XLVI CILAMCE
VITÓRIA - ES - BRAZIL - 2025
XLVI IBERO-LATIN AMERICAN CONGRESS ON COMPUTATIONAL METHODS IN ENGENEERING

234 - GPU-Accelerated High-Order Regularized Lattice Boltzmann Simulations of Lid-Driven Cavity Flows

Gustavo T. Choaire[1], Bruno Y. dos Anjos[1], Luiz A. Hegele Jr.[1]

[1]*Dept. of Petroleum Engineering, Santa Catarina State University*
Lourival Cesario Pereira Avenue, 88336-275, Santa Catarina, Brasil
gustavo.choaire@udesc.br, bruno.anjos@edu.udesc.br, luiz.hegele@udesc.br

Abstract. We use high-order moment approximations within the moment representation lattice Boltzmann method (MR-LBM) to simulate lid-driven cavity (LDC) flows, assessing the performance of D3Q27 velocity stencil. Numerical experiments were conducted for Reynolds numbers ranging from 3,200 to 10,000 as part of the validation process. The MR-LBM implementation utilizes a regularized formulation of the lattice Boltzmann equation combined with the Bhatnagar-Gross-Krook (BGK) collision operator. Particle populations are reconstructed based on moments of the distribution function, with third and fourth-order moments approximated through lower-order moment relationships. This approach enables extended regularization up to fourth-order terms, significantly enhancing numerical stability and accuracy without compromising computational efficiency. Additionally, the incompressible regularized boundary condition (IRBC) was adopted. The systematic expansion of the distribution function to arbitrary orders was facilitated by Hermite polynomials, which provided a recursive framework for computing high-order terms. Consequently, the developed LBM framework accurately captured flow features at high Reynolds numbers. Leveraging the explicit nature of LBM equations allowed straightforward and efficient implementation on graphics processing units (GPUs), enabling highly accurate large-scale simulations with fine grid resolutions. Mean and root mean square (RMS) velocity profiles from the simulations demonstrated good agreement with experimental data from the literature. Building upon this validation, the methodology was extended to higher Reynolds numbers—specifically 10,000, 15,000, and 50,000—highlighting the robustness and versatility of the proposed MR-LBM approach across a wide range of complex flow conditions.

Keywords: Lattice Boltzmann Method, High-order Regularization, Lid-driven Cavity, High-Reynolds numbers flows.

1 Introduction

The lattice Boltzmann method (LBM) has emerged as a widely adopted numerical approach for accurately simulating fluid flows. This method models flow dynamics through the time evolution of microscopic velocities, utilizing lattice stencils to represent the probability of particle interactions among neighboring lattice nodes. Since collisions are solely responsible for altering the velocity distribution function, a collision model is employed to capture this mechanism. The most commonly used model is the Bhatnagar–Gross–Krook (BGK) operator [1], which is suitable for simulating weakly compressible flows but exhibits numerical instability at high Reynolds numbers.

The lattice Boltzmann method (LBM) has undergone continuous development, with numerous advancements introduced to improve boundary conditions, accuracy, and stability—particularly when employing the BGK operator. Among these contributions, the moment representation lattice Boltzmann method (MR-LBM) stands out. Initially proposed by [2], prior to the introduction of the regularized LBM, it gained broader recognition following the work of [3]. The MR-LBM framework stores macroscopic moments rather than distribution function populations, thereby reducing memory usage and enabling the collision step to be executed in moment space. The regularization procedure, which enhances accuracy and stability in the hydrodynamic regime, was first introduced by [4] and later refined by [3].

In addition to the regularization procedure and moment representation, boundary condition formulations offer further opportunities for improvement. Hegele Jr et al. [5] proposed the regularized boundary condition (RBC) for Dirichlet-type problems, which employs second-order moments to analytically solve a system of equations, thereby improving stability at high Reynolds numbers. Building upon this approach, Ferrari et al. [6] introduced

the incompressible regularized boundary condition (IRBC), which incorporates the incompressibility constraint ($\boldsymbol{\nabla} \cdot \boldsymbol{u} = 0$) and reduces the number of unknown second-order moments, further enhancing the stability of the method.

Despite previous advancements, Malaspinas [7] introduced a high-order LBM scheme that enhances both stability and accuracy. This formulation is grounded in the observation that the velocity distribution function can be expressed recursively, enabling the reconstruction of non-equilibrium moments of the Boltzmann–BGK equation at arbitrary order using only lower-order moments.

This work presents a high-order MR-LBM framework, incorporating terms up to the fourth order, for simulating high Reynolds number flows using a D3Q27 lattice stencil and the incompressible regularized boundary condition (IRBC) for Dirichlet-type boundaries. The choice of the D3Q27 stencil in this work is motivated by the fact that, within the discrete Boltzmann framework, the D3Q19 stencil does not preserve Hermite orthogonality. This limitation may lead to artifacts in the simulation, since the calculation of moments may incorrectly include spurious contributions. For example, the summation $\sum_i w_i H_{xxy,i} H_{yzz,i} \neq 0$. Owing to the explicit nature of the LBM equations, graphics processing units (GPUs) have recently been adopted to facilitate efficient and straightforward implementations. Accordingly, an in-house GPU-based solver was developed, integrating all the LBM enhancements discussed above, to simulate high Reynolds number regimes within the lid-driven cavity (LDC) configuration.

2 Mathematical Modeling

The lattice Boltzmann method (LBM) is a discrete kinetic model based on the evolution of particle distribution functions. The lattice evolution is governed by the following expression:

$$f_i\left(\mathbf{r} + \mathbf{c}_i \Delta t, t + \Delta t\right) = f_i^{(eq)}\left(\mathbf{r}, t\right) + \left(1 - \frac{\Delta t}{\tau}\right) \hat{f}_i^{(neq)}\left(\mathbf{r}, t\right), \tag{1}$$

where $\mathbf{r}$, $\mathbf{c}_i$, and t denote the position, discrete microscopic velocity, and time, respectively, in lattice units [8, 9]. In the present formulation, the distribution function is regularized using macroscopic moments, which include local density, momentum, and second-order moments [4]:

$$\{\rho, \rho u_\alpha, \rho m_{\alpha\beta}^{(2)}\} = \sum_i g_i \{1, c_{i\alpha}, c_{i\alpha} c_{i\beta} - \delta_{\alpha\beta}/a_s^2\}, \tag{2}$$

where $\delta_{\alpha\beta}$ is the Kronecker delta, $a_s = \sqrt{3}$ is a scaling parameter, and g_i refers to either the raw distribution f_i or the regularized distribution $\hat{f}_i$. Combined with the regularization procedure, the moment representation LBM (MR-LBM) employs the regularized distribution function defined as:

$$\hat{f}_i^N = \rho w_i \left(1 + a_s^2 u_\alpha c_{i\alpha} + \frac{1}{2} a_s^2 m_{\alpha\beta}^{(2)} \mathscr{H}_{\alpha\beta,i}^{(2)} + \sum_{n=3}^{N} \frac{1}{b!} a_s^n m_{\boldsymbol{\alpha}}^{(n)} \mathscr{H}_{\boldsymbol{\alpha},i}^{(n)}\right), \tag{3}$$

where ρ is the fluid density, u_α is the macroscopic velocity, $m_{\alpha\beta}^{(2)}$ denotes the second-order moment, and N specifies the chosen order of approximation. The terms w_i are the lattice weights, $c_{i\alpha}$ are the discrete velocity components, and $\mathscr{H}_{\alpha\beta}^{(2)} = c_{i\alpha}\, c_{i\beta} - \delta_{\alpha\beta}/a_s^2$ denotes the second-order Hermite polynomial.

After the population reconstruction, the streaming step is applied:

$$f_i(\mathbf{x} + \mathbf{c}_i \Delta t, t + \Delta t) = f_i^*(\mathbf{x}, t), \tag{4}$$

where f_i^* refers to the post-collision distribution function.

The collision step is performed in moment space using the Bhatnagar–Gross–Krook (BGK) operator [1], which updates the distribution function according to:

$$f_i^* = \hat{f}i + \omega(fi^{(eq)} - \hat{f}_i), \tag{5}$$

where ω represents the relaxation frequency. The relationship between ω and the kinematic viscosity ν is given by:

$$\omega = \left(a_s^2 \nu + \frac{\Delta t}{2}\right)^{-1}. \tag{6}$$

The collision step modifies the velocity distribution function, guiding the system toward local equilibrium. After this step, the non-equilibrium moments are recalculated, and the regularization procedure is reapplied to

preserve consistency and enhance numerical stability.

$$\hat{f}_i\left(\mathbf{r},t\right)=f_i^{(eq)}\left(\mathbf{r},t\right)+\hat{f}_i^{(neq)}\left(\mathbf{r},t\right). \tag{7}$$

The definition of the non-equilibrium distribution function is crucial for implementing the high-order scheme, as this work adopts the method proposed by Malaspinas [7]. This approach defines a recursive formulation that expresses higher-order moments as functions of second-order moments, while preserving accuracy and even improving the computational stability of the method.

A key advantage of this technique is its strong compatibility with the moment representation framework, as it does not require the storage of additional moments in memory, resulting in a lightweight and more efficient implementation. The moments can be expressed as:

$$a^{(n)}_{1,\alpha_1\ldots\alpha_n}=a^{(n-1)}_{1,\alpha_1\ldots\alpha_{n-1}}u_{\alpha_n}+\left(u_{\alpha_1}\ldots u_{\alpha_{n-2}}a^{(2)}_{1,\alpha_{n-1}\alpha_n}+perm\left(\alpha_n\right)\right), \tag{8}$$

where $a_1^{(n)}$ represents the non-equilibrium moments of order n, and perm(α_n) denotes all permutations of the term in the α_n coordinate, as described by Malaspinas [7].

3 Validation

The validation procedure is conducted by solving the lid-driven cavity (LDC) benchmark for two Reynolds number regimes: $Re = 3,200$ and $Re = 10,000$. Multiple mesh resolutions are employed to assess both grid convergence and the ability of high-order schemes to accurately capture high Reynolds number flows on coarse grids.

Centerline mean velocity profiles are compared with the experimental data reported by Prasad and Koseff [10]. The lid velocity is set to $U_{\text{lid}} = 0.0256$ lattice units per time step (lu/ts), and simulations are run up to $t^* = 500$, where $t^* = tU_{\text{lid}}/L$.

The first parameter used to validate the model is the total kinetic energy per unit mass, defined as $E_k^* = \sum\left(u_\alpha u_\alpha\right)/2L^3U_{\text{lid}}^2$, with results presented in Figure 1.

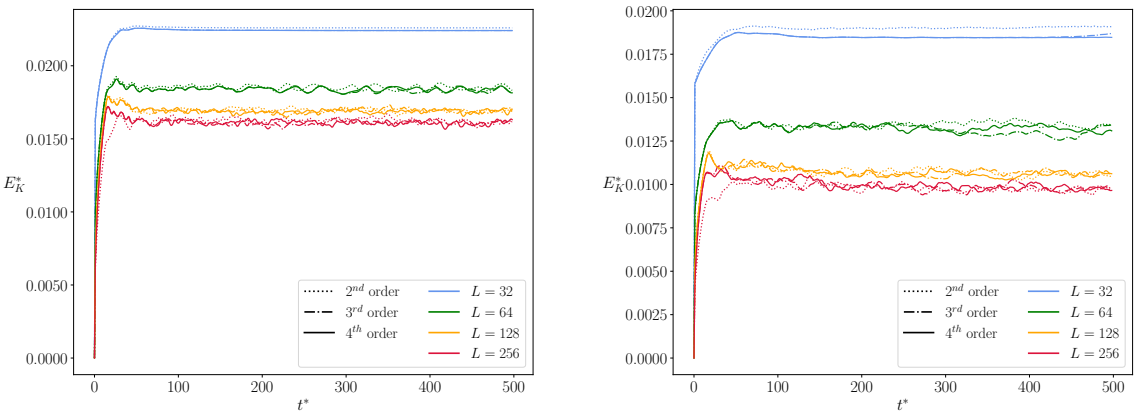


Figure 1. Temporal evolution of total kinetic energy E_k^* for various grid sizes and orders. (Left) Re = 3,200. (Right) Re = 10,000

As observed, the total kinetic energy decreases with increasing Reynolds number—an expected behavior that has been previously reported in the literature. Ferrari et al. [6] attributed this trend to the presence of more coherent structures within the cavity at lower Reynolds numbers, which help maintain higher kinetic energy and a thicker boundary layer along the lid surface.

Another important aspect of the validation process involves the grid convergence analysis. Total kinetic energy also offers valuable insights in this context, as increasing the grid resolution leads to a convergence of the E_k^* curves. As grid size increases, these curves tend to collapse onto one another, indicating that the numerical schemes are consistently capturing the same complex flow features within the cavity.

This behavior is evident in the proximity of the curves for $L = 128$ and $L = 256$, suggesting that further refinement (e.g., $L = 512$) would likely result in even closer agreement. Still referring to Figure 1, for the high

Reynolds number case, the energy curves are noticeably more dispersed compared to the lower Reynolds case, indicating that even finer grids may be required to accurately resolve such flows.

The second part of the validation consists of evaluating the L2 error associated with the total kinetic energy, E_k^*, using the following equation:

$$E_{L2}^* = \sqrt{\frac{\left(E_{t,b} - E_{t,s}\right)^2}{E_{t,b}^2}}, \tag{9}$$

where $E_{t,b}$ is the total kinetic energy of the reference benchmark (e.g., the highest resolution available $L = 512$), and $E_{t,s}$ is the total kinetic energy of a given resolution.

Since turbulent flows require the analysis of variables in a statistically stationary state—where the properties exhibit minimal variation over time—the L2 norm was evaluated within the interval $t^* = 300$ to $t^* = 500$, thereby avoiding the effects of the initial impulsive start of the flow.

Tables 1 and 2 present the L2 error results for different scheme orders and grid resolutions, across both Reynolds number regimes.

Table 1. E_{L2}^* for $Re = 3{,}200$

$L/Order$	32	64	128	256
2^{nd} order	41.67%	16.03%	6.12%	2.06%
3^{rd} order	40.52%	14.97%	6.53%	1.64%
4^{th} order	40.51%	15.27%	5.89%	1.29%

Table 2. E_{L2}^* for $Re = 10{,}000$

$L/Order$	32	64	128	256
2^{nd} order	101.86%	43.14%	12.87%	3.35%
3^{rd} order	95.99%	37.27%	13.22%	4.51%
4^{th} order	95.63%	39.91%	12.21%	4.58%

Coarser grids were intentionally selected to stress the model, despite not yielding fully accurate results. As shown in Tables 1 and 2, high-order schemes produced slightly lower errors on coarser grids. However, in some cases, higher-order schemes resulted in greater relative errors compared to the second-order scheme—for instance, the third-order scheme with $L = 128$ for both $Re = 3,200$ and $Re = 10,000$, and the third- and fourth-order schemes with $L = 256$ for $Re = 10,000$. Two possible explanations for this behavior are proposed.

The first explanation concerns the range of resolvable frequencies during the time-averaging period. Executing the time-averaging operation at a higher frequency (i.e., at every time step) could improve the convergence of the time-averaged solution, as a higher sampling frequency increases the upper cutoff frequency that can be captured by the averaging process. Additionally, extending the total sampling period would improve the representation of low-frequency components, since the duration over which time averaging is performed determines the lowest frequency that can be reliably resolved.

The second explanation concerns the benchmark used to evaluate the total kinetic energy results. The reference solution was obtained using a second-order scheme with a high-resolution grid of $L = 512$. Although grid refinement reduces numerical errors for all approximation orders, a more appropriate comparison would involve generating reference solutions at $L = 512$ for each scheme individually—i.e., computing results at $L = 512$ for both third- and fourth-order schemes and comparing them accordingly.

The last part of the validation presents the mean centerline and RMS velocity profiles for both Reynolds regimes and higher-order schemes. Figure 2 shows the mean centerline velocity profiles for $Re = 3{,}200$ and $Re = 10{,}000$, using a grid resolution of $L = 64$. Each numerical scheme is represented by different markers and line styles.

For the case of $Re = 3{,}200$, the second- and third-order schemes produce similar results at the bottom and right (impingement) walls, whereas the fourth-order scheme shows a small deviation from both the lower-order schemes and the experimental data. At the bottom wall, the second- and third-order approximations closely match the experimental results, while at the impingement wall the opposite occurs—only the fourth-order scheme aligns better with the experimental comparison.

A different trend is observed for the higher Reynolds number case of $Re = 10{,}000$. At the bottom and left walls of the LDC, the higher-order schemes better approximate the expected results. However, at the impingement wall, they exhibit a slightly larger deviation from the experimental data.

Figure 3 illustrates the RMS velocity profiles along the centerlines in the x and y directions at $z^* = 0.5$ (red and blue lines), where $z^* = z/L_z$, as well as the cross-product of velocity fluctuations, $u'_x u'_y$ (green lines). All results were obtained using a grid resolution of $L = 64$ to challenge the performance of the model. In order to separately assess the behavior on coarser grids and to better visualize the differences between approximation orders under reduced resolution, the results for $Re = 10{,}000$ obtained with the finest grid ($L = 256$) are presented in the results section.

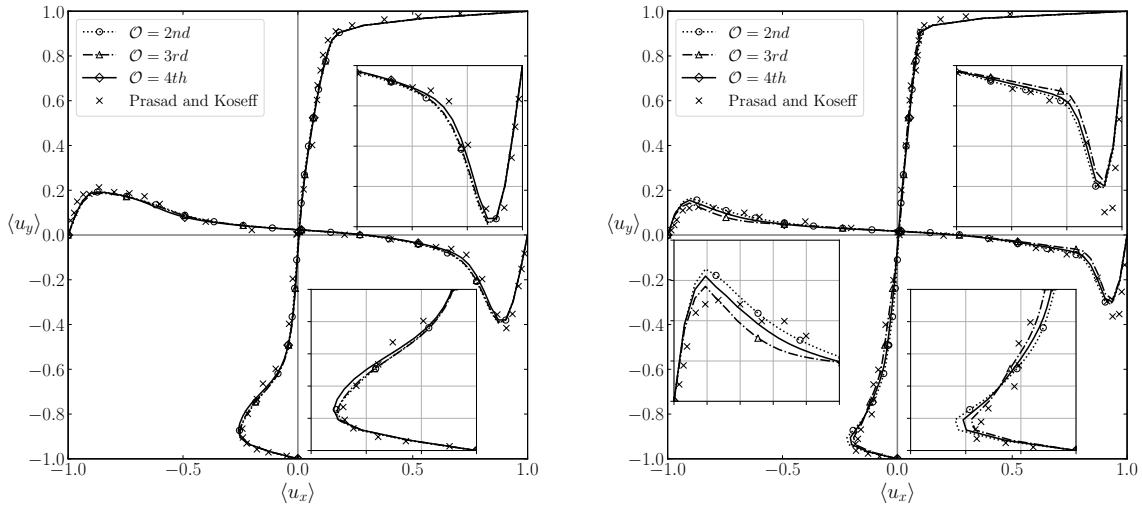


Figure 2. Mean velocity profiles for $Re = 3{,}200$ (left) and $Re = 10{,}000$ (right) at the centerlines using different orders and a $L = 64$ grid resolution.

As previously discussed for the mean velocity profiles, the inset zooms within the figure highlight the ability of higher-order schemes to capture steep gradients near the walls at higher Reynolds numbers. For the lower Reynolds number case, the second-order scheme outperforms both the third- and fourth-order approximations across all walls. Conversely, as the Reynolds number increases, higher-order schemes tend to capture the steeper gradients more accurately.

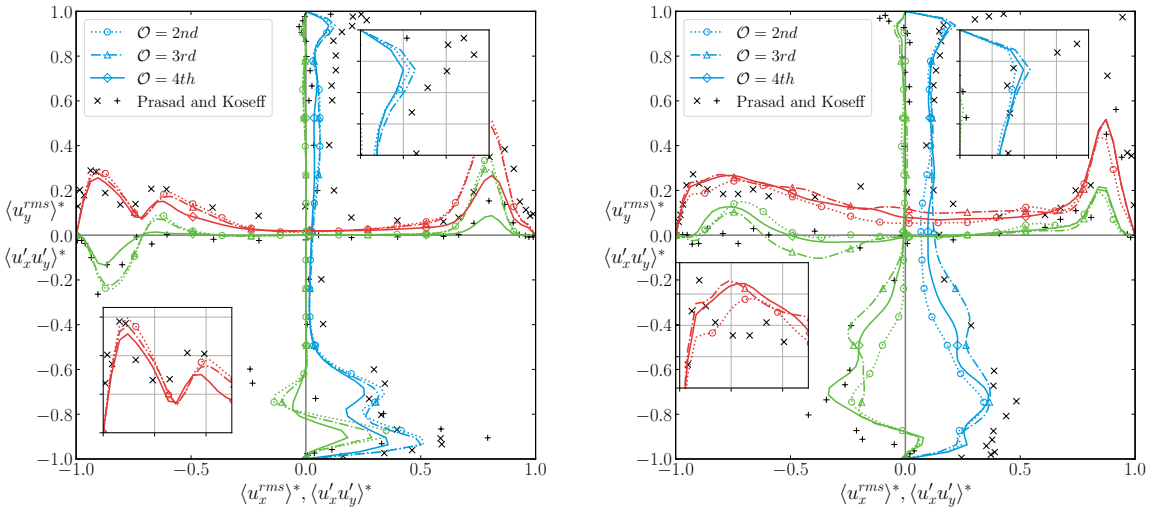


Figure 3. RMS profiles (red and blue) and cross products of velocities profiles (green), along the x and y coordinates at $z^* = 0.5$, for $Re = 3{,}200$ (left) and $Re = 10{,}000$ (right) and $L = 64$ grid resolution.

4 Results

The results presented correspond to higher Reynolds numbers of $Re = 15{,}000$ and $Re = 50{,}000$. For these cases, no experimental data are available for direct comparison. However, since the $L = 256$ grid resolution

yielded excellent agreement with experimental results for the $Re = 10{,}000$ case, it was adopted to compute the final results at higher Reynolds numbers.

Both mean centerline velocity profiles and RMS profiles were analyzed. Figure 4 shows the results for $Re = 10{,}000$, where experimental comparison is still possible. The mean velocity profiles demonstrate excellent agreement with the reference data, with all numerical schemes—regardless of order—performing similarly well in this case.

For the RMS and cross-product velocity profiles, the higher-order schemes exhibited superior performance near the walls, where steeper gradients are typically found. This behavior was observed not only for higher Reynolds numbers using coarser grids, but also with refined resolutions, reinforcing the suitability of high-order schemes for accurately resolving complex near-wall regions.

The results for $Re = 15{,}000$ are presented in Figure 5. Unlike the $Re = 10{,}000$ case, the solutions obtained using different order schemes show noticeable discrepancies in specific regions of the LDC, particularly near the bottom and right (impingement) walls. This suggests that the selected grid resolution may not be sufficiently fine to accurately resolve the flow at such a high Reynolds number.

A similar pattern is observed in the RMS and cross-product velocity profiles. As previously discussed, higher-order schemes are generally more effective at capturing near-wall behavior; however, in this case, the second-order scheme appears to perform comparably—or even slightly better—than the third- and fourth-order schemes.

This outcome may be attributed to the increased flow complexity at higher Reynolds numbers, where more intense turbulence and sharper gradients are expected to occur.

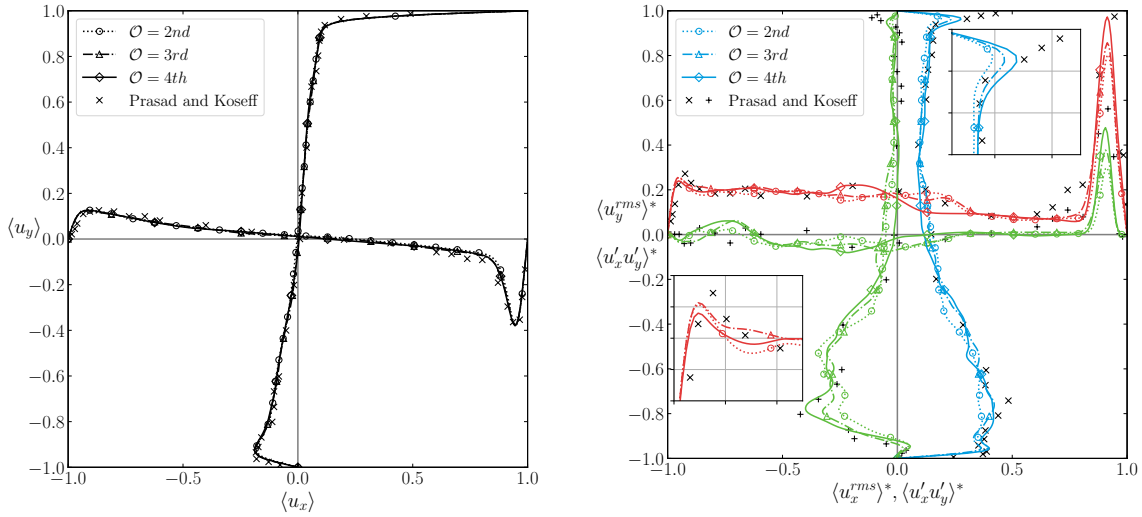


Figure 4. Mean velocity profiles (left) and RMS velocity profiles (right), for $Re = 10{,}000$ at the centerlines using different orders and $L = 256$ grid resolution.

The final set of results corresponds to the $Re = 50{,}000$ flow. Given the discrepancies observed among the different approximation orders for the mean velocity profiles at $Re = 15{,}000$ (Figure 5), similar behavior was expected at this higher Reynolds number. However, as shown in Figure 6, the results obtained with the various approximation orders collapse onto one another. This outcome may be attributed to the $L = 256$ grid resolution being insufficient to fully resolve the flow at this regime, or to the relatively short time-averaging interval used.

On the other hand, consistent with findings from Hegele Jr et al. [5] and Ferrari et al. [6], an increase in Reynolds number leads to a reduction in the minimum peak of u_x near the bottom wall—a feature captured by all schemes, regardless of order.

The RMS components and cross-products of velocity fluctuations also showed very similar results across the different approximation orders, with only minor differences observed near the bottom and right walls. As noted in Hegele Jr et al. [5], a transitional behavior occurs beyond $Re = 20{,}000$, altering the velocity profiles in near-wall regions—an effect also confirmed in Ferrari et al. [6]. The central region of the LDC continues to exhibit the lowest levels of velocity fluctuation. As previously pointed out by Bouffanais et al. [11], the experimental data were reported without error bars, which prevents assessing whether the numerical results fall within the uncertainty range of the experiments.

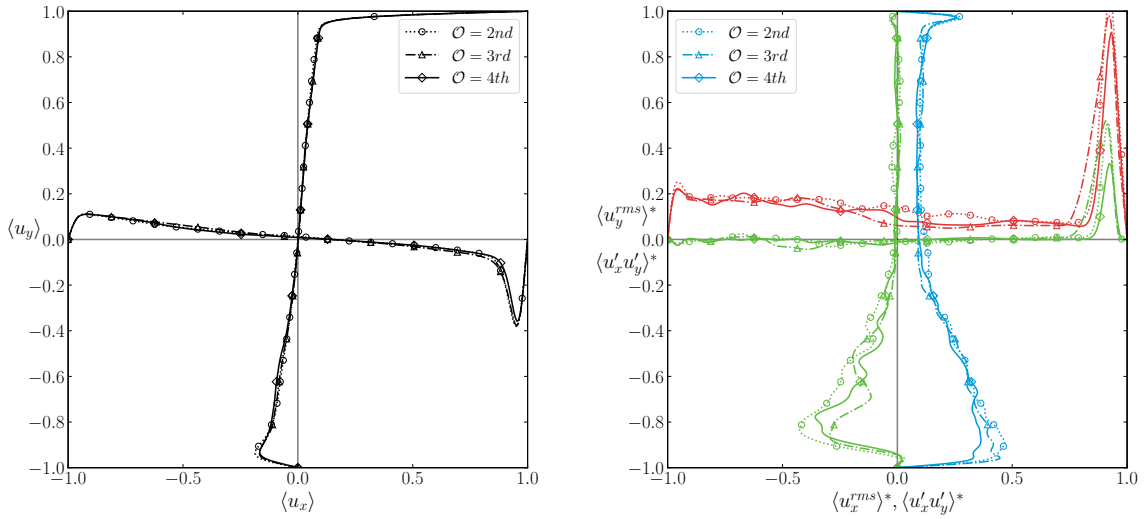


Figure 5. Mean velocity profiles (left) and RMS velocity profiles (right), for $Re = 15{,}000$ at the centerlines using different orders and $L = 256$ grid resolution.

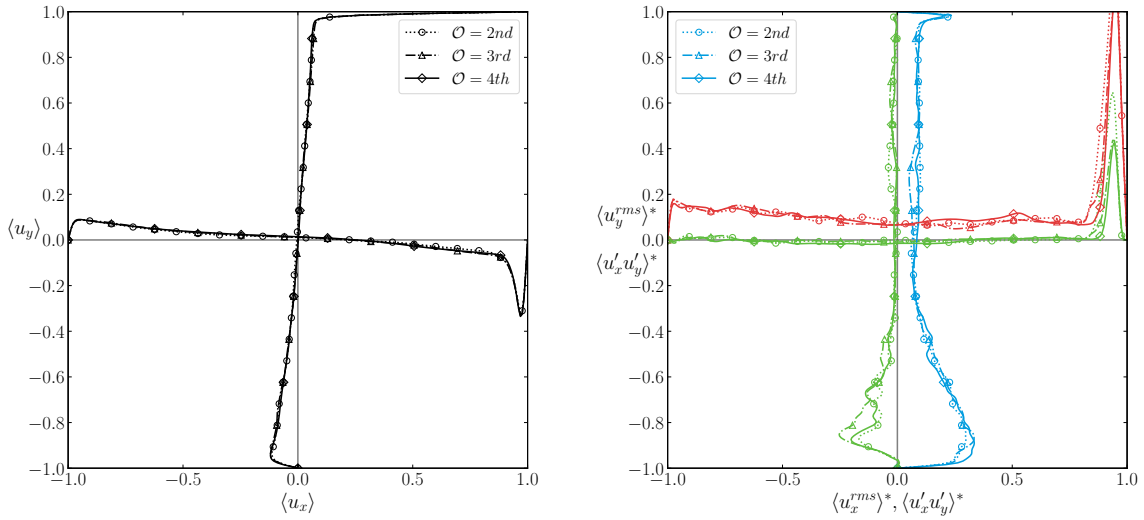


Figure 6. Mean velocity profiles (left) and RMS velocity profiles (right), for $Re = 50{,}000$ at the centerlines using different orders and $L = 256$ grid resolution.

Finally, the mega lattice updates per second (MLUPS) metric was employed to evaluate the computational performance of the proposed framework. The simulations were performed on an NVIDIA RTX 4090 GPU, using single-precision floating-point arithmetic. For the higher Reynolds number regimes presented in this section ($Re = 15{,}000$ and $Re = 50{,}000$), and using the $L = 256$ grid resolution, the results for $Re = 15{,}000$ were 4,700, 4,691, and 4,683 MLUPS for the second-, third-, and fourth-order schemes, respectively.

Similarly, for the $Re = 50{,}000$ case, performance remained consistent, yielding 4,623 (second-order), 4,705 (third-order), and 4,693 (fourth-order) MLUPS. These results indicate that the adoption of higher-order schemes within the proposed approach did not lead to any significant degradation in computational efficiency. Although the MR-LBM achieves higher performance overall, the D3Q27 stencil involves a larger velocity set, which inherently reduces efficiency compared to the D3Q19 stencil.

5 Conclusions

This study presented a numerical investigation of the three-dimensional lid-driven cavity (LDC) benchmark using the moment representation lattice Boltzmann method (MR-LBM), with an emphasis on high-order schemes, including up to fourth-order approximations. The in-house code was developed to run on CUDA-based GPU architectures, enabling accurate and efficient computations.

The results were validated against the experimental data of Prasad and Koseff [10], showing excellent agreement for Reynolds numbers up to $10{,}000$. Grid convergence analysis, supported by the L2 norm, confirmed that the error decreased monotonically across all approximation orders. In most cases, the higher-order schemes yielded lower error percentages.

The use of high-order schemes proved effective in capturing steep near-wall gradients in high Reynolds number flows, producing robust and reliable results. Furthermore, performance evaluation using mega lattice updates per second (MLUPS) indicated that higher-order approximations do not significantly impact computational efficiency, supporting their applicability in complex flow regimes. The proposed methodology demonstrated enhanced accuracy for high Reynolds number flows, without compromising computational performance.

Acknowledgements. This study is part of the project "Experiment-based physical and numerical modeling of subsea oil jet dispersion", a cooperation agreement between UDESC and PETROBRAS (number 0050.0125248.23.9), governed by the investment clauses in P, D&I of the Brazilian National Agency of Petroleum, Natural Gas and Biofuels (ANP) (Resolution ANP 918/2023).

Authorship statement. The authors hereby confirm that they are the sole liable persons responsible for the authorship of this work, and that all material that has been herein included as part of the present paper is either the property (and authorship) of the authors, or has the permission of the owners to be included here.

References

[1] P. L. Bhatnagar, E. P. Gross, and M. Krook. A model for collision processes in gases. i. small amplitude processes in charged and neutral one-component systems. *Physical review*, vol. 94, n. 3, pp. 511, 1954.

[2] R. Argentini, A. Bakker, and C. Lowe. Efficiently using memory in lattice boltzmann simulations. *Future Generation Computer Systems*, vol. 20, n. 6, pp. 973–980, 2004.

[3] M. Vardhan, J. Gounley, L. Hegele, E. W. Draeger, and A. Randles. Moment representation in the lattice boltzmann method on massively parallel hardware. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, pp. 1–21, 2019.

[4] J. Latt and B. Chopard. Lattice boltzmann method with regularized pre-collision distribution functions. *Mathematics and Computers in Simulation*, vol. 72, n. 2-6, pp. 165–168, 2006.

[5] L. Hegele Jr, A. Scagliarini, M. Sbragaglia, K. Mattila, P. Philippi, D. Puleri, J. Gounley, and A. Randles. High-reynolds-number turbulent cavity flow using the lattice boltzmann method. *Physical Review E*, vol. 98, n. 4, pp. 043302, 2018.

[6] M. Ferrari, A. Franco, and L. Hegele. Turbulent flow inside a cubic lid-driven cavity using moment representation lattice boltzmann method. *Physics of Fluids*, vol. 36, n. 7, 2024.

[7] O. Malaspinas. Increasing stability and accuracy of the lattice boltzmann scheme: recursivity and regularization. *arXiv preprint arXiv:1505.06900*, 2015.

[8] Y.-H. Qian, D. d'Humières, and P. Lallemand. Lattice bgk models for navier-stokes equation. *Europhysics letters*, vol. 17, n. 6, pp. 479, 1992.

[9] S. Succi. *The Lattice Boltzmann Equation for Fluid Dynamics and Beyond*. Oxford University Press, 2001.

[10] A. K. Prasad and J. R. Koseff. Reynolds number and end-wall effects on a lid-driven cavity flow. *Physics of Fluids A: Fluid Dynamics*, vol. 1, n. 2, pp. 208–218, 1989.

[11] R. Bouffanais, M. O. Deville, and E. Leriche. Large-eddy simulation of the flow in a lid-driven cubical cavity. *Physics of Fluids*, vol. 19, n. 5, 2007.