

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

Original

Thermal link-wise artificial compressibility method: GPU implementation and validation of a double-population model / Obrecht, Christian; Asinari, Pietro; Kuznik, Frédéric; Roux, Jean Jacques. - In: COMPUTERS & MATHEMATICS WITH APPLICATIONS. - ISSN 0898-1221. - ELETTRONICO. - 72:2(2016), pp. 375-385. [10.1016/j.camwa.2015.05.022]

Availability:

This version is available at: 11583/2624289 since: 2016-06-30T14:17:13Z

Publisher:

Elsevier Ltd

Published

DOI:10.1016/j.camwa.2015.05.022

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

(Article begins on next page)

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

^aINSA de Lyon, CETHIL UMR 5008, Université de Lyon, France

^bSMaLL, Dipartimento Energia, Politecnico di Torino, Italy

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 hybrid D3Q19 MRT-LBM. Using a single graphics card with 6 GiB¹ 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.

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

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

*To whom correspondence should be addressed. E-mail: christian.obrecht@insa-lyon.fr.

¹Instead 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 and 1 KiB = 2^{10} B.

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.

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 [9] 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 and a memory consumption reduced by a factor of 5 with respect to a state-of-the-art three-dimensional isothermal LBM solver. Furthermore, 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 [3], which is generally considered as the most elaborate version of LBM.

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 [4]. 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 [10]. 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 [5].

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 [11], as well as for results obtained using hybrid thermal LBM (HTLBM) [6, 7] and hybrid thermal LW-ACM (HTLA) [10]. 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 $\{\xi_\alpha\}$ with $\xi_0 = \mathbf{0}$. 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.




Figure 1: The D3Q19 stencil — The arrows represent the ξ_α velocities. This stencil links any bulk node to 18 of its nearest neighbours.

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 ρ and \mathbf{u} the density and velocity of the fluid, the variables f_α obey:

$$\rho = \sum_{\alpha} f_{\alpha} \quad \text{and} \quad \rho \mathbf{u} = \sum_{\alpha} f_{\alpha} \xi_{\alpha}. \quad (1)$$

The fundamental updating rule of isothermal LW-ACM is given by:

$$f_{\alpha}(\mathbf{x}, t+1) = f_{\alpha}^{(e)}(\mathbf{x} - \xi_{\alpha}, t) + 2 \left(\frac{\omega - 1}{\omega} \right) \left(f_{\alpha}^{(e,o)}(\mathbf{x}, t) - f_{\alpha}^{(e,o)}(\mathbf{x} - \xi_{\alpha}, t) \right), \quad (2)$$

where ω is a relaxation frequency related to the kinematic viscosity ν , $f_{\alpha}^{(e)}$ are the local equilibrium functions and $f_{\alpha}^{(e,o)}$ are the odd parts of the equilibrium functions defined as:

$$f_{\alpha}^{(e,o)}(\rho, \mathbf{u}) = \frac{1}{2} \left(f_{\alpha}^{(e)}(\rho, \mathbf{u}) - f_{\alpha}^{(e)}(\rho, -\mathbf{u}) \right). \quad (3)$$

The formal similarity of LW-ACM and LBM becomes more explicit when reformulating Eq. 2 in:

$$f_{\alpha}(\mathbf{x}, t+1) = f_{\alpha}^*(\mathbf{x} - \xi_{\alpha}, t) + 2 \left(\frac{\omega - 1}{\omega} \right) f_{\alpha}^{(e,o)}(\mathbf{x}, t), \quad (4)$$

$$f_{\alpha}^*(\mathbf{x}, t+1) = f_{\alpha}^{(e)}(\mathbf{x}, t+1) - 2 \left(\frac{\omega - 1}{\omega} \right) f_{\alpha}^{(e,o)}(\mathbf{x}, t+1), \quad (5)$$

the former relation being the equivalent of the streaming step in its pull version, whereas the latter is the equivalent of the collision step.

It should however be noted that to apply the updating rule of LW-ACM at a given node x , it is only necessary to compute the equilibria at x and some of its neighbours, these equilibria being known functions of the macroscopic variables of the fluid. Contrary to LBM where the distribution functions are stored at each time step, the LW-ACM only requires to save the density and the velocity, thus leading to significant memory savings. The artificial speed of sound being set to $c = 1/\sqrt{3}$, the equilibrium functions are given by:

$$f_\alpha^{(e)}(\rho, \mathbf{u}) = w_\alpha \rho \left(1 + 3\xi_\alpha \cdot \mathbf{u} + \frac{9}{2} (\xi_\alpha \cdot \mathbf{u})^2 - \frac{3}{2} \mathbf{u}^2 \right). \quad (6)$$

with the weights w_α depending on the stencil. For the D3Q19 stencil, we have:

$$w_\alpha = \begin{cases} 1/3 & \alpha = 0, \\ 1/18 & \alpha = 1, \dots, 6, \\ 1/36 & \alpha = 7, \dots, 18. \end{cases} \quad (7)$$

As shown in appendix B of [1], under the assumption of diffusive scaling and with appropriate scaling of the moments, asymptotic expansion of Eq. 2 yields the artificial compressibility equation and the momentum equation up to second order in space and first order in time. The kinematic viscosity obeys:

$$\nu = \frac{1}{3} \left(\frac{1}{\omega} - \frac{1}{2} \right). \quad (8)$$

As isothermal LBM, isothermal LW-ACM is therefore an effective numerical procedure to solve the isothermal incompressible Navier–Stokes equations.

2.2. Double-population link-wise artificial compressibility method

Double-population thermal LW-ACM is an extension of isothermal LW-ACM involving a second set of Q dependent variables $\{g_\alpha\}$ such that:

$$T = \sum_\alpha g_\alpha \quad (9)$$

where T is the temperature. Although one might use different stencils for the two variable sets, in our work we have chosen to use the D3Q19 stencil for both, which is more convenient in terms of implementation.

The updating rule for the second set is:

$$g_\alpha(\mathbf{x}, t+1) = g_\alpha^{(e)}(\mathbf{x} - \xi_\alpha, t) + 2 \left(\frac{\omega_t - 1}{\omega_t} \right) \left(g_\alpha^{(e,e)}(\mathbf{x}, t) - g_\alpha^{(e,e)}(\mathbf{x} - \xi_\alpha, t) \right) \quad (10)$$

where ω_t is a relation frequency linked to the thermal diffusivity κ , $g_\alpha^{(e)}$ are the local equilibrium functions and $g_\alpha^{(e,e)}$ are the even parts of the equilibrium functions defined as:

$$g_\alpha^{(e,e)}(T, \mathbf{u}) = \frac{1}{2} \left(g_\alpha^{(e)}(T, \mathbf{u}) + g_\alpha^{(e)}(T, -\mathbf{u}) \right). \quad (11)$$

The definition of the equilibrium functions for g_α is similar to the one of the equilibrium functions for f_α , namely:

$$g_\alpha^{(e)}(T, \mathbf{u}) = w_\alpha T \left(1 + 3\xi_\alpha \cdot \mathbf{u} + \frac{9}{2} (\xi_\alpha \cdot \mathbf{u})^2 - \frac{3}{2} \mathbf{u}^2 \right). \quad (12)$$

It should be noted that, in principle, because energy equation is linear with regards to \mathbf{u} , one could remove all terms proportional to \mathbf{u}^2 in this equilibrium.

With appropriate scaling of the moments, asymptotic expansion of Eq. 10 shows that the temperature field is governed by an advection-diffusion equation (see appendix B of [1] for details), the thermal diffusivity being:

$$\kappa = \frac{1}{3} \left(\frac{1}{\omega_t} - \frac{1}{2} \right). \quad (13)$$

Under the assumption of weak coupling between fluid-dynamic and energy equations, the double-population thermal LW-ACM is thus an effective approach to perform thermal flow simulations.

3. Implementation

3.1. Algorithmic aspects

As shown in [9], an implementation of isothermal LW-ACM may use either the single-step formulation (Eq. 2) or the two-step formulation (Eqs. 4 and 5) of the updating rule. The latter approach is interesting since it makes possible to easily adapt an existing LBM code to LW-ACM. It proves also to be well-suited to LBM specific boundary conditions such as simple bounce-back or interpolated bounce-back. The data access pattern is almost similar to LBM except in Eq. 4 where there is an additional access to $\rho(\mathbf{x}, t)$ and $\mathbf{u}(\mathbf{x}, t)$ in order to compute $f_\alpha^{(e,o)}(\mathbf{x}, t)$. In a memory bound context, a two-step LW-ACM implementation might therefore be slightly less efficient than the corresponding LBM solver.

In terms of memory consumption however, the two-step version is equivalent to LBM, whereas the single-step version may lead to a considerable diminution. As already stated, to apply the updating rule expressed in Eq. 2, it is only necessary to store the density and velocity fields at each time step. Let us consider the simple method to avoid read-after-write issues, involving two instances of the simulation data, one for even time steps and the other for odd time steps. When following this approach and using the D3Q19 stencil, the ratio of memory consumption between LBM and single-step LW-ACM is $(19 \times 2 + 4)/(4 \times 2) = 5.25$. As an example, in the case of single-precision D3Q19, a standard GPU computing device fitted with 6 GiB of memory is able to handle about 38 million nodes using LBM and over 200 million nodes using single-step LW-ACM.

From a local point of view, in the case of three-dimensional simulations, updating a single node using single-step LW-ACM requires $4Q$ read operations and 4 write operations whereas LBM requires Q read operations and Q write operations (or possibly $Q + 4$ write operations when saving the density and velocity fields). In a memory bound context, the main optimisation goal for single-step LW-ACM is therefore the reduction of read redundancy. This might be achieved for CPU implementations by an appropriate tiling, in order to optimise cache reuse, and for GPU implementations by an efficient use of shared memory.

In the case of the double-population thermal LW-ACM, the updating rule of the second population expressed in Eq. 10 may also be split in two steps. As for isothermal LW-ACM, a double-population thermal LBM solver may thus easily be adapted to DPTLA. The memory requirements are however even higher than for isothermal LBM. We therefore chose to follow the single-step approach for both updating rules which only requires the additional storing of the temperature field. The general formulation of our implementation is outlined in Algorithm 1.

Compared to single-precision D3Q19 HTLBM, the memory consumption is reduced by a factor of $(20 \times 2 + 4)/(5 \times 2) = 4.4$, and the aforementioned GPU computing device is able to handle about 37 million nodes using HTLBM and over 160 million nodes using DPTLA. Considering a cubic cavity, the largest available size is therefore 320 for HTLBM and 536 for DPTLA. Moreover, when using HTLBM with simple bounce-back, the thermal boundary conditions must take into account that the location of the physical interface between fluid and solid is located halfway between two nodes. In our implementation, we chose finite difference formulations for both the fluid-dynamic and thermal boundary conditions. The fluid-solid interface may therefore be on-grid, leading to simpler expressions for the thermal boundary conditions.

3.2. GPU implementation

Our implementation of the DPTLA was developed using the CUDA technology, introduced by the Nvidia company in 2007 as a framework for general purpose computing on graphics processing units. A detailed description of CUDA programming principles being beyond the scope of this article, we refer the interested reader to [5]. There are however a few aspects regarding the CUDA framework that should be mentioned before proceeding to the description of our program. A CUDA program mainly consists of host code, which are portions of the program run by the host system, and *kernels*, which are functions run in parallel by the GPU device in several *threads*. Each thread has access to its own instance of the local variables of the kernel. The whole of the threads running a kernel is called the *execution grid*, which is subdivided into *blocks* of threads of identical size. On modern Nvidia hardware, grids and blocks may have up to three dimensions. Threads within a block are executed synchronously in groups named *warps* and may exchange data through a local shared memory, which is fast but limited in size. At grid level, blocks are executed asynchronously, and threads belonging to different blocks may only exchange data through the off-chip device memory. Although cached, the device memory is by far slower than registers and shared memory. Its size ranges from 1 GiB on basic commodity graphics cards to 12 GiB on high-end dedicated computing devices.

For our GPU implementation of the DPTLA, we followed a strategy similar to the one described in [9]. We chose to map the computation domain by the execution grid, each thread being assigned to a single node of the mesh, which enables to take advantage of the massive parallelism of GPUs. Each block of threads is therefore associated to a block of nodes in the computation domain. For the sake of clarity, let us call *halo* of a block the set of all next neighbour nodes of this block. When processing the updating rules on a block of nodes, the associated threads must first load from device memory into shared memory the density, velocity, and temperature for the block and its halo. Each thread carries out the read operations for its assigned node. In addition, as illustrated by Fig. 2, the threads at the faces of the block are responsible for the nodes at the faces of the halo and the threads

1. **for all** time step t **do**
2. **for all** mesh node \mathbf{x} **do**
3. **for all** index α **do**
4. **if** $\mathbf{x} - \xi_\alpha$ is a boundary node **then**
5. apply boundary conditions
6. **else**
7. load $\rho(\mathbf{x} - \xi_\alpha, t)$, $\mathbf{u}(\mathbf{x} - \xi_\alpha, t)$, and $T(\mathbf{x} - \xi_\alpha, t)$
8. **end if**
9. compute $f_\alpha^{(e)}(\mathbf{x} - \xi_\alpha, t)$ and $f_\alpha^{(e,o)}(\mathbf{x} - \xi_\alpha, t)$
10. compute $g_\alpha^{(e)}(\mathbf{x} - \xi_\alpha, t)$ and $g_\alpha^{(e,e)}(\mathbf{x} - \xi_\alpha, t)$
11. **end for**
12. **for all** index α **do**
13. compute $g_\alpha(\mathbf{x}, t + 1)$
14. **end for**
15. compute $T(\mathbf{x}, t + 1)$ and buoyancy force $\mathbf{F}(\mathbf{x}, t)$
16. **for all** index α **do**
17. compute $f_\alpha(\mathbf{x}, t + 1)$
18. **end for**
19. compute $\rho(\mathbf{x}, t + 1)$ and $\mathbf{u}(\mathbf{x}, t + 1)$
20. store $\rho(\mathbf{x}, t + 1)$, $\mathbf{u}(\mathbf{x}, t + 1)$, and $T(\mathbf{x}, t + 1)$
21. **end for**
22. **end for**

Algorithm 1: General formulation of the double-population thermal LW-ACM.

at the edges of the block are responsible for the nodes at the edges of the halo. The density and velocity components are grouped into a `float4` structure, whereas an additional `float` number is used for the temperature. After all load operations have completed, which is enforced by a synchronisation barrier, the threads proceed to the computations and store the updated values for the block in device memory.

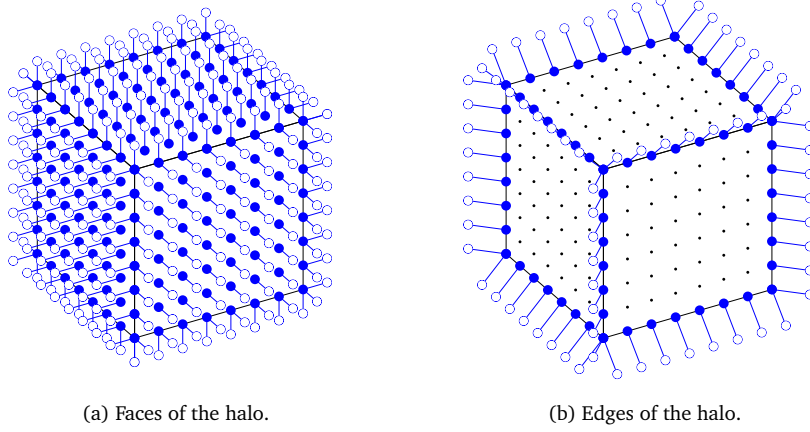


Figure 2: Data access to the halo of a block for the main computation kernel — The plain discs represent the active threads whereas the hollow ones represent the nodes from which ρ , \mathbf{u} , and T are read.

The shape and size of the blocks must be chosen carefully in order to minimise read redundancy. On one hand, to obtain the minimal ratio between halo and block, the blocks must be cubic and as large as possible. On the other hand, the corresponding data must fit into shared memory and number of threads within the block should be a multiple of the warp size (which is equal to 32 for all hardware generations up to now). Our experiments show that, with a maximal shared memory size of 48 KiB as with our test hardware, the optimal dimensions for a block are $8 \times 8 \times 8$. In this case, the read redundancy ratio is $(10^3 - 8)/8^3 \approx 1.94$ (the values at the vertices of the halo are not needed when using the D3Q19 stencil). In the bulk of the computation domain, the average amount of data read is therefore about 9.69 words per node for D3Q19 DPTLA, whereas it is 26 words per node for the D3Q19 HTLBM GPU implementation described in [6].

4. Methodology

4.1. Differentially heated cubic cavity

To validate our DPTLA solver, we simulated the differentially heated cubic cavity (DHCC) represented in Fig. 3. This test case consists of a cubic cavity in a constant gravity field \mathbf{g} , with two opposite isothermal vertical walls, one at $T = +T_0$ and the other at $T = -T_0$, and four adiabatic walls. In our simulations, the kinematic viscosity is set to $\nu = 0.05$ in lattice units, i.e. using δx and δt as units of length and time respectively, and the Prandtl number $\text{Pr} = \nu/\kappa$ is set to $\text{Pr} = 0.71$.

The Rayleigh number $Ra = Pr \cdot Gr$ ranges from 10^4 to 10^7 , Gr being the Grashof number defined as:

$$Gr = \frac{2T_0\beta g N^3}{\nu^2}, \quad (14)$$

where β is the thermal expansion coefficient and N is the size of the cavity in lattice units. We assume the Boussinesq approximation for the buoyancy force F , which is therefore expressed as:

$$F = -\rho T \beta g. \quad (15)$$




Figure 3: The DHCC test case — The cavity consists of two isothermal vertical walls and four adiabatic walls.

4.2. Convergence criterion

Let us denote $T_{i,j,k}^l$ the temperature at time step l and node (i, j, k) in lattice units. In numerous works devoted to the DHCC, the convergence criterion chosen for the temperature field is equivalent to:

$$\frac{1}{N^3} \sum_{i,j,k} \left(T_{i,j,k}^{l+p} - T_{i,j,k}^l \right)^2 < B, \quad (16)$$

where p is a constant number of time steps and B is a constant convergence threshold. However, such criterion is not satisfactory since it is resolution-dependent. As a matter of fact, in the case of diffusive scaling, setting $\varepsilon = 1/N$, leads to $\delta x = \varepsilon = 1/N$ and $\delta t = \varepsilon^2 = 1/N^2$. Hence, the duration in physical units corresponding to p time steps decreases as the resolution increases. In other words, the higher the resolution, the less severe the convergence criterion.

We chose instead to use the L^2 -norm of the temperature evolution rate, which is expressed as:

$$\|\dot{T}^l\|_2 = \left[\delta x^3 \sum_{i,j,k} \left(\frac{T_{i,j,k}^{l+1} - T_{i,j,k}^l}{\delta t} \right)^2 \right]^{1/2} = \left[N \sum_{i,j,k} \left(T_{i,j,k}^{l+1} - T_{i,j,k}^l \right)^2 \right]^{1/2} \quad (17)$$

and we declare convergence when the criterion $\|\dot{T}^l\|_2 < 10^{-2}$ is met.

4.3. Performance evaluation

We carried out performance measures using an Nvidia GeForce GTX Titan Black graphics card featuring a GK110 GPU with 2880 cores and 6 GiB of device memory. The host system is an Intel Xeon workstation running the Debian 7.7 GNU/Linux operating system and the CUDA 6.5 framework. We tested our DPTLA solver as well as the HTLA solver presented in [10] and the single-GPU HTLBM solver described in [6], for cavities of size ranging from 96 to 512 for the LW-ACM based programs, and from 96 to 320 for the HTLBM program. Due to adaptive power management, the results obtained with this computing device for short runs significantly depends on the initial temperature of the GPU. In order to get reproducible results, we performed simulations lasting approximately 60 s whatever the size of the cavity. We estimated the required number of time steps by launching preliminary runs on cavities of various sizes.

5. Results and discussion

5.1. Performance

The performance results using single-precision on the DHCC for increasing cavity size are reported in Fig. 4. The three programs have close performance, ranging from 1197 to 1394 million lattice node updates per second (MLUPS). On average, the DPTLA and the HTLA implementations perform better than the HTLBM solver by 5 % and 7 % respectively. These small differences seem to be mostly hardware-related, since former tests on a GeForce GTX Titan, which is slightly less efficient than the device used in the present work, showed almost identical performance results for the three solvers.

Although all three thermal GPU solvers achieve similar performance, the corresponding data throughput for the DPTLA and HTLA codes is considerably lower than for the HTLBM code. For the GeForce GTX Titan Black, the maximum sustained throughput reported by the `bandwidthTest` program from the CUDA software development kit is around 223.9 GiB/s. As shown in Tab. 1, the relative data throughput with respect to this maximum throughput is on average 32.0 % and 32.8 % for the DPTLA and HTLA implementations respectively, whereas the HTLBM implementation reaches 83.2 % on average.

A close inspection of the LW-ACM based programs shows that this significantly lower parallelisation efficiency is due to the local synchronisation barriers required after data fetching. A cure to this issue would be to decrease the size of the thread blocks, however this would in turn increase read redundancy. Our experiments show that the retained size of $8 \times 8 \times 8$ is the best trade-off for the targeted hardware but this might be different for future CUDA hardware generations.

5.2. Nusselt numbers

Tab. 2 provides the Nusselt numbers at the isothermal hot wall obtained with DPTLA as well as the results for HTLA given in [10] for N ranging from 128 to 512. For HTLBM, in order to apply the same convergence criterion, we used a modified version of the multi-GPU solver described in [7] running on one to four Tesla C2075 computing devices for $N = 128, 256$, and 384 . All the computations were carried out in single-precision. The temperature gradient in the direction




Figure 4: Performance comparison for the DPTLA, HTLA, and HTLBM solvers — Performance is reported in million lattice node updates per second (MLUPS). The programs were tested on the same GeForce GTX Titan Black graphics card.

normal to the wall is estimated by second-order finite differences. We furthermore report the relative discrepancy of these results with respect to the reference values published by Tric *et al.* in [11]. These benchmark data were obtained using an accurate pseudo-spectral method. We selected the values computed on a 81^3 grid for our comparisons.

For Rayleigh numbers ranging from 10^4 to 10^6 , all three models are in good agreement with our reference. The relative discrepancy being typically less than 1%, and close to 0.1% at the highest resolution. For $Ra = 10^7$, all simulations diverge at $N = 128$. In terms of agreement with the reference values, the DPTLA performs slightly better than the two other models, but this is inconclusive since at these levels of spatial resolution, the thermal boundary layer is unlikely to be accurately resolved.

5.3. Flow features

In addition to heat fluxes, the work of Tric *et al.* provides some informations regarding the flow features in the DHCC, namely the value and location of the maxima of the velocity components u_m , v_m , and w_m . Tabs 3, 4, and 5 report these results for all three models at $N = 384$, as well as comparisons with the reference data. For this purpose, the velocity components are scaled by κ/N and the coordinates range from $-1/2$ to $1/2$. Again, we selected the reference values computed on a 81^3 grid. It should be mentioned that there is a probable misprint in the table providing results for $Ra = 10^4$ (see Tab. 4 of [11]). In the part concerning w_m , the line labelled x actually provides the y -coordinate and the line providing the x -coordinate is missing. There is therefore no comparison for this case.

Again, for Rayleigh numbers ranging from 10^4 to 10^6 , all three models are in good agreement with our reference. For both the maximal velocity components

N	DPTLA (%)	HTLA (%)	HTLBM (%)
96	33.4	34.1	83.5
128	32.4	31.1	85.1
160	30.8	31.3	84.1
192	31.7	32.2	84.3
224	31.4	32.2	83.7
256	31.6	32.3	82.1
288	31.6	32.8	83.5
320	31.7	33.4	79.7
352	32.1	33.0	—
384	32.0	33.3	—
416	32.3	33.4	—
448	32.1	33.5	—
480	32.3	33.5	—
512	32.2	33.7	—
Average	32.0	32.8	83.2

Table 1: Relative data throughput of the three solvers with respect to the maximum sustained throughput of the GeForce GTX Titan Black

and their location, the relative discrepancy is typically less than 1 %. The same conclusion holds for $Ra = 10^7$, except for u_m where all models disagree with the reference data.

6. Conclusions

In this contribution, we describe an effective approach to implement a double-population LW-ACW thermal flow solver on CUDA GPUs. Performing large scale simulations of the differentially heated cubic cavity, we show that DPTLA is comparable to HTLA and HTLBM, i.e. hybrid LW-ACW or LBM thermal flow solvers, both in terms of accuracy with respect to reference data, and in terms of performance on recent GPUs. LW-ACM based thermal flow solvers are promising, since in both cases the device memory consumption is divided by 4.4 compared to HTLBM, allowing to manage substantially larger computation domains. Performance of the DPTLA implementation on a computing device featuring a GK110 GPU is about 1300 MLUPS on average, and might be significantly higher on future hardware such as the GK210 which will provide up to 112 KiB of local shared memory instead of 48 KiB with the GK110. This larger shared memory should result in a higher occupancy rate and therefore may have a positive impact on the local synchronisation issue.

In future, we will focus on designing and implementing boundary conditions for DPTLA, with the aim of performing thermal flow simulations involving complex geometries. Special care will be taken in order to minimise the impact of these boundary conditions on performance. We also plan to extend our single-GPU DPTLA solver to multi-GPU using an MPI-based approach similar to the one presented in [8].

Ra	N	DPTLA		HTLA		HTLBM	
		Nu	Δ	Nu	Δ	Nu	Δ
10^4	128	2.06360	0.46 %	2.06422	0.49 %	2.05984	0.27 %
	256	2.06131	0.35 %	2.06147	0.35 %	2.05916	0.24 %
	384	2.05974	0.27 %	2.05995	0.28 %	2.05862	0.22 %
	512	2.05683	0.13 %	2.05836	0.20 %	—	—
10^5	128	4.34411	0.16 %	4.33908	0.05 %	4.34320	0.14 %
	256	4.33501	−0.05 %	4.33364	−0.08 %	4.33120	−0.13 %
	384	4.33288	−0.10 %	4.33243	−0.11 %	4.33034	−0.15 %
	512	4.33250	−0.10 %	4.33248	−0.10 %	—	—
10^6	128	8.68767	0.54 %	8.56856	−0.83 %	8.82661	2.15 %
	256	8.65080	0.12 %	8.61264	−0.32 %	8.66098	0.23 %
	384	8.64038	0.00 %	8.62276	−0.21 %	8.64082	0.00 %
	512	8.63726	−0.04 %	8.62787	−0.15 %	—	—
10^7	256	16.27012	−0.44 %	15.94519	−2.43 %	16.71278	2.26 %
	384	16.37431	0.19 %	16.14409	−1.22 %	16.46645	0.76 %
	512	16.36820	0.16 %	16.22397	−0.73 %	—	—

Table 2: Nusselt numbers (Nu) at the isothermal hot wall obtained by DPTLA, HTLA, and HTLBM, and relative discrepancy (Δ) with respect to reference data

Ra	10^4		10^5		10^6		10^7	
u_m	16.6728	0.28 %	44.2835	0.86 %	128.4355	1.14 %	406.1065	5.50 %
	0.0195		-0.1836		-0.3060		-0.3763	
	0.0039	0.0040	0.2227	0.0032	0.3008	0.0026	0.3711	0.3711
	0.3242		0.3893		0.4388		0.4674	
v_m	2.1720	0.08 %	9.7936	0.98 %	25.6421	0.30 %	85.9780	3.00 %
	0.3841		0.4180		0.4518		-0.3320	
	0.2826	0.0018	0.3398	0.0017	0.3971	0.0017	0.4102	0.0026
	0.3451		0.3815		0.4180		0.3971	
w_m	18.9823	0.71 %	71.3399	0.38 %	237.0305	0.13 %	770.2316	0.28 %
	0.3841		0.4310		0.4622		0.4779	
	0.2357	—	0.3737	0.0039	0.4310	0.0023	0.4596	0.0388
	0.0195		0.0039		0.0273		-0.0065	

Table 3: Value and location of the maxima of the velocity components obtained by DPTLA, relative discrepancy of the value and normalised distance of the location with respect to reference data.

Ra	10^4	10^5	10^6	10^7
u_m	16.6800	0.24 %	128.6368	1.29 %
	0.0195	−0.1836	−0.3060	−0.3763
	0.0039	0.0040	0.3008	0.0026
	0.3242	0.3893	0.4388	0.0013
v_m	0.4701			
	2.1737	0.79 %	25.8183	0.98 %
	0.3841	0.4180	0.4518	−0.3294
	0.2826	0.0018	0.3971	0.0017
w_m	0.3451	0.3815	0.4180	0.4023
	18.9843	0.00 %	237.2971	0.24 %
	0.3841	0.4310	0.4622	0.4779
	0.2357	—	0.4310	0.0023
	0.0195	0.0039	0.0273	0.4596
				0.0117

Table 4: Value and location of the maxima of the velocity components obtained by DPTLA, relative discrepancy of the value and normalised distance of the location with respect to reference data.

Ra	10^4		10^5		10^6		10^7	
u_m	16.6707	0.29%	44.2512	0.79%	127.5759	0.47%	377.2339	1.74%
	0.0195		-0.1836		-0.3060		-0.3737	
	0.0013	0.0015	0.2201	0.0008	0.2982	0.0016	0.3659	0.3659
	0.3242		0.3867		0.4362		0.4648	
v_m	2.1661	0.44%	9.7902	0.95%	25.7571	0.75%	82.8791	0.63%
	0.3815		0.4180		0.4518		-0.3320	
	0.2826	0.0009	0.3398	0.0015	0.3971	0.0018	0.4076	0.0015
	0.3451		0.3789		0.4154		0.3945	
w_m	18.9468	0.19%	71.2223	0.0022	236.7128	0.00%	766.0296	0.27%
	0.3841		0.4310		0.4596		0.4779	
	0.2331	—	0.3737	0.0065	0.4310	0.0023	0.4596	0.0055
	0.0195		0.0065		0.0247		0.0378	

Table 5: Value and location of the maxima of the velocity components obtained by HTLBM, relative discrepancy of the value and normalised distance of the location with respect to reference data.

References

- [1] P. Asinari, T. Ohwada, E. Chiavazzo, and A.F. Di Rienzo. Link-wise Artificial Compressibility Method. *Journal of Computational Physics*, 231(15):5109–5143, 2012.
- [2] A. J. Chorin. A numerical method for solving incompressible viscous flow problems. *Journal of Computational Physics*, 2(1):12–26, 1967.
- [3] D. d’Humières, I. Ginzburg, M. Krafczyk, P. Lallemand, and L.S. Luo. Multiple-relaxation-time lattice Boltzmann models in three dimensions. *Philosophical Transactions of the Royal Society A*, 360:437–451, 2002.
- [4] P. Lallemand and L.-S. Luo. Theory of the lattice Boltzmann method: Acoustic and thermal properties in two and three dimensions. *Physical Review E*, 68(3):36706(1–25), 2003.
- [5] Nvidia. *Compute Unified Device Architecture Programming Guide version 6.5*, August 2014.
- [6] C. Obrecht, F. Kuznik, B. Tourancheau, and J.-J. Roux. The TheLMA project: a thermal lattice Boltzmann solver for the GPU. *Computers & Fluids*, 54:118–126, 2012.
- [7] C. Obrecht, F. Kuznik, B. Tourancheau, and J.-J. Roux. Multi-GPU Implementation of a Hybrid Thermal Lattice Boltzmann Solver using the TheLMA Framework. *Computers & Fluids*, 80:269–275, 2013.
- [8] C. Obrecht, F. Kuznik, B. Tourancheau, and J.-J. Roux. Scalable Lattice Boltzmann Solvers for CUDA GPU Clusters. *Parallel Computing*, 39(6-7):259–270, 2013.
- [9] C. Obrecht, P. Asinari, F. Kuznik, and J.-J. Roux. High-performance implementations and large-scale validation of the link-wise artificial compressibility method. *Journal of Computational Physics*, 275:143–153, 2014.
- [10] C. Obrecht, F. Kuznik, G. Rusaouën, and J.-J. Roux. Towards high-performance thermal flow solvers based on the link-wise artificial compressibility method. In *Proceedings of the 15th International Heat Transfer Conference*, 2014.
- [11] E. Tric, G. Labrosse, and M. Betrouni. A first incursion into the 3D structure of natural convection of air in a differentially heated cubic cavity, from accurate numerical solutions. *International Journal of Heat and Mass Transfer*, 43(21):4043–4056, 2000.