ICCFD13 MILAN ITALY 2026

ICCFD13-2026-175

Probabilistic Lattice Boltzmann Methods for Statistical Solutions of Turbulent Flows

S. Simonis[*,**]
Corresponding author: ssimonis@ethz.ch, stephan.simonis@kit.edu
[*] Seminar for Applied Mathematics, ETH Zurich, Rämistrasse 101, 8092 Zürich, Switzerland.
[**] Institute for Applied and Numerical Mathematics, Karlsruhe Institute of Technology, Englerstrasse 2, 76131 Karlsruhe, Germany.

Abstract: Global well-posedness for three-dimensional fluid flow equations remains a profound open problem. Recent efforts have shifted toward statistical solutions as a robust and physically meaningful framework for describing turbulence; yet, efficient computational tools to explore these solutions in three dimensions are scarce. To this end, we develop novel probabilistic lattice Boltzmann methods to compute and analyze statistical solutions for three-dimensional incompressible fluid flows. We combine various types of lattice Boltzmann methods with probabilistic approaches, such as Monte Carlo and stochastic Galerkin methods. Through accelerated implementations on heterogeneous high-performance computers, we efficiently target application-relevant flow configurations. We present numerical results demonstrating the computational exploration of statistical solutions and their associated observables. We validate the Wasserstein convergence of these statistical measures in periodic regimes and demonstrate the computational feasibility of extracting statistical solutions for initial-boundary value problems where deterministic simulations are nonunique. Finally, we verify that the same pushforward construction extends to the compressible, inviscid regime as well. Specific parts of the developed methodologies are implemented in the open-source framework OpenLB-UQ to ensure public accessibility and sustainable reusability, enabling applications such as training generative diffusion models.

Keywords: Statistical Solutions, Lattice Boltzmann Methods, Uncertainty Quantification, Turbulence, High-Performance Computing.

1. Introduction and problem statement

Viscous ($\varepsilon > 0$), incompressible Newtonian fluid flows can be described by the incompressible Navier–Stokes equations

$$\begin{cases} \mathrm{div}_{\boldsymbol{x}}\,(\boldsymbol{u}) = 0, \\ \partial_t \boldsymbol{u} + \mathrm{div}_{\boldsymbol{x}}\,(\boldsymbol{u}\otimes\boldsymbol{u}) - \varepsilon\boldsymbol{\Delta}_{\boldsymbol{x}}\boldsymbol{u} + \boldsymbol{\nabla}_{\boldsymbol{x}} p = \boldsymbol{F}, \\ \boldsymbol{u}|_{t=0} = \overline{\boldsymbol{u}}. \end{cases} \tag{1}$$

Under an inviscid flow assumption ($\varepsilon \searrow 0$), (1) reduces to the incompressible Euler equations. In both cases, we refer to (1) below. Here, $\boldsymbol{u}\colon \Omega\times I \to U := \mathbb{R}^d$ is the flow velocity dependent on space $\boldsymbol{x}\in\Omega\subseteq\mathbb{R}^d$ and time $t \in I := (0,T] \subseteq \mathbb{R}$, where $d\in\{2,3\}$, and $p\colon \Omega\times I\to\mathbb{R}$ is the pressure (Lagrange multiplier). The given external forces $\boldsymbol{F}$ and the pressure are rescaled with a constant density. The initial data is defined by $\overline{\boldsymbol{u}}\colon \Omega\to\mathbb{R}^d$, which is weakly divergence-free and square-integrable, i.e., $\overline{\boldsymbol{u}} \in L^2_{\mathrm{div}}(\Omega; U)$. The kinematic viscosity $\varepsilon > 0$ is given. In the case of $\Omega \neq \mathbb{R}^d$, the system (1) is supplied with boundary conditions to form an initial-boundary value problem.

Even for spatial periodicity ($\Omega = \mathbb{R}^d$), the global (in time) well-posedness of classical solutions of the 3D incompressible Navier–Stokes and Euler equations is unknown [1, 2]. Although the existence of Leray–Hopf weak solutions [3, 4] has been complemented with uniqueness in the two-dimensional (2D) case, the uniqueness of weak solutions to the incompressible Navier–Stokes equations in 3D remains an open problem [1]. In contrast, supported by the high sensitivity of the flow field $\boldsymbol{u}$ to its initial conditions $\overline{\boldsymbol{u}}$ in experiments, using statistical observables is more appropriate than relying on single sample data [5], especially in applications with advection-dominated (turbulent) flow regimes ($\varepsilon \searrow 0$). Foiaş and Prodi [6, 7] proposed the concept of statistical solutions: Given that the initial velocity is probabilistic with a particular distribution $\overline{\mu}$, the statistical solutions μ are time-parameterized probability measures on $L^2(\Omega; U)$. Considering the incompressible Navier–Stokes equations, whereas the existence and uniqueness of statistical solutions are proven in 2D [6, 8] and existence is proven in 3D, global uniqueness is unknown in 3D. For the incompressible Euler equations, the existence and uniqueness of dissipative statistical solutions have been proven by Mishra and coworkers with suitable initial data in 2D (globally) [2] and in 3D (locally) [2, 9]. Besides, vanishing viscosity limits of statistical solutions of incompressible Navier–Stokes equations towards the Euler equations have been proven for up to 3D [10], assuming a weakened Kolmogorov's hypothesis on small-scale self-similarity [5].

From a computational viewpoint, the concept of statistical solutions offers a mathematically rigorous framework for classical ensemble averaging in engineering applications. The latter has been found to unlock reduced computational effort for computing averaged statistics in ergodic flows [11]. In addition, in non-ergodic flows, where the classical ergodicity assumption (see, e.g., [12] and references therein) fails, using statistical solutions or ensemble-averaging is the only valid approach capable of recovering the time-dependent stochastic response from fluid flows to date.

Though statistical solutions allow for rigorously mapping various features of turbulence, the concept has been rarely studied in application-relevant fluid flows. Computing statistical solutions to 3D incompressible flows with numerical simulations poses an extreme challenge due to the vast number of floating-point operations required and the increased energy-to-solution. Nevertheless, looking back, the effort is necessary for progress in terms of bringing forth theoretical indications, i.e., finding or ruling out possibly unique statistical solution candidates. Many studies have been conducted by Mishra and coworkers: A multi-level Monte Carlo (MC) finite difference discretization of the 2D periodic incompressible Navier–Stokes equations has been devised by Leonardi et al. [13]. Statistical solutions of the 2D incompressible Navier–Stokes equations with boundaries have been computed with single-level MC finite element methods by Bansal [14]. Statistical solutions to the 2D and 3D incompressible Euler equations with periodic boundaries have been approximated with an MC spectral hyper-viscosity method [2, 15]. Finite volume methods have been combined with single and multi-level MC, as well as quasi-Monte Carlo (QMC) methods for the computation of 3D statistical solutions to the compressible Euler equations in *Alsvinn* [16]. Lattice Boltzmann methods (LBMs) have been combined with MC sampling to computationally validate the inviscid limit of periodic statistical solutions to the Navier–Stokes equations in [17, 18]. Notably, prior to the present work, neither numerical methods nor implementations had been proposed for computing statistical solutions to 3D incompressible Navier–Stokes and Euler equations with non-periodic boundaries.

Despite the rapid progress of massively parallel computing hardware, integrated, efficiently scalable numerical methods that can saturate modern-day HPC machinery when computing statistical solutions for canonical flows in 3D have not yet been established. For this purpose, the LBM offers distinct advantages in terms of stability and parallelizability. LBMs have thus become an established alternative to conventional approximation tools for the Navier–Stokes equations where optimized scalability in HPC is crucial. Using LBM in large eddy simulations (LES) provides significant speedup over traditional methods for industry-scale turbulent flows [19]. Moreover, the mesoscopic derivation of LBMs naturally allows for thermodynamically consistent extensions towards multi-physics modeling [20–22] and simulating compressible flows [23–25]. Thus, its unique combination of robustness and near-perfect scalability on current HPC hardware renders LBM a promising numerical scheme for computing statistical solutions. Although combinations of LBMs with both intrusive and non-intrusive methods have rarely been studied in the past, the results are consistently promising. An intrusive stochastic Galerkin LBM is developed in [26] to approximate stochastic convection–diffusion equations with complex boundary conditions. In [27], a non-intrusive collocation method has been combined with LBMs for fluid flows through porous media, reducing the computational effort by a factor of 100 in 2D compared to single-level MC LBMs. Despite these promising features, to the best of the author's knowledge, neither non-intrusive nor intrusive methods have been combined with LBMs for computing statistical solutions in 3D initial-boundary value problems.

2. Objective

The aim of this paper is to introduce the concept of computing statistical solutions for application-relevant turbulent fluid flows. To this end, we propose probabilistic LBMs that combine UQ methods with advanced and efficient kinetic methods. This paper is structured as follows. In Section 3, we define the methodology for both statistical solutions and the numerical approximation of the pushforward measures based on LBM operators. In Section 4, we present specific PLBMs developed to date and include a selection of results. Finally, Section 5 draws a conclusion and suggests future research directions.

3. Methodology

3.1 Statistical solutions of incompressible fluid flows

Let $D \subset \mathbb{R}^3$ be a bounded spatial domain with Lipschitz boundary ∂D and outward unit normal $\boldsymbol{n}$. We consider the evolution of an incompressible, wall-bounded turbulent flow governed by the incompressible Navier–Stokes equations under uncertainty in the initial data.

3.1.1 The phase space and measure space

We define the phase space of physically admissible fluid states, $\mathcal{X}$, as the space of square-integrable, weakly divergence-free vector fields tangent to the boundary (the Leray–Hopf space)

$$\mathcal{X} := \{\boldsymbol{u} \in L^2(D;\mathbb{R}^3) \,|\, \nabla \cdot \boldsymbol{u} = 0 \text{ in } D,\ \boldsymbol{u} \cdot \boldsymbol{n} = 0 \text{ on } \partial D\}, \tag{2}$$

where the divergence constraint and the boundary condition are understood in the weak (distributional and trace) sense, respectively. We equip $\mathcal{X}$ with the norm $\|\boldsymbol{u}\|_{\mathcal{X}} := \|\boldsymbol{u}\|_{L^2(D;\mathbb{R}^3)}$ inherited from the ambient L^2 space, rendering $(\mathcal{X}, \|\cdot\|_{\mathcal{X}})$ a separable Hilbert space. The tangential boundary condition in the definition of $\mathcal{X}$ corresponds to impermeable walls and is stated for definiteness; for configurations with inflow and outflow boundaries, as in Sections 4.3 to 4.5, the phase space is adapted to incorporate the respective inhomogeneous boundary data, while the pushforward construction below remains unchanged. To describe the statistics of the flow, we consider the space of Borel probability measures on $\mathcal{X}$ with finite second moment,

$$\mathcal{P}_2(\mathcal{X}) := \left\{\mu \in \mathcal{P}(\mathcal{X}) \,\middle|\, \int_{\mathcal{X}} \|\boldsymbol{u}\|_{\mathcal{X}}^2 \,\mathrm{d}\mu(\boldsymbol{u}) < \infty\right\}, \tag{3}$$

which we metrize by the p-Wasserstein distance. For $p \in [1,\infty)$ and $\mu, \eta \in \mathcal{P}_p(\mathcal{X})$, the latter is

$$W_p(\mu,\eta) := \left(\inf_{\gamma\in\Pi(\mu,\eta)} \int_{\mathcal{X}\times\mathcal{X}} \|\boldsymbol{u}-\boldsymbol{v}\|_{\mathcal{X}}^p \,\mathrm{d}\gamma(\boldsymbol{u},\boldsymbol{v})\right)^{1/p}, \tag{4}$$

where $\Pi(\mu,\eta) \subset \mathcal{P}(\mathcal{X}\times\mathcal{X})$ denotes the set of all couplings (transport plans) with marginals μ and η. Throughout this work we employ the case $p = 1$, i.e., the 1-Wasserstein distance W_1, both as the analytical metric and as the basis of the computable diagnostics introduced below. A measure $\mu^\varepsilon \in \mathcal{P}_2(\mathcal{X})$ describes the probability distribution of the fluid state, where $\varepsilon > 0$ denotes the kinematic viscosity.

3.1.2 The statistical solution via push-forward

Let $S_t : \mathcal{X} \to \mathcal{X}$, $t \in [0, T]$, be the solution operator (semigroup) associated with the Navier–Stokes equations, such that $\boldsymbol{u}(t) = S_t(\boldsymbol{u}_0)$ for an initial state $\boldsymbol{u}_0 \in \mathcal{X}$. Due to the turbulent nature of the flow, S_t may exhibit chaotic sensitivity to the initial conditions. Given an initial probability measure $\mu_0^\varepsilon \in \mathcal{P}_2(\mathcal{X})$ encoding the uncertainty in the initial state (and potentially in further parameters), the statistical solution at time $t \in (0, T]$ is defined as the push-forward measure

$$\mu_t^\varepsilon := (S_t)_\# \mu_0^\varepsilon, \tag{5}$$

characterized by the change-of-variables identity

$$\int_{\mathcal{X}} \Phi(\boldsymbol{u}) \, \mathrm{d}\mu_t^\varepsilon(\boldsymbol{u}) = \int_{\mathcal{X}} \Phi\big(S_t(\boldsymbol{u}_0)\big) \, \mathrm{d}\mu_0^\varepsilon(\boldsymbol{u}_0), \qquad \forall\, \Phi \in C_b(\mathcal{X}), \tag{6}$$

where Φ ranges over all bounded continuous observables (functionals) in $C_b(\mathcal{X})$ of the flow.

Following the correlation-measure framework of [2, 10], the measure μ_t^ε is equivalently and uniquely characterized by its hierarchy of k-point correlation marginals $\{\nu_{t,\boldsymbol{x}}^{\varepsilon,k}\}_{k\geq 1}$. For $k \in \mathbb{N}$ and $\boldsymbol{x} = (\boldsymbol{x}_1, \ldots, \boldsymbol{x}_k) \in D^k$, the k-point correlation marginal $\nu_{t,\boldsymbol{x}}^{\varepsilon,k} \in \mathcal{P}((\mathbb{R}^3)^k)$ is the joint law of the velocity values $(\boldsymbol{u}(\boldsymbol{x}_1), \ldots, \boldsymbol{u}(\boldsymbol{x}_k))$ under μ_t^ε, i.e., the push-forward of μ_t^ε under the evaluation map $\boldsymbol{u} \mapsto (\boldsymbol{u}(\boldsymbol{x}_1), \ldots, \boldsymbol{u}(\boldsymbol{x}_k))$. The case $k = 1$ yields the single-point law $\nu_{t,\boldsymbol{x}}^{\varepsilon,1}$, while $k = 2$ encodes two-point correlations and thus the structure functions central to the statistical theory of turbulence [5].

Recent analytical results [10] show that, under a weakened Kolmogorov K41 self-similarity (weak scaling) assumption, the statistical solutions of the Navier–Stokes equations converge in the inviscid limit ($\varepsilon \searrow 0$) to a statistical solution of the incompressible Euler equations in the Wasserstein topology induced by (4). In particular, in periodic 3D configurations the statistical Navier–Stokes solutions exhibit weak-strong convergence toward the (weak-strong) unique statistical Euler solution.

3.2 Probabilistic lattice Boltzmann methods

3.2.1 The lattice Boltzmann approximation

The LBM operates on a discrete kinetic phase space $\mathcal{L} := [L^2(D; \mathbb{R})]^q$, where $q \in \mathbb{N}$ denotes the number of discrete velocities and a state $\boldsymbol{f} = (f_0, \ldots, f_{q-1}) \in \mathcal{L}$ collects the discrete particle distribution functions. For a fixed grid spacing $\triangle x > 0$ and associated time step $\triangle t > 0$, the LBM evolution operator (itself being a composition of collision and streaming steps) is denoted by

$$E_t^{\triangle x} \colon \mathcal{L} \to \mathcal{L}. \tag{7}$$

The collision step depends on the kinematic viscosity ε through the relaxation time τ of the collision operator via the Chapman–Enskog relation $\varepsilon = c_s^2 \left(\tau - \frac{1}{2}\right) \frac{\triangle x^2}{\triangle t}$; we suppress this dependence in the notation of $E_t^{\triangle x}$ and of the induced operators below. To link the kinetic and macroscopic descriptions, we define the moment projection map

$$P \colon \mathcal{L} \to \mathcal{X}, \qquad \boldsymbol{f} \mapsto P(\boldsymbol{f}) = \boldsymbol{u} \tag{8}$$

with

$$\boldsymbol{u}(\boldsymbol{x}) := \frac{\sum_{i=0}^{q-1} \boldsymbol{c}_i \, f_i(\boldsymbol{x})}{\sum_{i=0}^{q-1} f_i(\boldsymbol{x})}, \tag{9}$$

where $\{\boldsymbol{c}_i\}_{i=0}^{q-1}$ are the discrete lattice velocities. Note that (9) conforms to the classical concept in LBM that the zeroth (denominator in (9)) and first order moments (numerator in (9)) recover the density ρ and the momentum density $\rho\boldsymbol{u}$, respectively. Conversely, we define an equilibrium initialization map

$$P^*: \mathcal{X} \to \mathcal{L}, \qquad \boldsymbol{u} \mapsto P^*(\boldsymbol{u}) = \boldsymbol{f}^{\mathrm{eq}} \tag{10}$$

with, for example, the classical second order expansion

$$f_i^{\mathrm{eq}}(\boldsymbol{x}) := \hat{w}_i\, \rho(\boldsymbol{x}) \left(1 + \frac{\boldsymbol{c}_i \cdot \boldsymbol{u}(\boldsymbol{x})}{c_s^2} + \frac{(\boldsymbol{c}_i \cdot \boldsymbol{u}(\boldsymbol{x}))^2}{2c_s^4} - \frac{\boldsymbol{u}(\boldsymbol{x}) \cdot \boldsymbol{u}(\boldsymbol{x})}{2c_s^2}\right), \qquad i = 0, \dots, q-1, \tag{11}$$

where $\{\hat{w}_i\}_{i=0}^{q-1}$ are the lattice weights and c_s is the lattice speed of sound, assigning a discrete equilibrium population $\boldsymbol{f}^{\mathrm{eq}} \in \mathcal{L}$ to a macroscopic state. Note that we use a constant reference density in the initialization (11). By construction, P^* satisfies the consistency relation $P \circ P^* = \mathrm{id}_{\mathcal{X}}$. The approximate solution operator on the physical phase space $\mathcal{X}$ induced by the LBM is then

$$S_t^{\triangle x} := P \circ E_t^{\triangle x} \circ P^* : \mathcal{X} \to \mathcal{X}. \tag{12}$$

Note that we regard the choices (9) and (11) as flexible modules within the PLBM class. In the numerical experiments presented in Section 4, a variety of moment and equilibrium models is used.

3.2.2 The combined probabilistic lattice Boltzmann approximation

PLBMs approximate the statistical solution μ_t^{ε} on $(0, T]$ by propagating samples of the input measure μ_0^{ε} through the deterministic LBM operator (12) and reconstructing a discrete empirical measure $\mu_t^{\varepsilon,\triangle x,M}$. Given a set of M samples $\{\boldsymbol{u}_0^{(i)}\}_{i=1}^{M} \subset \mathcal{X}$ drawn from μ_0^{ε} (MC or QMC) or a set of quadrature nodes with associated weights (SC), the PLBM approximation reads

$$\mu_t^{\varepsilon} \approx \mu_t^{\varepsilon,\triangle x,M} := \sum_{i=1}^{M} \alpha_i\, \delta_{S_t^{\triangle x}(\boldsymbol{u}_0^{(i)})}, \qquad \sum_{i=1}^{M} \alpha_i = 1, \tag{13}$$

where $\delta_{\boldsymbol{v}}$ is the Dirac measure centered at $\boldsymbol{v} \in \mathcal{X}$, the propagated samples are denoted by

$$\boldsymbol{u}_t^{\triangle x}(\omega_i) := S_t^{\triangle x}\big(\boldsymbol{u}_0^{(i)}\big) \in \mathcal{X}, \qquad i = 1, \dots, M, \tag{14}$$

and the weights reduce to $\alpha_i = 1/M$ for the (Q)MC estimators and to the quadrature weights for SC.

The statistical observables visualized in Section 4 are the low-order moments of this empirical measure, evaluated point-wise in space. For a bounded continuous observable $\Phi \in C_b(\mathcal{X})$, its expectation under (13) is the weighted sample average

$$\mathbb{E}_{\mu_t^{\varepsilon,\triangle x,M}}[\Phi] := \int_{\mathcal{X}} \Phi(\boldsymbol{v})\, \mathrm{d}\mu_t^{\varepsilon,\triangle x,M}(\boldsymbol{v}) = \sum_{i=1}^{M} \alpha_i\, \Phi\big(\boldsymbol{u}_t^{\triangle x}(\omega_i)\big). \tag{15}$$

Choosing for Φ the (component-wise) point evaluation $\boldsymbol{v} \mapsto \boldsymbol{v}(\boldsymbol{x})$ at a fixed location $\boldsymbol{x} \in D$ yields the empirical mean velocity field, denoted by the ensemble-average bracket $\langle\cdot\rangle$,

$$\langle \boldsymbol{u} \rangle_t^{\triangle x}(\boldsymbol{x}) := \mathbb{E}_{\mu_t^{\varepsilon,\triangle x,M}}\big[\boldsymbol{v} \mapsto \boldsymbol{v}(\boldsymbol{x})\big] = \sum_{i=1}^{M} \alpha_i\, \boldsymbol{u}_t^{\triangle x}(\omega_i)(\boldsymbol{x}), \tag{16}$$

which approximates the first correlation moment $\int_{\mathcal{X}} \boldsymbol{v}(\boldsymbol{x})\, \mathrm{d}\mu_t^{\varepsilon}(\boldsymbol{v})$ of the statistical solution and corresponds to the barycenter of the single-point marginal $\nu_{t,\boldsymbol{x}}^{\varepsilon,1}$. Analogously, the second centered moment defines the empirical (component-wise) variance field

$$\big(\sigma_t^{\triangle x}(\boldsymbol{x})\big)_j^2 := \sum_{i=1}^{M} \alpha_i \left(\big[\boldsymbol{u}_t^{\triangle x}(\omega_i)(\boldsymbol{x})\big]_j - \big[\langle \boldsymbol{u} \rangle_t^{\triangle x}(\boldsymbol{x})\big]_j \right)^2, \qquad j \in \{1,2,3\}, \tag{17}$$

taken per Cartesian velocity component j, with the associated standard deviation (std) field obtained component-wise as $\big(\sigma_t^{\triangle x}(\boldsymbol{x})\big)_j = \big[(\sigma_t^{\triangle x}(\boldsymbol{x}))_j^2\big]^{1/2}$. Together, the mean field $\langle \boldsymbol{u} \rangle_t^{\triangle x}$ (16) and the standard-deviation field $\sigma_t^{\triangle x}$ (17) are the statistics reported in the numerical experiments; their spatial profiles and iso-volumes constitute the primary statistical signature of the computed solution (13). For the (Q)MC estimators ($\alpha_i = 1/M$), the empirical mean (16) is an unbiased estimator of the true mean field with root-mean-square error of order $O(M^{-1/2})$, whereas the SC estimators attain higher-order accuracy under sufficient regularity of the parameter-to-solution map.

For intrusive PLBMs based on stochastic Galerkin methods (SG), the empirical construction (13) is replaced by a spectral (generalized polynomial chaos) expansion of the random state, such that the approximate statistical solution is obtained as the push-forward $\mu_t^{\varepsilon,\mathrm{SG}} := (S_t^{\mathrm{SG}})_{\#}\mu_0^{\varepsilon}$ under the SG-propagated operator S_t^{SG}, rather than through the sample-based moments (16) and (17).

3.2.3 Strong versus statistical convergence

The weak-strong uniqueness principle states that whenever a classical strong solution $\boldsymbol{u}^{\star}$ exists on $[0,t]$, every admissible statistical solution emanating from $\delta_{\boldsymbol{u}^{\star}(0)}$ must collapse to the Dirac mass $\mu_t = \delta_{\boldsymbol{u}^{\star}(t)}$, i.e., a state of zero variance. In highly turbulent, high-Reynolds-number regimes, such strong solutions are not known to exist globally, and the statistical solution may remain genuinely non-Dirac. To distinguish these two scenarios numerically, without recourse to an analytical or fixed global reference solution, we contrast a strong (sample-wise, trajectory-level) error against a statistical (law-level) error, both formulated either as Cauchy differences between consecutive grid refinements $\triangle x$ and $\triangle x/2$ or as classical experimental orders of convergence against the most spatially refined resolution. Note that we only define the Cauchy version here; for the latter, we refer the reader to [18].

Strong sample-wise error. Let Θ denote the underlying sample space and $\omega_m \in \Theta$, $m = 1,\dots,M$, the drawn realizations, with identical pseudo-random seeding maintained across resolutions so that the error isolates numerical convergence rather than sample variance. We define the relative strong sample-wise Cauchy L^1-error

$$\mathcal{E}_{\mathrm{strong}}(t) := \frac{1}{M} \sum_{m=1}^{M} \frac{\|\boldsymbol{u}_t^{\triangle x}(\omega_m) - \mathcal{R}\big(\boldsymbol{u}_t^{\triangle x/2}(\omega_m)\big)\|_{L^1(D)}}{\|\mathcal{R}\big(\boldsymbol{u}_t^{\triangle x/2}(\omega_m)\big)\|_{L^1(D)}}, \tag{18}$$

where $\boldsymbol{u}_t^{\triangle x}(\omega_m) = S_t^{\triangle x}(\boldsymbol{u}_0^{(m)})$ is the coarse-grid solution as in (14), $\boldsymbol{u}_t^{\triangle x/2}(\omega_m)$ the corresponding solution on the refined grid, and $\mathcal{R}$ a conservative block-averaging restriction operator mapping the fine-grid field onto the coarse-grid geometry to enable a point-wise spatial comparison.

Statistical error via correlation marginals. The natural law-level metric is the 1-Wasserstein distance (4) between the empirical measures $\mu_t^{\triangle x}$ and $\mu_t^{\triangle x/2}$, where we abbreviate $\mu_t^{\triangle x} := \mu_t^{\varepsilon,\triangle x,M}$ from (13) whenever ε and M are fixed. Since this distance lives on the infinite-dimensional space $\mathcal{X}$ and is computationally intractable to calculate directly [2], we exploit the correlation-marginal characterization and measure convergence through the k-point marginals $\nu_{t,\boldsymbol{x}}^{\triangle x,k}$ of $\mu_t^{\triangle x}$. Following [2, Eq.(5.7), Proposition E.1], the spatially averaged Wasserstein distance between k-point marginals lower-bounds the full distance up to a constant $C_k > 0$,

$$\int_{D^k} W_1\big(\nu_{t,\boldsymbol{x}}^{\triangle x,k}, \nu_{t,\boldsymbol{x}}^{\triangle x/2,k}\big)\,\mathrm{d}\boldsymbol{x} \;\le\; C_k\, W_1\big(\mu_t^{\triangle x}, \mu_t^{\triangle x/2}\big), \qquad \text{a.e. } t \in (0,T], \tag{19}$$

so that the convergence of the (computable) left-hand side is implied by, and serves as a verifiable diagnostic for, the convergence of the (intractable) right-hand side.

We therefore introduce the family of k-point Wasserstein diagnostics $W_{1,k}$. In the notation $W_{p,k}$, the first index p is the order of the underlying Wasserstein metric (here $p = 1$, i.e., the ground cost is an L^1 distance), and the second index k is the order of the correlation marginal (the number of spatial points). Concretely, for $k = 1$ we define the spatially averaged one-point distance

$$W_{1,1}(t) := \frac{1}{|D|} \int_D W_1\big(\nu_{t,\boldsymbol{x}}^{\triangle x,1}, \nu_{t,\boldsymbol{x}}^{\triangle x/2,1}\big)\, \mathrm{d}\boldsymbol{x}, \tag{20}$$

and for $k = 2$ the spatially averaged two-point distance

$$W_{1,2}(t) := \frac{1}{|D|^2} \int_{D^2} W_1\big(\nu_{t,(\boldsymbol{x}_1,\boldsymbol{x}_2)}^{\triangle x,2}, \nu_{t,(\boldsymbol{x}_1,\boldsymbol{x}_2)}^{\triangle x/2,2}\big)\, \mathrm{d}\boldsymbol{x}_1\, \mathrm{d}\boldsymbol{x}_2, \tag{21}$$

where the evaluation of the local distance W_1 inside the integrands follows the conventions specified next.

Component-wise and vector-valued evaluation. For a scalar observable, the k-point marginal is a probability measure on $\mathbb{R}^k$ and the local W_1 in (20) and (21) is unambiguous. For the velocity, the k-point marginal is a measure on $(\mathbb{R}^3)^k$, and we additionally introduce the component marginals $\nu_{t,\boldsymbol{x}}^{\triangle x,k,(j)}$, defined as the joint law of $\big(u_j(\boldsymbol{x}_1), \dots, u_j(\boldsymbol{x}_k)\big) \in \mathbb{R}^k$ for each Cartesian component $j \in \{1, 2, 3\}$. The component-wise evaluation replaces the local W_1 by the average $\frac{1}{3}\sum_{j=1}^{3} W_1\big(\nu_{t,\boldsymbol{x}}^{\triangle x,k,(j)}, \nu_{t,\boldsymbol{x}}^{\triangle x/2,k,(j)}\big)$, whereas the vector-valued evaluation computes W_1 directly between the joint laws on $(\mathbb{R}^3)^k$ with Euclidean ground cost. Since coordinate projections are 1-Lipschitz, every component distance, and hence their average, is bounded by the vector-valued distance, which in turn is controlled by the full distance through (19); the component-wise evaluation decouples the transport problem into scalar subproblems and is computationally cheaper, whereas the vector-valued evaluation retains inter-component correlations at the cost of a higher-dimensional optimal transport problem.

Throughout this work, the one-point diagnostic $W_{1,1}$ is always evaluated component-wise, so that no notational distinction is required at $k = 1$. At $k = 2$, both evaluations are computed for the velocity, distinguished as $W_{1,2}$ (component-wise) and $W_{1,2}^{\mathrm{v}}$ (vector-valued); this applies to the randomized Taylor–Green vortex in Section 4.2. For scalar observables, such as the density in the astrophysical jet configuration in Section 4.5, the component marginal coincides with the full marginal, the two evaluations agree, and we write $W_{1,2}$ without a superscript.

Computable empirical estimator. For $k = 1$, the component-wise diagnostic (20) admits an exact closed-form evaluation on the empirical ensembles. Because the local one-point component marginals are one-dimensional empirical distributions of equal size M, the optimal transport problem at each grid point and for each scalar component reduces exactly to the L^1 distance between the order statistics of the two samples. Writing $\mathcal{S}(\cdot)_m$ for the operator that sorts the M ensemble values at a given coordinate in ascending order and returns the m-th order statistic, we approximate (20) by

$$W_{1,1}(t) \approx \frac{1}{N_{\boldsymbol{x}}} \sum_{l=1}^{N_{\boldsymbol{x}}} \frac{1}{3} \sum_{j=1}^{3} \frac{1}{M} \sum_{m=1}^{M} \left| \mathcal{S}\Big(\big[\boldsymbol{u}^{\triangle x}(\Theta, \boldsymbol{x}_l, t)\big]_j\Big)_m - \mathcal{S}\Big(\big[\mathcal{R}(\boldsymbol{u}^{\triangle x/2}(\Theta, \boldsymbol{x}_l, t))\big]_j\Big)_m \right|, \tag{22}$$

where $N_{\boldsymbol{x}}$ is the number of discrete spatial cells of the restricted coarse domain, $\boldsymbol{x}_l$ the corresponding cell coordinate, $\boldsymbol{u}^{\triangle x}(\Theta, \boldsymbol{x}_l, t)$ the list of the M ensemble values at $\boldsymbol{x}_l$, and $[\,\cdot\,]_j$ extracts the j-th Cartesian component. For scalar observables, the component average is void and (22) applies verbatim to the scalar ensemble values. We emphasize that the index m in (22) enumerates order statistics of the sorted ensemble, in contrast to its role as a realization index in (18). At $k = 2$, the local transport problems are genuinely multi-dimensional and no sorting shortcut applies; here we employ the same statistical error approximations as in [18]. Convergence of $W_{1,1}$ (and, analogously, of the higher-order $W_{1,k}$) as $\triangle x \searrow 0$ empirically certifies that the law of the flow converges to an admissible statistical state, even where individual realizations diverge through small-scale vortex dislocations.

3.2.4 Convergence rates and error structure

The convergence of the PLBM approximation to the true statistical solution is quantified in the 1-Wasserstein distance W_1 on $\mathcal{X}$ from (4), with the k-point diagnostics $W_{1,k}$ serving as its computable surrogates via the bound (19). While global uniqueness for the 3D Navier–Stokes equations is an open problem, we define μ_t^{ε} constructively as the limit, provided it exists, of the unique statistical solutions of the LBM-regularized dynamics, where the LBM provides a unique selection principle at finite $\triangle x$.

A triangle inequality for W_1 separates the total error into a modeling and a sampling contribution,

$$W_1\big(\mu_t^{\varepsilon}, \mu_t^{\varepsilon,\triangle x,M}\big) \ \leq \ \underbrace{W_1\big((S_t)_{\#}\mu_0^{\varepsilon}, (S_t^{\triangle x})_{\#}\mu_0^{\varepsilon}\big)}_{\text{modeling error}} + \underbrace{W_1\big((S_t^{\triangle x})_{\#}\mu_0^{\varepsilon}, \mu_t^{\varepsilon,\triangle x,M}\big)}_{\text{sampling error}} . \tag{23}$$

The modeling error depends on the consistency of the LBM with the Navier–Stokes equations as $\triangle x \searrow 0$, and on the stability of the solution operator $S_t^{\triangle x}$. Here, consistency is understood under diffusive time-step scaling, $\triangle t \propto \triangle x^2$, under which the lattice Mach number obeys $Ma \propto \triangle x$, so that the leading-order compressibility error $O(Ma^2) = O(\triangle x^2)$ vanishes at second order along the refinement. Holding the relaxation time τ fixed under this scaling would keep the viscosity ε constant; instead, in the numerical experiments of Section 4.2, the inviscid limit is coupled to the grid refinement by prescribing $\varepsilon \propto \triangle x$, equivalently $Re \propto N \propto \triangle x^{-1}$ and $\tau \searrow \frac{1}{2}$ with $\tau - \frac{1}{2} \propto \triangle x$, where $\tau = \frac{1}{2}$ marks the stability boundary of the scheme. Thus, the limits $\triangle x \searrow 0$ and $\varepsilon \searrow 0$ are taken jointly along a coupled path, consistent with the vanishing-viscosity statistical limit under investigation, and the proximity of τ to the stability boundary motivates the entropy-controlled collision model of Section 4.2.

The sampling error depends on the chosen UQ method and the regularity of the parameter-to-solution map; it scales as $O(M^{-1/2})$ for MC and faster for SC/QMC under sufficient regularity. For PLBMs based on stochastic Galerkin methods (SG) (see Section 4.1), the sampling error is expected to decay spectrally, $O(e^{-cM'})$ with $c > 0$ and M' denoting the number of retained gPC modes. This spectral decay holds provided the parameter-to-solution map is sufficiently smooth, in which case the approximation $\mu_t^{\varepsilon,\mathrm{SG}} := (S_t^{\mathrm{SG}})_{\#}\mu_0^{\varepsilon}$ is expected to inherit the spectral accuracy. By testing Wasserstein convergence along the inviscid limit ($\varepsilon \searrow 0$), we lay the groundwork for the exploration of wall-bounded configurations.

4. Numerical experiments

Below, we summarize the specific PLBMs developed to date and include selected results.

4.1 Stochastic Galerkin lattice Boltzmann method

In [28], a PLBM based on stochastic Galerkin methods using polynomial chaos has been proposed and validated through formal analysis, as well as numerical tests for 2D randomized benchmark flows. The results demonstrate the expected spectral convergence and a significantly reduced computational cost, with a speedup factor of 5.72 compared to MC sampling for the randomized Taylor–Green vortex (RTGV) in 2D. The extension of this method to 3D configurations is planned for future work.

4.2 Entropic Monte Carlo lattice Boltzmann method

In exploratory computations of periodic statistical solutions to the 3D Navier–Stokes equations with OpenLB [17], the expected scaling of the energy spectra and the structure functions in terms of Kolmogorov's K41 theory along the inviscid limit ($\varepsilon \searrow 0$) is observed. The computations are carried out for the 3D RTGV, which is based on a perturbed initial velocity $\mathfrak{u}_0 = \boldsymbol{u}_0 + \mathfrak{s} \in \mathbb{R}^3$, involving 24 i.i.d. random variables in $\mathfrak{s}$. We combine MC sampling with a reduced Karlin–Bösch–Chikatamarla (KBC) collision scheme [0] with entropy-controlled relaxation of higher-order moments for nonlinear stability when approximating the incompressible Navier–Stokes equations with LBM along the inviscid limit ($Re \nearrow \infty$). Single samples and mean fields of the RTGV flow are illustrated in Figure 1 for increasing Reynolds numbers, including a Wasserstein convergence validation. The complete computational campaign required approximately 2.5×10^3 GPUh on HoreKa Green (NVIDIA A100 GPUs), where the finest grid contained approximately 1.34×10^8 cells (excluded in Figure 1). Based on PLBM approximations, we have thus computationally confirmed the Wasserstein convergence of statistical Navier–Stokes solutions toward weak-strong unique statistical Euler solutions in a periodic 3D configuration [18].

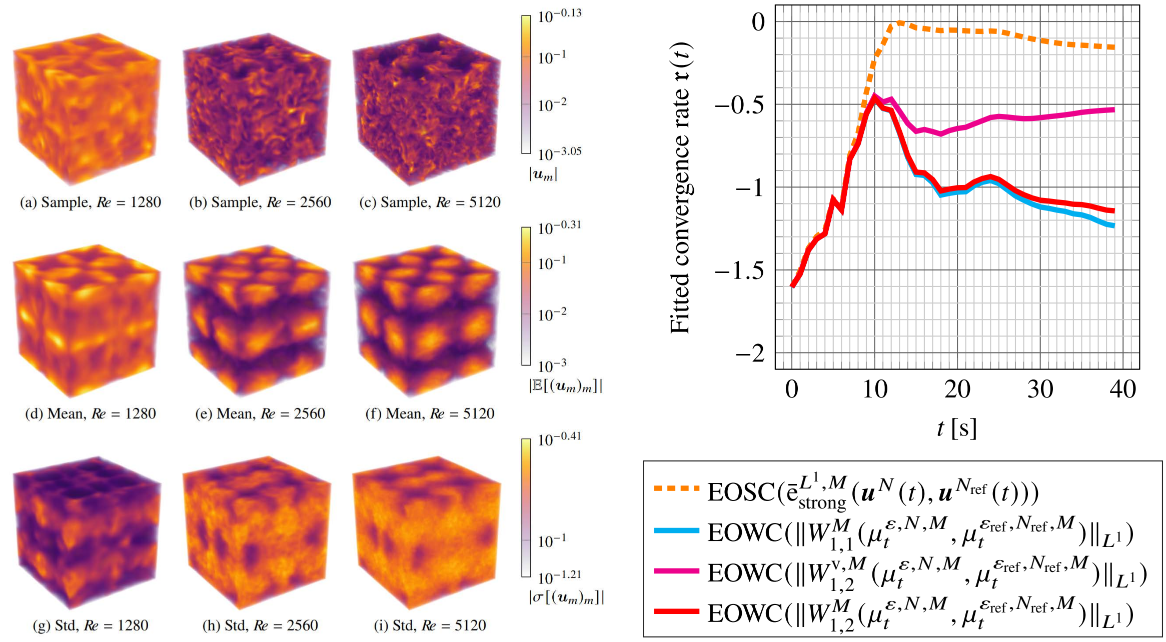

Figure 1: Left panel: Iso-volume velocity fields (colored by magnitude) of 3D RTGV flow simulation with MC KBC LBM at $t \approx 15$s with increasing Reynolds number Re. Single sample computations (top row), mean (16) (middle row), and standard deviation (17) (bottom row) for $Re = 1280, 2560, 5120$ on NVIDIA A100 GPUs. Right panel: Convergence rates $O(N^{\mathfrak{r}(t)})$ calculated at several timesteps via log-log linear fits across grid resolutions N. Experimental order of Wasserstein convergence (EOWC) plotted over time for $W_{1,1}$, $W^{\mathrm{v}}_{1,2}$, and $W_{1,2}$ for $M = 1000$ samples with respect to spatial resolution $N = 256$ at $Re = 10240$. The predicted bound at -0.5 for EOWC is confirmed for $W^{\mathrm{v}}_{1,2}$ (21), whereas the averaged trajectory-wise experimental order of sample convergence (L^1 EOSC, similar to Cauchy error (18)) diverges as expected once turbulence is developed. Figure and notation from [18] (with permission from the authors).

4.3 Recursively regularized stochastic collocation lattice Boltzmann method

In [29], we developed a recursively regularized SC LBM for urban LES with uncertain data assimilation of wind measurements (see Figure 2). The circular lateral boundary condition is perturbed with randomized measurement data for wind speed and direction. The scalable ensemble implementation achieves faster-than-real-time performance (assuming sufficient availability of compute resources, simulating 48 real hours can be reduced to 15 hours of wall-clock time) and provides key statistics such as mean fields and confidence intervals without modifying the deterministic solver. Notably, in this framework, the concept of statistical solutions based on PLBMs is coupled with time-averaging to account for varying wind speeds and directions over the simulated time horizon.

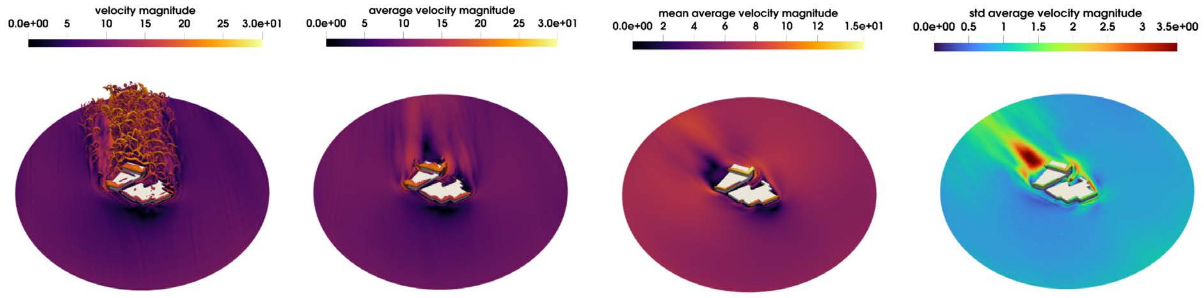

Figure 2: Measurement data-assimilated SC LBM simulation ($M = 121$ samples) of urban wind flow around a real building geometry. Hourly measurements of wind speed and direction are used as continuously updated input data. From left to right: sample of velocity, time-average of single sample, mean of time-averaged velocity, standard deviation of time-averaged velocity. Per sample: $Re = 1.8 \times 10^6$, 10 million grid cells, simulated physical time $T = 48$h in about 15h wall clock time on 4 nodes with 2x Intel Xeon Platinum 8368 each. Figure adapted from [29] (with permission from the authors).

4.4 Incompressible Monte Carlo large eddy lattice Boltzmann method

To conduct probabilistic turbulent round jet simulations, we have developed an incompressible MC LES LBM. The inflow condition is randomized by using an abstracted vortex method [30]. The geometry mimics an injection nozzle exiting into a larger cylindrical container. Figure 3 shows the velocity fields of 32 samples, the mean, the standard deviation, and the velocity component magnitudes plotted over the cross-section centerline. The Reynolds number is defined from the nozzle diameter and is set to $Re = 5000$. Per sample, 10^7 grid cells are used, and the physical simulation time is $T = 150$ [s], requiring approximately 2 hours on a workstation with 38 CPU cores and 2 GPUs. These results demonstrate the capability of PLBMs to stably sample high-Reynolds-number wall-bounded turbulence, paving the way for statistical convergence studies in non-periodic geometries.

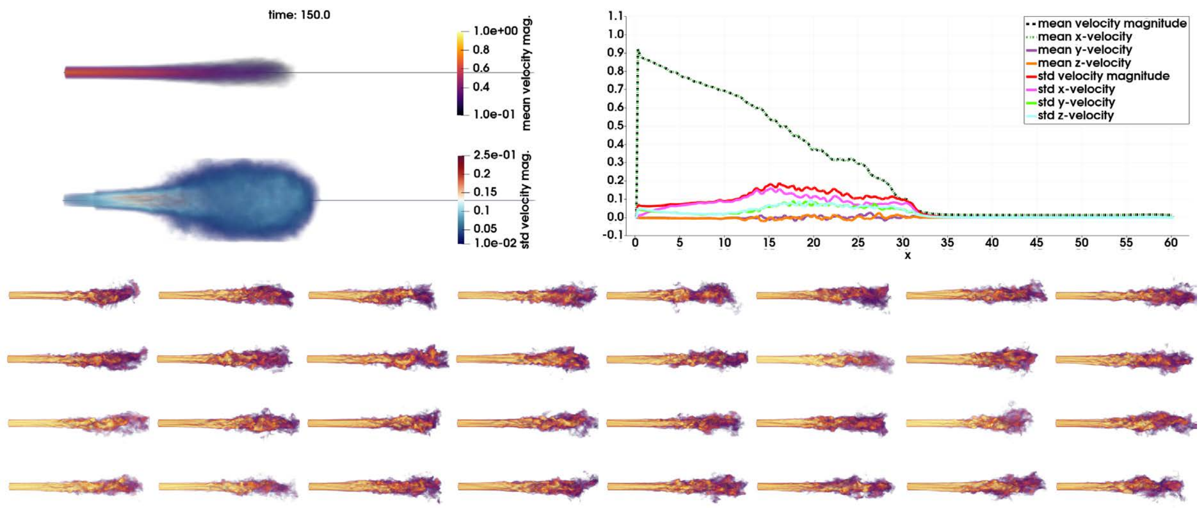

Figure 3: MC LES LBM simulation of a turbulent round jet with random inlet condition ($M = 32$ samples). Top left: Mean (16) and standard deviation (17) of the statistical solution. Top right: Mean and standard deviation velocity probes (magnitude and individual components) along the centerline of the domain. Per sample: $Re = 5000$, 10 million grid cells, simulated physical time $T = 150$s simulated in approximately 2h using 38 CPU cores (Intel Xeon Platinum 8368) and 2x NVIDIA A100 GPUs.

4.5 Compressible Monte Carlo vectorial lattice Boltzmann method

The proposed methodology extends to compressible and inviscid flows as well if a suitable vectorial LBM is used as a deterministic operator in the pushforward. In [31], we computationally explore the limits of weak-strong uniqueness of a Mach 2000 jet by defining the statistical solution as the pushforward of a probability measure through a vectorial LBM (VLBM) operator. The jet inlet at $x = 0$ is randomized with a truncated Karhunen–Loève density perturbation. We compute an ensemble of $M = 1000$ Monte Carlo samples across a sequence of highly refined spatial grids of up to 12.8 million cells and subsequently post-process the empirical measures (see Figure 4). In Figure 5, we contrast the strong sample-wise L^1 Cauchy error divergence with the convergence of the probability measure in the Wasserstein metric with respect to the highest grid resolution. Our results demonstrate that while individual flow realizations physically diverge due to chaotic shear-layer instabilities, the statistical solution converges to an admissible limit measure at a temporally stable rate (cf. Figure 5). Conclusively, we have successfully used PLBMs to provide numerical evidence that the statistical solution of the astrophysical jet problem is non-Dirac and remains stable in the extreme compressible regime [31].

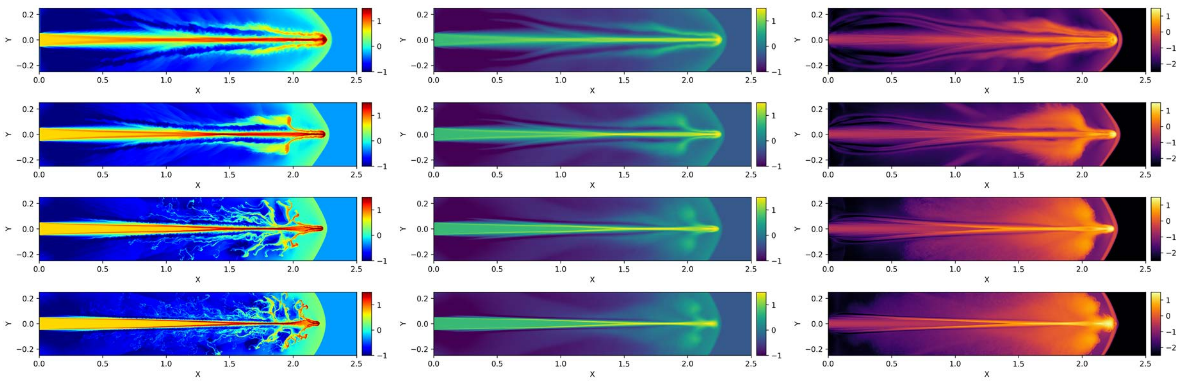

Figure 4: MC VLBM simulation of a Mach 2000 astrophysical jet in two spatial dimensions with Karhunen–Loève-type randomized density inlet condition. Single samples and statistical moments (mean (16) and standard deviation (17)) of the density at $t = 0.0035$ (left to right). The layout displays x-resolutions 500, 1000, 2000, and 4000 (top to bottom) on the y-clipped domain $[0, 2.5] \times [-0.25, 0.25]$. Colorbars are in logarithmic scale. Single samples are computed on NVIDIA H200 GPUs for all resolutions in less than six minutes each. The complete ensemble is distributed over 20 nodes with 50 samples per batch job in the array. Figure from [31] (with permission from the authors).

4.6 Large-scale sampling for learning statistical solutions

Among other data, we include round jet samples (Figure 3) in the order of $M = 10^5$ to train GenCFD [30], a conditional score-based diffusion model. Unlike deterministic MSE-based ensembles that regress to the mean, GenCFD generates spectrally accurate, realistic turbulence samples. Furthermore, [30] provides rigorous theoretical justification for this generative capability, enabled by the efficient production of training data from statistical solutions approximated with PLBMs, among other methods.

4.7 Implementations and code releases

4.7.1 OpenLB-UQ framework

All simulations in Sections 4.2, 4.3, and 4.4 have been implemented in the open-source C++ library OpenLB [20], which is designed for efficient and extensible simulations of complex fluid dynamics on high-performance computers. Specifically, in OpenLB-UQ [32], we leverage the efficiency of OpenLB for large-scale flow sampling with a dedicated and integrated UQ module, where the focus is placed on non-intrusive SC methods based on generalized polynomial chaos and QMC sampling. The OpenLB-UQ framework is extensively validated through convergence tests concerning statistical metrics and sample

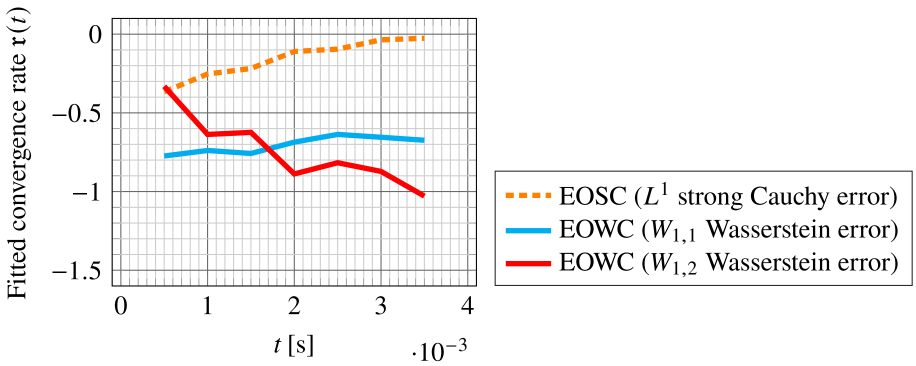

Figure 5: Convergence rates $O(N^{\mathfrak{r}(t)})$ calculated at several timesteps via log-log linear fits across axis grid resolutions N of Mach 2000 astrophysical jet simulations (see Figure 4). As chaotic vortex shedding develops, the magnitude of the strong Cauchy error convergence order (EOSC), averaged over single sample trajectories, decays significantly to $\mathfrak{r} \approx -0.03$. In contrast, the Wasserstein metric with respect to the highest grid resolution $N = 4000$ stabilizes in order magnitude for 1-point marginals ($W_{1,1}$ EOWC), while the order magnitude further increases for 2-point marginals ($W_{1,2}$ EOWC), respectively. Figure adapted from [31] (with permission from the authors).

efficiency using selected benchmark cases [32]. The results in [32] confirm the expected convergence rates and demonstrate promising scalability, robust statistical accuracy, and computational efficiency across thousands of samples and thousands of CPU cores; cf. Figure 6. A structured refactoring of the module towards complete GPU-support, using the existing platform-transparent LBM-kernels in OpenLB, is currently being carried out.

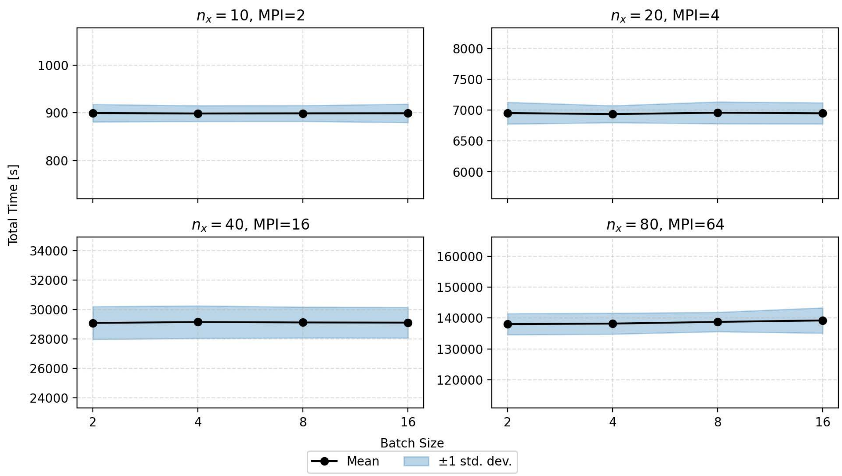

Figure 6: Weak scaling performance of sample-level parallelization for a 2D flow around a cylinder case at several resolutions n_x without postprocessing. The black line shows the mean total time per batch size, with the blue shaded area indicating the standard deviation across repeated runs. Each batch is assigned 100 samples, resulting in overall sample sizes of $M = 200, 400, 800, 1600$ for $2, 4, 8, 16$ batches, respectively. This batch configuration is applied uniformly across all resolutions. The samples are parallelized on domain-level with 2, 4, 16, and 64 MPI processes for the respective resolutions. The batches are then computed in parallel, using up to 1024 CPU cores for the maximum sample size and resolution. Figure from [32] (with permission from the authors).

4.7.2 Python-based implementation

The simulations in Section 4.5 have been performed using Python and CuPy. Conclusively, the entire statistical data generation and processing can be contained in a single short Python file, maximizing portability and user friendliness. An open-source release of the code is planned.

5. Conclusion and future work

We have presented a unified framework for computing statistical solutions to the 3D incompressible Navier–Stokes equations, combining the efficiency of LBM with rigorous probabilistic schemes. We successfully validated Wasserstein convergence in periodic turbulence and extended the methodology to complex, wall-bounded fluid flows. The concept of PLBM further carries over seamlessly to the compressible, inviscid regime. Our results demonstrate that PLBMs are not only computationally feasible on heterogeneous HPC systems but also essential for capturing physical observables in regimes where deterministic solutions are nonunique. In addition, PLBMs provide a mathematically tractable and computationally feasible framework in non-ergodic fluid flow configurations. The high-fidelity datasets generated by this framework provide a crucial foundation for training state-of-the-art generative AI models, such as GenCFD. Future work will focus on the complete GPU porting of the OpenLB-UQ module and the further exploration of inviscid limits in wall-bounded and compressible turbulence to establish new benchmarks for statistical fluid dynamics.

Funding

This work is supported by the PRIME programme of the German Academic Exchange Service (DAAD) with funding from the Federal Ministry of Research, Technology and Space (BMFTR). The presentation of this work at the ICCFD13 2026 was supported by a DAAD Congress Travel Grant.

The author gratefully acknowledges the computing time provided on the high-performance computer HoreKa by the National High-Performance Computing Center at KIT (NHR@KIT). This center is jointly supported by the Federal Ministry of Research, Technology and Space and the Ministry of Science, Research and the Arts of Baden-Württemberg, as part of the National High-Performance Computing (NHR) joint funding program (https://www.nhr-verein.de/en/our-partners). HoreKa is partly funded by the German Research Foundation (DFG).

Acknowledgements

The author thanks S. Mishra, M.J. Krause, and M. Frank for providing the excellent research infrastructure at ETH Zurich and KIT.

Declaration of Generative AI and AI-assisted technologies in the writing process

During the preparation of this work, the author used Anthropic Claude Fable to assist with text formatting. After using this tool, the author reviewed and edited the content as needed and takes full responsibility for the content of the publication.

Data availability statement

Parts of the codes used in this work are openly available within the OpenLB-UQ framework; for further code references, see [32]. The remaining codes and data are available upon reasonable request.

Declaration of competing interest statement

The author declares that he has no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

License statement

References

[1] C. L. Fefferman. “Existence and smoothness of the Navier-Stokes equation”. In: *The Millennium Prize Problems* (2000). URL: https://www.claymath.org/wp-content/uploads/2022/06/navierstokes.pdf.

[2] S. Lanthaler, S. Mishra, and C. Parés-Pulido. “Statistical solutions of the incompressible Euler equations”. In: *Mathematical Models and Methods in Applied Sciences* 31.02 (2021), pp. 223–292. DOI: 10.1142/S0218202521500068.

[3] J. Leray. “Sur le mouvement d’un liquide visqueux emplissant l’espace”. In: *Acta Mathematica* 63.1 (1934), pp. 193–248. DOI: 10.1007/BF02547354.

[4] E. Hopf. “Über die Anfangswertaufgabe für die hydrodynamischen Grundgleichungen”. In: *Mathematische Nachrichten* 4.1-6 (1950), pp. 213–231. DOI: 10.1002/mana.3210040121.

[5] A. Kolmogorov. “The Local Structure of Turbulence in Incompressible Viscous Fluid for Very Large Reynolds Numbers”. In: *Proceedings of the USSR Academy of Sciences* 30 (1941), pp. 301–305.

[6] C. Foias and G. Prodi. “Sur les solutions statistiques des équations de Navier-Stokes”. In: *Annali di Matematica Pura ed Applicata* 111.1 (1976), pp. 307–330. DOI: 10.1007/BF02411822.

[7] C. Foias, O. Manley, R. Rosa, and R. Temam. *Navier-Stokes equations and turbulence*. Vol. 83. Cambridge University Press, 2001.

[8] D. Chae. “The vanishing viscosity limit of statistical solutions of the Navier-Stokes equations. I. 2-D periodic case”. In: *Journal of Mathematical Analysis and Applications* 155.2 (1991), pp. 437–459. DOI: 10.1016/0022-247X(91)90012-O.

[9] K. O. Lye. “Computation of statistical solutions of hyperbolic systems of conservation laws”. Doctoral Thesis. Zurich: ETH Zurich, 2020. DOI: 10.3929/ethz-b-000432014.

[10] U. S. Fjordholm, S. Mishra, and F. Weber. “On the Vanishing Viscosity Limit of Statistical Solutions of the Incompressible Navier–Stokes Equations”. In: *SIAM Journal on Mathematical Analysis* 56.4 (2024), pp. 5099–5143. DOI: 10.1137/23M1566261.

[11] V. Makarashvili, E. Merzari, A. Obabko, A. Siegel, and P. Fischer. “A performance analysis of ensemble averaging for high fidelity turbulence simulations at the strong scaling limit”. In: *Computer Physics Communications* 219 (2017), pp. 236–245. DOI: 10.1016/j.cpc.2017.05.023.

[12] B. Galanti and A. Tsinober. “Is turbulence ergodic?” In: *Physics Letters A* 330.3 (2004), pp. 173–180. DOI: 10.1016/j.physleta.2004.07.009.

[13] F. Leonardi, S. Mishra, and C. Schwab. “Numerical approximation of statistical solutions of planar, incompressible flows”. In: *Mathematical Models and Methods in Applied Sciences* 26.13 (2016), pp. 2471–2523. DOI: 10.1142/S0218202516500597.

[14] P. Bansal. “Numerical approximation of statistical solutions of the incompressible Navier-Stokes Equations”. In: *arXiv* preprint (2021). DOI: 10.48550/ARXIV.2107.06073.

[15] T. Rohner and S. Mishra. “Efficient Computation of Large-Scale Statistical Solutions to Incompressible Fluid Flows”. In: *Proceedings of the Platform for Advanced Scientific Computing Conference*. Pasc ’24. Zurich, Switzerland: Association for Computing Machinery, 2024. DOI: 10.1145/3659914.3659922.

[16] K. O. Lye. “Alsvinn: A Fast multi-GPGPU finite volume solver with a strong emphasis on reproducibility”. In: *arXiv* preprint (2019). DOI: 10.48550/ARXIV.1912.07645.

[17] S. Simonis and S. Mishra. “Computing statistical Navier–Stokes solutions”. In: *Hyperbolic Balance Laws: Interplay between Scales and Randomness*. Ed. by R. Abgrall, M. Garavello, M. Lukáčová-Medvid’ová, and K. Trivisa. Oberwolfach Reports 21, no. 1. EMS Press, 2024, pp. 634–637. DOI: 10.4171/OWR/2024/10.

[18] J. L. Grafen, T. Rohner, and S. Simonis. “Computing weak-strong uniqueness in three dimensions: The statistical Euler limit of Navier–Stokes”. In: *To appear* (2026).

[19] R. Löhner. “Towards overcoming the LES crisis”. In: *International Journal of Computational Fluid Dynamics* 33.3 (2019), pp. 87–97. DOI: 10.1080/10618562.2019.1612052.

[20] M. J. Krause et al. “OpenLB–Open source lattice Boltzmann code”. In: *Computers & Mathematics with Applications* 81 (2021), pp. 258–288. DOI: 10.1016/j.camwa.2020.04.033.

[21] S. A. Hosseini and I. V. Karlin. “Lattice Boltzmann for non-ideal fluids: Fundamentals and Practice”. In: *Physics Reports* 1030 (2023), pp. 1–137. DOI: 10.1016/j.physrep.2023.07.003.

[22] F. Bukreev, A. Kummerländer, and M. J. Krause. “Lattice Boltzmann Methods for Compressible (Magneto)hydrodynamics”. In: *arXiv* preprint (2026). DOI: 10.48550/arXiv.2606.00641.

[23] C. Coreixas and J. Latt. “Compressible lattice Boltzmann methods with adaptive velocity stencils: An interpolation-free formulation”. In: *Physics of Fluids* 32.11 (2020), p. 116102. DOI: 10.1063/5.0027986.

[24] N. G. Kallikounis, B. Dorschner, and I. V. Karlin. “Particles on demand for flows with strong discontinuities”. In: *Physical Review E* 106 (1 2022), p. 015301. DOI: 10.1103/PhysRevE.106.015301.

[25] G. Wissocq, Y. Liu, and R. Abgrall. “A Positive- And Bound-Preserving Vectorial Lattice Boltzmann Method in Two Dimensions”. In: *SIAM Journal on Scientific Computing* 47.6 (2025), A3276–A3302. DOI: 10.1137/24M1712412.

[26] W. Zhao, J. Huang, and W.-A. Yong. “Lattice Boltzmann Method for Stochastic Convection-Diffusion Equations”. In: *SIAM/ASA Journal on Uncertainty Quantification* 9.2 (2021), pp. 536–563. DOI: 10.1137/19M1270665.

[27] L. Zhao and H. Li. “Stochastic modeling of the permeability of randomly generated porous media via the lattice Boltzmann method and probabilistic collocation method”. In: *Transport in Porous Media* 128.2 (2019), pp. 613–631. DOI: 10.1007/s11242-019-01261-7.

[28] M. Zhong, T. Xiao, M. J. Krause, M. Frank, and S. Simonis. “A stochastic Galerkin lattice Boltzmann method for incompressible fluid flows with uncertainties”. In: *Journal of Computational Physics* 517 (2024), p. 113344. DOI: 10.1016/j.jcp.2024.113344.

[0] S. Simonis, B. Dorschner, I. Karlin, and M. J. Krause. “Spectral Effects of Entropic Multi-Relaxation for Lattice Boltzmann Methods”. In: *SSRN Preprint* (2024). DOI: 10.2139/ssrn.4972419.

[29] M. Zhong, D. Teutscher, A. Kummerländer, M. J. Krause, M. Frank, and S. Simonis. “Uncertain Data Assimilation for Urban Wind Flow Simulations with OpenLB-UQ”. In: *arXiv preprint* (2025). DOI: 10.48550/arXiv.2508.18202.

[30] R. Molinaro et al. “Generative AI for fast and accurate statistical computation of fluids”. In: *arXiv preprint* (2025). DOI: 10.48550/arXiv.2409.18359.

[31] S. Simonis and G. Wissocq. “Computing statistical solutions of a Mach 2000 astrophysical jet”. In: *arXiv preprint* (2026). DOI: 10.48550/arXiv.2605.25282.

[32] M. Zhong, A. Kummerländer, S. Ito, M. J. Krause, M. Frank, and S. Simonis. “OpenLB-UQ: An Uncertainty Quantification Framework for Incompressible Fluid Flow Simulations”. In: *arXiv preprint* (2025). DOI: 10.48550/arXiv.2508.13867.