. These techniques, however, lead to fairly more intricate codes than the straightforward two-grid method and might be challenging to implement in conjunction with complex boundary conditions or couplings.

Regarding thermal simulations, it is well known that energy conserving LBM models suffer from spurious couplings between energy and shear modes of the linearised collision operator [6]. To overcome this flaw, thermal LBM models must separate the resolution of the energy equation from the resolution of the continuity and momentum equations, which is usually achieved by using either a finite difference scheme or a second set of distribution functions. Both of these approaches also apply to LW-ACM. The former has been successfully tested in a recent work [7]. In the present article, we report our experiments following the later approach in order to implement and validate a three-dimensional thermal flow solver on GPUs using the CUDA technology [8].

The remainder of the paper is organised as follows. Section 2 presents the implemented model starting with a summary of isothermal LW-ACM, followed by a description of the chosen double-population thermal LW-ACM model (DPTLA). Section 3 describes the algorithmic aspects of DPTLA as well as the GPU implementation principles. Section 4 presents the methodology we followed in our validation study for both our performance results and our simulation data for differentially heated cubic cavity. Section 5 reports and discusses the results. A throughout comparison of thermal fluxes and characteristic flow features is carried out for the present model and reference data [9], as well as for results obtained using hybrid thermal LBM (HTLBM) [10,11] and hybrid thermal LW-ACM (HTLA) [7]. Section 6 provides some concluding remarks.

2. Model

2.1. Link-wise artificial compressibility method

Similarly to standard isothermal LBM, the isothermal LW-ACM operates on a *lattice*, i.e. a regular Cartesian grid of mesh size δx with a constant time step δt associated to a finite set of Q velocities $\{\boldsymbol{\xi}_\alpha\}$ with $\boldsymbol{\xi}_0 = \mathbf{0}$. In the present work, we assume *diffusive scaling*, i.e. we have $\delta x = \varepsilon$ and $\delta t = \varepsilon^2$ for some small parameter ε. The set of velocities, which we will refer to as the *stencil* of the model, is chosen such as to link neighbouring nodes of the mesh in one time step. For our simulations, we used the three-dimensional D3Q19 stencil displayed in Fig. 1.

The simulated fluid is represented at each node by a set of Q independent variables $\{f_\alpha\}$ which are the formal equivalents of the LBM distribution functions. Let us denote by ρ and $\boldsymbol{u}$ the density and velocity of the fluid, the variables f_α