

TURBULENT FLOW SIMULATION WITH HIGH-ORDER REGULARIZED LATTICE BOLTZMANN METHOD USING D3Q27

Bruno Yan dos Anjos*, Gustavo Choaie* and Luiz A. Hegele Jr.*

* Department of Petroleum Engineering
Santa Catarina State University
Balneário Camboriú - SC 88336-275
e-mail: luiz.hegele@udesc.br

Abstract: We employ the lattice Boltzmann method (LBM) to perform flow simulations with the objective of evaluating the performance of the D3Q27 stencil in capturing complex flow structures. We used a specific variation of the method known as the Moment Representation Lattice Boltzmann Method (MR-LBM), which employs a second-order moment representation to regularize the distribution functions, thereby enhancing computational performance. The implementation incorporates the Bhatnagar-Gross-Krook (BGK) collision operator, which is embedded in the moment evaluation process. As benchmark cases, we considered the lid-driven cavity (LDC) flow and a turbulent jet, both of which require robust numerical treatment to handle boundary conditions and resolve small-scale flow features. Dirichlet boundary conditions are enforced at solid walls, and the implementation adopts the incompressible regularized boundary condition (IRBC), which reduces the number of constraints on the second-order moments, thereby improving the stability and efficiency of the simulations. Additionally, we incorporated a high-order regularization technique and enforced a zero-trace condition on the momentum flux tensor to further enhance numerical stability and accuracy. Numerical experiments were conducted for Reynolds numbers of 10,000, 15,000, 25,000, and 50,000, employing mesh sizes up to 256 lattice nodes per dimension, ensuring a comprehensive assessment of the performance of the method under different flow conditions. The results were evaluated based on root mean square (RMS) and mean velocity profiles across different mesh sizes and Reynolds numbers. The findings were validated with results from the literature and showed good agreement.

Keywords. High-Order Regularization, High Reynolds

1 INTRODUCTION

The lattice Boltzmann method (LBM) is a widely adopted computational fluid dynamics (CFD) framework that solves the Navier–Stokes equations (NSE) through a mesoscopic scheme, capable of simultaneously capturing microscopic effects and describing macroscopic fluid behavior. This method has proven effective for complex problems, particularly turbulent flows, as it provides a detailed physical representation of fluid motion [1].

Although historically associated with high computational demands, recent advances in GPU parallel programming, well aligned with the intrinsic parallelism of LBM, together with methodological improvements developed in recent years, have established LBM as a robust CFD approach, able to reproduce complex fluid mechanics scenarios with high fidelity [2].

The LBM framework originates from the continuous Boltzmann equation (BE), where a density distribution function f is advected by the particle velocity ξ , subject to external forces and governed by a collision operator Ω . When the BE is discretized, the resulting formulation is known as the lattice Boltzmann equation (LBE). In this approach, the continuous space is subdivided into a finite set of distributions, each associated with a discrete direction denoted by i . The number of these directions is determined by the velocity stencil selected for the problem. By combining the discrete distributions f_i with the microscopic velocities \mathbf{c}_i , it becomes possible to reconstruct the macroscopic physical properties of the flow, such as density, momentum, and the strain rate tensor, all consistent with the NSE, as demonstrated through the Chapman-Enskog expansion [3].

Since the LBM is derived from a discretization of a continuous equation, the resulting formulation is inherently approximate and inevitably introduces some degree of error. To improve accuracy, stability, and computational performance, numerous studies have proposed a variety of approaches to LBM. Among these, the regularized LBM is particularly relevant to the present work. In this formulation, the particle populations at each lattice site are expressed as combinations of their moments and Hermite polynomials, thereby enhancing the numerical stability of the method. In the present methodology, this approach is coupled with the Bhatnagar–Gross–Krook (BGK) collision operator [4].

Based on the principle of regularization, [5] proposed a system of equations that enables the calculation of moments at the boundaries. This approach, referred to as regularized boundary conditions (RBC), can be further improved by assuming incompressibility at the boundaries, which enforces the trace of the strain rate tensor to be zero. This assumption reduces the number of equations that must be solved at the boundaries, resulting in the incompressible regularized boundary conditions (IRBC) [6]. Using the RBC approach, [5] successfully simulated flows at Reynolds numbers up to 50,000, demonstrating its potential to enhance the stability of the method.

The moment representation of the LBM is employed as a computational strategy to enhance performance. In a three-dimensional simulation using the D3Q27 stencil, the classical formulation requires storing 27 populations per lattice site. By contrast, the moment-based LBM stores only 10 moments, which are sufficient to regularize the distributions during the pre-streaming step [7]. With this representation, the collision step is performed directly on the moments rather than on the populations, thereby eliminating the need to explicitly compute the equilibrium distribution. This approach further reduces the memory footprint and significantly improves computational efficiency. Moreover, the moment-based LBM integrates seamlessly with CUDA parallel programming, making it particularly well suited for high-performance implementations [8].

Building upon these developments, the present work applies the high-order regularization proposed by [9]. The objective is to investigate the computational impact of

high-order regularization on execution time and to assess its influence on simulation accuracy when compared with the classical second-order formulation. To conduct this study, the lid-driven cavity (LDC) flow was selected because it combines geometric simplicity with complex flow characteristics, and has been extensively examined in the literature, providing a reliable benchmark for comparison.

2 METHODS

To address the proposed problem, the LBM scheme was employed to obtain the results presented in this paper. The method calculates fluid behavior on a set of discrete points, known as lattice nodes. Each lattice node contains a set of density distribution functions, denoted as f_i , arranged in a predefined number of discrete directions. The number and orientation of these directions depend on the selected velocity stencil [9].

The temporal evolution of f_i within the domain is referred to as the streaming process, during which the distribution function propagates from one lattice node to its neighbor. When distributions meet at a lattice site after streaming, they interact through a collision process, altering the behavior of f_i . This collision process is denoted by Ω_i [5]. The combined dynamics of streaming and collision are expressed as:

$$f_i(\mathbf{x} + \mathbf{c}_i \Delta t, t + \Delta t) = f_i(\mathbf{x}, t) + \Omega_i(\mathbf{x}, t), \quad (1)$$

where \mathbf{c}_i is the microscopic velocity vector and \mathbf{x} represents the lattice position coordinate. The formulation of the collision operator varies in the literature, ranging from single-relaxation-time to multi-relaxation-time (MRT) models. In this work, the Bhatnagar–Gross–Krook (BGK) collision operator is employed. This model describes the relaxation of the distribution function toward its equilibrium state, with the rate of relaxation governed by a characteristic relaxation time constant, denoted as τ [10]. The BGK collision operator is expressed as:

$$\Omega_i(\mathbf{x}, t) = -\frac{\Delta t}{\tau} (f_i(\mathbf{x}, t) - f_i^{eq}(\mathbf{x}, t)). \quad (2)$$

The relaxation time τ can be determined through its correlation with the kinematic viscosity, denoted as ν . This relationship is established by the Chapman–Enskog expansion, which links the Lattice Boltzmann Equation (LBE) to the Navier–Stokes Equations (NSE) [3]. Accordingly, the kinematic viscosity is defined as:

$$\nu = a_s^{-2} \left(\tau - \frac{\Delta t}{2} \right), \quad (3)$$

In the previous equation, a_s denotes the scaling factor, defined to ensure that the microscopic velocities \mathbf{c}_i are unit vectors, and its value depends on the chosen stencil. For classical LBM stencils such as D2Q9, D3Q19, or D3Q27, this factor is equal to $\sqrt{3}$. In this case, $a_s = 1/c_s$, where c_s represents the speed of sound in the classical LBM scheme [5].

In the LBM framework, macroscopic moments can be obtained directly from the distribution functions by combining them with the microscopic velocities. In the classical

formulation, three fundamental moments are identified: the zeroth-order, first-order, and second-order moments. These are physically associated with density, momentum, and the strain rate tensor, respectively [4]. The first-order moment is given by ρ , followed by ρu_α and $\rho m_{\alpha\beta}^{(2)}$, which are expressed as:

$$\left\{ \rho, \rho u_\alpha, \rho m_{\alpha\beta}^{(2)} \right\} = \sum_i f_i \left\{ 1, c_{i\alpha}, H_{\alpha\beta,i}^{(2)} \right\}, \quad (4)$$

where $H_{\alpha\beta} = c_{i\alpha}c_{i\beta} - \delta_{\alpha\beta}/a_s^2$ is the Hermite polynomial function, employed in the LBM due to its orthogonality property, which allows the distribution functions to be expressed in terms of their moments combined with these polynomials. The term $\delta_{\alpha\beta}$ denotes the Kronecker delta [4]. Since this study adopts the regularized LBM, the use of this formulation is particularly important. Both the regularized distributions, denoted by \hat{f}_i , and the equilibrium distributions are expressed in terms of their moments. The primary difference between them is that the second-order moments are represented as combinations of the velocity components in the α and β directions, as follows:

$$\hat{f}_i = \rho w_i \left(1 + a_s^2 u_\alpha c_{i\alpha} + \frac{a_s^4}{2} m_{\alpha\beta}^{(2)} H_{\alpha\beta,i}^{(2)} \right), \quad (5)$$

$$f_i^{eq} = \rho w_i \left(1 + a_s^2 u_\alpha c_{i\alpha} + \frac{a_s^4}{2} u_\alpha u_\beta H_{\alpha\beta,i}^{(2)} \right). \quad (6)$$

The variable w_i represents the weight associated with a given direction, describing the probability of that direction occurring when predicting the flow behavior at a lattice node. The values of w_i are determined using Gauss quadrature and vary according to the selected velocity stencil [3].

2.1 Moment representation

The moment representation is a technique designed to reduce the computational cost of simulations. Its core principle is to store the moments instead of the full set of particle populations at each lattice site. For a three-dimensional simulation, the memory savings are directly proportional to the reduction in the number of stored variables per lattice node [7].

The moment-based LBM provides significant memory savings and demonstrates greater efficiency as the complexity of the geometry and the velocity stencil definition increases. These advantages motivate its adoption in the present work. Before describing the details of the formulation, it is noted that a collided distribution function is denoted with an asterisk (*), and its behavior is expressed as:

$$f_i^*(\mathbf{x}, t) = f_i(\mathbf{x} + \mathbf{c}_i \Delta t, t + \Delta t). \quad (7)$$

Therefore, Eq. (1) combined with Eq. (2) is used to calculate the collided moments. As noted previously, the primary difference between equilibrium and regularized distributions arises in the second-order expansion. By subtracting the distributions in the collision

operator, only second-order terms remain. Comparing both sides of the equation then reveals that only the second-order moments are modified. As a result, the collision in moment space can be expressed as:

$$m_{\alpha\beta}^* = (1 - \omega)m_{\alpha\beta}^{(2)} + \omega u_\alpha u_\beta, \quad (8)$$

where $\omega = \Delta t/\tau$. Once the collided moments are determined, all particle populations can be reconstructed using Eq. (5), after which they are streamed to neighboring nodes, and the process is repeated.

2.2 High Order Regularization

A high-order expansion occurs when the distribution function is not truncated at the second-order term [9]. In this case, the definition of the distribution function can be further exploited, allowing the regularized distribution to be expressed as:

$$\hat{f}_i^N = w_i \sum_{n=0}^N \frac{a_s^{2n}}{n!} \mathbf{M}^{(n)} : H_i^{(n)}. \quad (9)$$

where $\mathbf{M}^{(n)}$ represents the moment of order n , obtained by summing the distribution functions weighted by the corresponding n -th order Hermite polynomial, as shown in Eq. 4. Up to first order, the classical moments are recovered: ρ and ρu_α . In general, the second- and higher-order moments will be expressed as $m_{\alpha_1 \dots \alpha_n}^{(n)}$.

First, the concept of the non-equilibrium distribution function is introduced, denoted as $f_i^{(neq)}$. Each distribution function can be expressed as the sum of the equilibrium distribution and a deviation term, referred to as the non-equilibrium distribution [9, 5], as shown in the following expression:

$$\hat{f}_i = f_i^{eq} + \hat{f}_i^{(neq)}, \quad (10)$$

Using assumptions similar to those applied in the moment collision step, the non-equilibrium second-order moment can be explicitly written as:

$$m_{\alpha\beta}^{(neq)} = m_{\alpha\beta}^{(2)} - u_\alpha u_\beta. \quad (11)$$

This definition is necessary because, as proposed by [9], third-order and higher moments can be expressed as functions of the second-order moments. To achieve this, the approach makes use of the non-equilibrium distribution functions. The generalized recurrence relation for the moments is given by:

$$m_{\alpha_1 \dots \alpha_n}^{(neq)} = u_{\alpha_n} m_{\alpha_1 \dots \alpha_{n-1}}^{(neq)} + (u_{\alpha_1} \dots u_{\alpha_{n-2}} m_{\alpha_{n-1} \alpha_n}^{(neq)} + perm(\alpha_n)), \quad (12)$$

where $perm(\alpha_n)$ denotes the permutation of the last index with respect to α_n . When expanding the distribution function to fourth order, it is necessary to account for the fact that a Hermite polynomial cannot be expressed if three identical coordinates are present,

such as in H_{xxx} or H_{yyy} . This limitation arises because linear independence of the system cannot be guaranteed [9].

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 $\sum_i w_i H_{xxy,i} H_{yzz,i} \neq 0$.

2.3 Boundary Conditions

In the case of the three-dimensional lid-driven cavity (LDC), the geometry consists of six walls with prescribed velocities, referred to as Dirichlet boundary conditions. Several approaches have been proposed to handle such boundaries. Among them, the regularized boundary condition (RBC) method provides a robust way to account for the physical constraints of the flow. This approach was first introduced by [5], who applied the regularization process to determine the missing directions at the boundary by formulating a system of equations. Solving this system yields the boundary moments, which in turn allow the populations at the boundary nodes to be defined. The proposed expressions for this system of equations are:

$$\rho_I = \sum_{i \in I_s} f_i = (1 - \omega) \sum_{i \in O_s} \hat{f}_i + \omega \sum f_i^{eq}, \quad (13)$$

$$\rho_I m_{\alpha\beta,I}^{(2)} = \sum_{i \in I_s} f_i H_{\alpha\beta,i}^{(2)} = \sum_{i \in I_s} \hat{f}_i H_{\alpha\beta,i}^{(2)}, \quad (14)$$

where I denotes the incoming moments, I_s the incoming populations, and O_s the outgoing distributions at each boundary node. Further details on the definition of incoming and outgoing directions can be found in [5].

An incompressible flow assumption was also employed at the boundaries, which simplifies the system by allowing the second-order moments to be expressed solely in terms of the velocity components, such as $m_{xx}^{(2)} = u_x^2$ and $m_{yy}^{(2)} = u_y^2$ [8]. This approximation is based on the definition of the strain rate tensor, specifically on the condition that its trace is equal to zero. The strain rate tensor can be expressed as:

$$S_{\alpha\beta} = \frac{a_s^2}{2\tau} \left(u_\alpha u_\beta - m_{\alpha\beta}^{(2)} \right) = \frac{1}{2} \left(\frac{\partial u_\alpha}{\partial \beta} + \frac{\partial u_\beta}{\partial \alpha} \right), \quad (15)$$

Since we are considering the trace of the strain rate tensor, it follows that $\alpha = \beta$. This condition implies that Eq. (15) can be rewritten as:

$$S_{\alpha\alpha} = \frac{a_s^2}{2\tau} \left(u_\alpha^2 - m_{\alpha\alpha}^{(2)} \right) = \frac{\partial u_\alpha}{\partial \alpha}, \quad (16)$$

By the definition of an incompressible fluid, $\nabla \mathbf{u} = 0$. When a lattice node lies on a Dirichlet flat-wall boundary, the velocity is constant along two coordinates, so the sum of their spatial derivatives must vanish. Under the incompressibility assumption, the third

derivative also becomes zero. This reasoning can be extended to edges and corners, which are intersections of two flat walls, where all three spatial velocity derivatives are zero. Consequently, $S_{\alpha\alpha} = 0$ for all boundary nodes, and thus:

$$m_{\alpha\alpha}^{(2)} = u_{\alpha}^2. \quad (17)$$

The RBC framework enhances the stability of LBM simulations, improving performance and enabling simulations at higher Reynolds numbers. With the IRBC assumption, the implementation becomes even simpler.

3 VALIDATION

The validation procedure considers two Reynolds regimes, $Re = 3,200$ and $Re = 10,000$, evaluated across several grid resolutions and two distinct orders of regularization. The results are compared with the experimental data reported by [11]. All simulations were run for a total nondimensional time of $t^* = 500$, where $t^* = tU_{lid}/L$ defines the turnover time. The turnover time incorporates the characteristic length and velocity, representing the number of simulation steps required for the fluid to traverse the entire domain. This ensures that all results correspond to the same physical time, thereby validating comparisons among different grid resolutions.

First, the grid convergence is analyzed to determine the refinement level required for an accurate physical representation of the phenomena. The main variable used to assess convergence is the nondimensional total kinetic energy per unit mass.

Figure 1 presents the results for both Reynolds regimes and four grid resolutions ($L = \{32, 64, 128, 256\}$), with an additional case of $L = 512$ included for $Re = 10,000$. As the resolution increases, the total kinetic energy decays until it reaches a value that changes only marginally compared with the previous resolution. For both Reynolds cases, the high-order scheme (6th order) appears to converge earlier, as its total kinetic energy decays slightly faster than that of the second-order scheme. In the case of $Re = 10,000$, convergence occurs even more rapidly with the high-order scheme. As the Reynolds number increases, the maximum value of the total kinetic energy decreases. Because higher Reynolds flows are more complex, finer grid resolutions are required. This observation is reinforced when comparing the decay rates at low and high Reynolds numbers.

To further validate convergence, a quantitative analysis is performed by calculating the L_2 error, which measures the variance of the nondimensional total kinetic energy within the statistically stationary state ($300 \leq t^* \leq 500$). This metric provides a suitable parameter to compare convergence across different grid resolutions and to assess the rate of error reduction. The L_2 error is defined as:

$$E_{L2}^* = \sqrt{\frac{(E_{t,b} - E_{t,s})^2}{E_{t,b}^2}}, \quad (18)$$

To further validate convergence, a quantitative analysis is performed by calculating the L_2 error, which measures the variance of the nondimensional total kinetic energy within the statistically stationary state ($300 \leq t^* \leq 500$). This metric provides a suitable

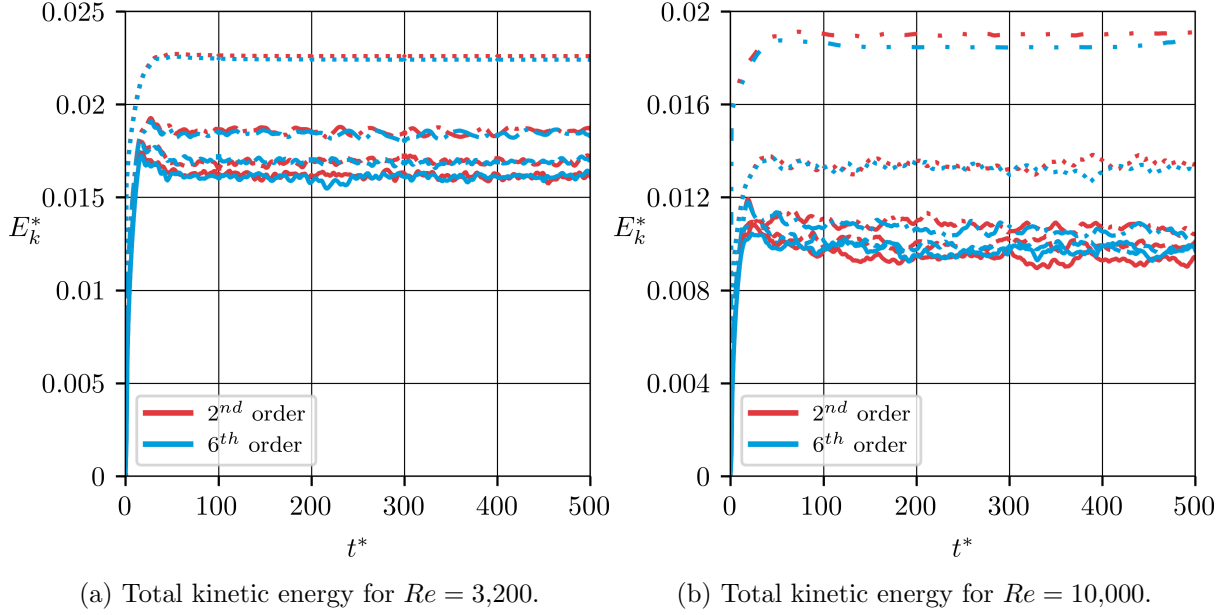


Figure 1: Total kinetic energy for 3,200 and 10,000 Reynolds numbers and different grid resolutions, comparing both 2^{nd} and 6^{th} order of approximation.

parameter to compare convergence across different grid resolutions and to assess the rate of error reduction. The L_2 error is defined as:

Figure 2 shows that for $Re = 3,200$, the convergence rates of the second- and sixth-order frameworks are relatively similar. However, as the Reynolds number increases to 10,000, the difference between the two schemes becomes more pronounced, with the sixth-order formulation converging significantly faster than the second-order scheme. This result suggests that at higher Reynolds numbers, the effects of high-order regularization become more pronounced and contribute more effectively to capturing the flow behavior.

The next step involves the validation of velocity profiles. Since the simulations consider turbulent Reynolds numbers, mean velocity profiles must be computed to account for the fluctuations along the centerline caused by turbulence. The averaging is performed over the interval $300 \leq t^* \leq 500$, ensuring that the fluid has reached a statistically steady state. Under this condition, the resulting velocity profiles can be directly compared with the benchmark data.

Figure 3 shows the velocity profiles along the horizontal and vertical centerlines for $Re = 3,200$. The results indicate that grid refinement improves the accuracy of the outcomes, as the simulated velocities exhibit a closer agreement with the experimental data. Beyond 128 grid points, however, the results change only marginally up to 256 grid points, suggesting that the grid resolution is sufficient for this Reynolds regime. In this case, the difference between the second- and sixth-order schemes is not significant.

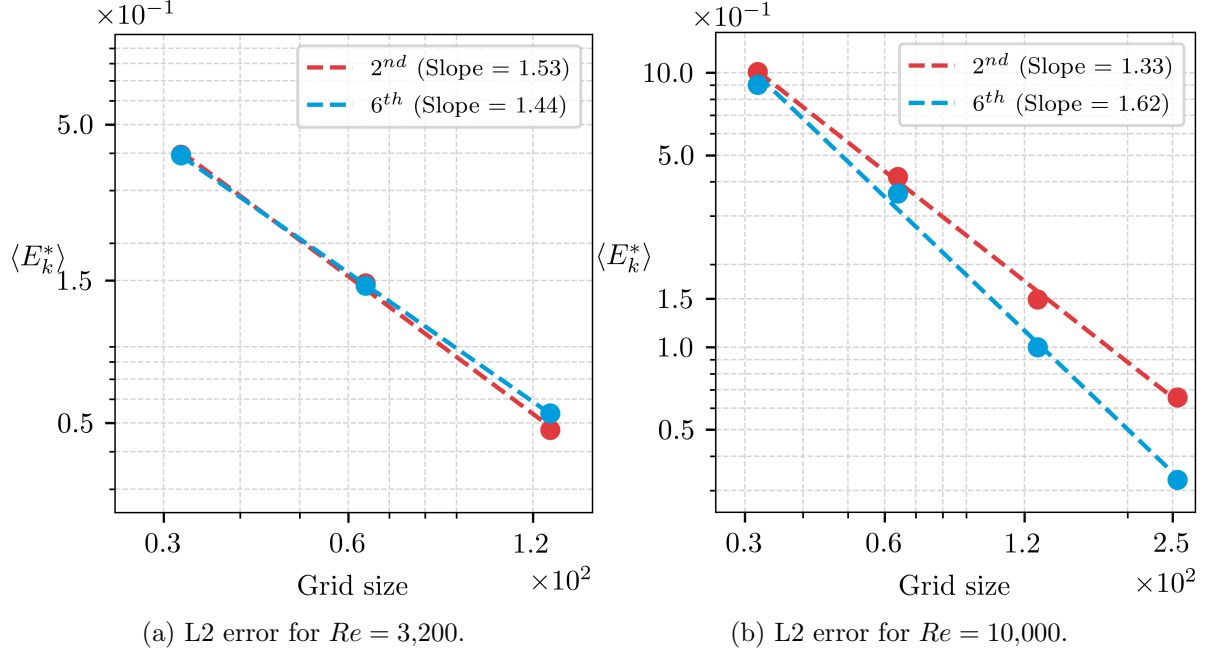


Figure 2: L2 error for 3,200 and 10,000 Reynolds numbers and different grid resolutions, comparing both 2nd and 6th order of approximation.

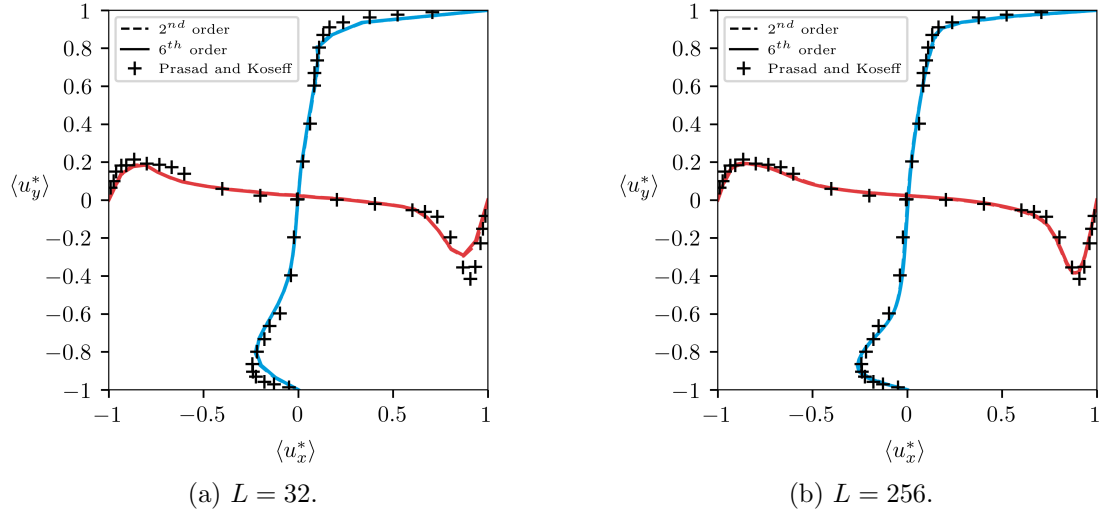
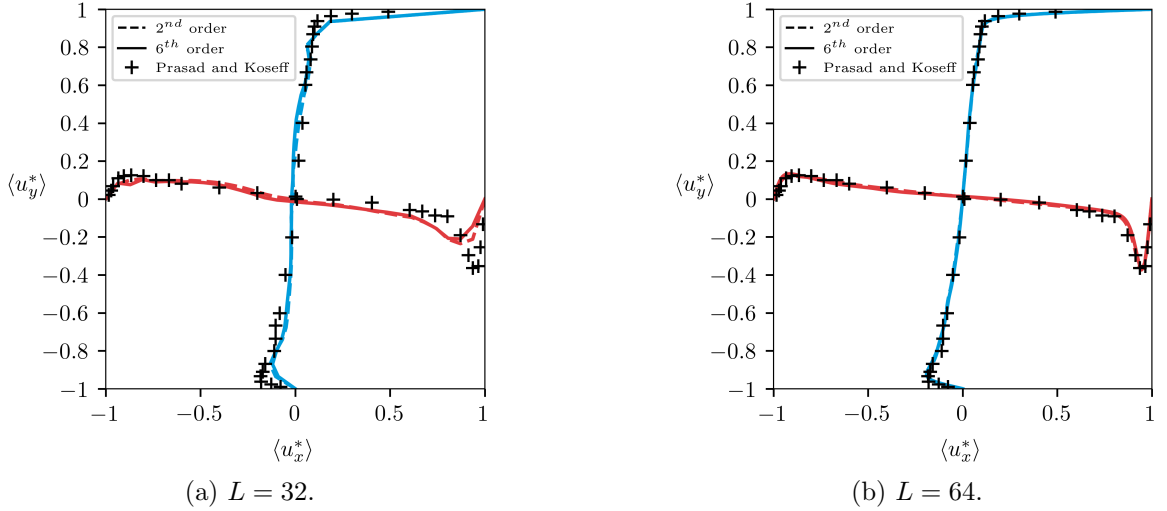


Figure 3: Mean velocity profiles for $Re = 3,200$ across x and y -axis at $z = 0.5$.

For $Re = 10,000$, the influence of grid resolution on the results is much more evident than in the previous case, as shown in Fig. 4. This indicates that high-Reynolds-number

Figure 4: Mean velocity profiles for $Re = 10,000$ across x and y -axis at $z = 0.5$.

flows require finer grid resolutions to achieve an accurate physical representation, which is expected given the greater complexity of turbulent flows. An interesting observation is that the second-order scheme provides a slightly better agreement with the reference data than the sixth-order scheme, which is counterintuitive since a higher-order expansion of the regularized distribution would typically be expected to reduce the error. This outcome suggests that further investigation is required.

In this section, it was shown that both the second- and sixth-order schemes exhibit good agreement with the experimental results of [11], demonstrating that both regularization approaches successfully capture the physical phenomena. This finding supports the consistency of the framework employed in this work. Furthermore, the results proved to be sensitive to grid resolution, which motivated the use of considerably refined grids in the simulations of higher-Reynolds-number flows to ensure accurate representation of the proposed scenarios.

4 RESULTS

To conduct the analysis, simulations were performed for Reynolds numbers of 10,000, 15,000, 25,000, and 50,000 using a grid size of 256, as the validation results indicated that this resolution provides adequate physical convergence. High Reynolds numbers were selected because the flow becomes more complex and increasingly sensitive to numerical instabilities, where the effect of the high-order scheme could be significant. The flow study compares mean velocity profiles and root mean square (RMS) values between the second- and sixth-order frameworks, evaluating whether the change in the regularization order produces any noticeable differences.

As shown in Fig. 5a, both frameworks reproduce the expected physical behavior. On

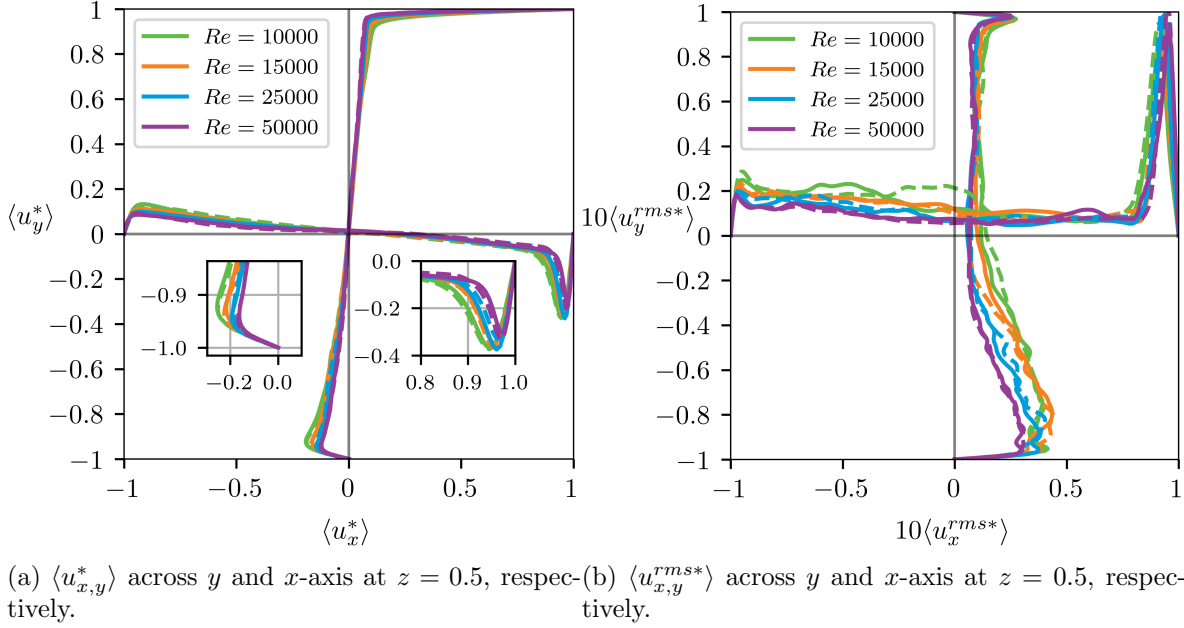


Figure 5: Mean and RMS velocity profiles for several Reynolds regimes across x and y -axis at $z = 0.5$.

the bottom wall, the minimum velocity peak in the x -direction decreases consistently with the observations of [5]. This occurs because the increase in turbulence disrupts the large-scale circulation. For the high-order regularization, a small deviation is observed; however, due to the absence of experimental data for high Reynolds numbers, it is not possible to determine which scheme is closer to reality. Further validation strategies will therefore be required in future work.

Figure 5b highlights the point at which the differences between the two frameworks are most pronounced. From a physical perspective, as the Reynolds number increases, momentum is redistributed from the bottom section of the cavity toward the center. At the same time, the location of the strongest velocity fluctuations shifts from the near-wall region toward the cavity core, reflecting the increasing complexity of turbulent structures. Furthermore, the results reveal noticeable discrepancies between the frameworks. In the absence of experimental data for these cases, a more detailed theoretical investigation will be required to clarify these differences.

The simulations were executed on a GeForce RTX 4090 using the CUDA programming language. Performance was measured in mega lattice updates per second (MLUPS). The implementation is based on a moment-representation LBM framework. While a sixth-order scheme would typically require storing 27 moments instead of the 10 used in a second-order simulation, the recursive moment relation proposed by [9] Eq. (12) allows the same 10 moments to be used to represent the sixth-order formulation. This efficiency is reflected in the MLUPS results: for a sixth-order simulation on a 512^3 grid at $Re = 10,000$, a performance of 4,792 MLUPS was achieved, compared to 4,606 MLUPS for the second-

order scheme. These results demonstrate that the computational performance remains essentially unchanged when employing high-order regularization.

5 CONCLUSION

This study compared second-order and high-order regularization within the LBM framework to evaluate potential improvements in accuracy as well as possible deviations. The validation and results covered a wide range of Reynolds numbers, from $Re = 3,200$ to $Re = 50,000$.

The L_2 error analysis further highlights this behavior. For $Re = 3,200$, the error slopes of the second and sixth-order schemes are slightly similar, indicating comparable convergence rates. However, at $Re = 10,000$, the slope becomes steeper for the sixth-order formulation, showing that the error decreases more rapidly with grid refinement. This behavior suggests that the high-order regularization improves convergence efficiency, particularly in flows with greater turbulence complexity. From a theoretical perspective, this is consistent with the expectation that high-order schemes retain more terms in the Hermite expansion, thereby reducing numerical error relative to the second-order formulation.

When examining the description of the physical phenomena, both the second and sixth-order frameworks demonstrated good agreement with the available experimental results. Moreover, at higher Reynolds numbers, where experimental data are scarce, the simulations followed the expected theoretical predictions, indicating that both regularization approaches are capable of accurately capturing the underlying flow dynamics.

As the Reynolds number increases, the differences between the second and sixth-order frameworks become more evident. However, due to the lack of experimental data at high Reynolds numbers, it is not possible to determine which framework is closer to the true physical behavior, and the results remain inconclusive. Future work should therefore focus on developing theoretical approaches to validate these results and to better understand the deviations observed between the two schemes, particularly in highly turbulent regimes.

Finally, regarding computational efficiency, the implementation based on the [9] recursive-moment approach shows no performance degradation relative to the second-order formulation. This observation suggests that future work should benchmark an alternative implementation in which all higher-order moments are computed and stored explicitly, without recursive reconstruction, to compare numerical accuracy and to quantify any trade-offs between error reduction and runtime/memory usage. Such an ablation study would also help isolate whether the deviations observed at high Reynolds numbers arise from recurrence/truncation effects rather than grid resolution or boundary conditions.

REFERENCES

- [1] Hudong Chen, Satheesh Kandasamy, Steven Orszag, Rick Shock, Sauro Succi, and Victor Yakhot. Extended boltzmann kinetic equation for turbulent flows. *Science*, 301(5633):633–636, 2003.

- [2] Christophe Coreixas and Jonas Latt. Gpu-based compressible lattice boltzmann simulations on non-uniform grids using standard c++ parallelism: From best practices to aerodynamics, aeroacoustics and supersonic flow simulations. *arXiv preprint arXiv:2504.04465*, 2025.
- [3] Timm Krüger, Halim Kusumaatmaja, Alexandr Kuzmin, Orest Shardt, Goncalo Silva, and Erlend Magnus Vigen. The lattice boltzmann method. *Springer International Publishing*, 10(978-3):4–15, 2017.
- [4] Jonas Latt and Bastien Chopard. Lattice boltzmann method with regularized pre-collision distribution functions. *Mathematics and Computers in Simulation*, 72(2-6):165–168, 2006.
- [5] LA Hegele Jr, A Scagliarini, M Sbragaglia, KK Mattila, PC Philippi, DF Puleri, J Gounley, and A Randles. High-reynolds-number turbulent cavity flow using the lattice boltzmann method. *Physical Review E*, 98(4):043302, 2018.
- [6] MA Ferrari, AT Franco, and LA Hegele. Turbulent flow inside a cubic lid-driven cavity using moment representation lattice boltzmann method. *Physics of Fluids*, 36(7), 2024.
- [7] Madhurima Vardhan, John Gounley, Luiz Hegele, Erik W Draeger, and Amanda 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*, pages 1–21, 2019.
- [8] Marco A Ferrari, Waine B de Oliveira Jr, Alan Lugarini, Admilson T Franco, and Luiz A Hegele Jr. A graphic processing unit implementation for the moment representation of the lattice boltzmann method. *International Journal for Numerical Methods in Fluids*, 95(7):1076–1089, 2023.
- [9] Orestis Malaspinas. Increasing stability and accuracy of the lattice boltzmann scheme: recursivity and regularization. *arXiv preprint arXiv:1505.06900*, 2015.
- [10] Christophe Coreixas, Bastien Chopard, and Jonas Latt. Comprehensive comparison of collision models in the lattice boltzmann framework: Theoretical investigations. *Physical Review E*, 100(3):033305, 2019.
- [11] Ajay K Prasad and Jeffrey R Koseff. Reynolds number and end-wall effects on a lid-driven cavity flow. *Physics of Fluids A: Fluid Dynamics*, 1(2):208–218, 1989.