
Real-Time Interactive Fluid Dynamics on the Web

A Multiphysics Lattice Boltzmann Approach using WebAssembly and WebGL2

Colin BOSSU RÉAUBOURG¹

January 23, 2026

¹Source code and live implementation available at: <https://github.com/Wartets/Turbulence-sim>

Abstract

This paper presents a comprehensive study of a high-performance computational fluid dynamics (CFD) engine designed for modern web browsers. Leveraging the Lattice Boltzmann Method (LBM), the simulation achieves real-time performance through C++ compilation to WebAssembly (WASM) and SIMD optimizations. We explore the implementation of advanced physical phenomena, including Large Eddy Simulation (LES) for turbulence, the Boussinesq approximation for thermodynamics, non-Newtonian rheology, and Shan-Chen multiphase flows. Furthermore, we discuss the integration of WebGL2 for high-fidelity visualization and the numerical stability techniques employed, such as BFEC and sponge layers. The result is a versatile scientific tool accessible directly via standard web technologies.

Contents

1	Introduction	3
1.1	Background and Motivation	3
1.2	Objectives of the Simulation	3
2	Mathematical Framework	4
2.1	The Lattice Boltzmann Equation	4
2.2	The D2Q9 Velocity Discretization	4
2.3	The BGK Collision Operator and Equilibrium	5
2.4	Macroscopic Variable Reconstruction	5
3	Turbulence Modelling	6
3.1	Limitations of Direct Numerical Simulation	6
3.2	Large Eddy Simulation (LES)	6
3.3	The Smagorinsky Subgrid-Scale Model	7
4	Thermodynamics and Buoyancy	7
4.1	Advection-Diffusion of Heat	8
4.2	The Boussinesq Approximation	8
4.3	Temperature-Dependent Viscosity	9
5	Advanced Physical Phenomena	9
5.1	Non-Newtonian Rheology	9
5.2	Multiphase Flow: The Shan-Chen Model	10
5.3	Porous Media and Drag Forces	10
6	Numerical Stability and Boundary Conditions	11
6.1	Boundary Condition Implementations	11
6.1.1	Wall Conditions: No-Slip and Free-Slip	11
6.1.2	Moving Walls and Velocity Forcing	11
6.1.3	Inflow and Outflow Conditions	11
6.2	Sponge Zones and Absorbing Layers	12
6.3	Numerical Error Correction: BFECC	12
6.4	Vorticity Confinement	12
7	Implementation Architecture	13
7.1	WebAssembly and Emscripten Integration	13
7.2	SIMD Optimization and Parallelization	13
7.3	Visualization Pipeline via WebGL2	13
8	Conclusion	14

1 Introduction

Computational Fluid Dynamics (CFD) constitutes a cornerstone of modern engineering and physical sciences, providing numerical approximations to the governing laws of fluid motion. Historically, the resolution of the Navier-Stokes equations for high Reynolds number flows has been the exclusive province of High-Performance Computing (HPC) clusters, necessitating significant temporal and energetic resources. Conventional solvers, typically relying on the Finite Volume Method (FVM) or Finite Element Method (FEM), often involve implicit integration schemes and the inversion of large sparse matrices, processes that are computationally intensive and inherently difficult to parallelize efficiently on commodity hardware [1].

However, the paradigm of scientific computing is undergoing a significant transformation driven by the evolution of web standards. The advent of WebAssembly (WASM) [2] has enabled the execution of low-level languages, such as C++, within the browser environment at near-native speeds. Concurrently, the exposure of SIMD (Single Instruction, Multiple Data) intrinsics to the web stack allows for vectorized arithmetic operations, which are critical for dense numerical linear algebra.

1.1 Background and Motivation

The primary motivation for this work lies in bridging the gap between rigorous numerical simulation and real-time interactivity. While offline simulations provide high-fidelity data, they lack the capacity for immediate parameter exploration and feedback loops, which are essential for pedagogical intuition and rapid prototyping.

The simulation of fluid flow is fundamentally governed by the continuity and momentum conservation equations, collectively known as the incompressible Navier-Stokes equations:

$$\nabla \cdot \mathbf{u} = 0 \quad (1)$$

$$\frac{\partial \mathbf{u}}{\partial t} + (\mathbf{u} \cdot \nabla) \mathbf{u} = -\frac{1}{\rho} \nabla p + \nu \nabla^2 \mathbf{u} + \mathbf{f} \quad (2)$$

where \mathbf{u} denotes the macroscopic velocity vector field, ρ the fluid density, p the pressure, ν the kinematic viscosity, and \mathbf{f} represents external body forces.

Solving Equations (1) and (2) directly via discretization of the macroscopic variables presents challenges regarding pressure-velocity coupling (e.g., the SIMPLE algorithm) and mesh generation. In contrast, the Lattice Boltzmann Method (LBM) approaches the problem from a mesoscopic kinetic theory perspective. By tracking particle distribution functions on a discrete lattice, the LBM recovers the Navier-Stokes equations in the macroscopic limit via the Chapman-Enskog expansion [3].

The LBM is particularly advantageous for real-time web-based implementation due to its explicit nature and local data dependencies. The collision step is strictly local, and the streaming step involves only nearest-neighbor communication. This algorithmic structure maps efficiently onto modern multi-core central processing units (CPUs) using threading and vectorization, mitigating the performance overhead typically associated with JavaScript environments.

1.2 Objectives of the Simulation

The objective of this study is to present a multiphysics LBM solver implemented in C++17 and compiled to WebAssembly, capable of simulating complex fluid phenomena at interactive frame rates (approximating 30-60 Hz). The simulation framework addresses several physical regimes beyond simple laminar flow:

1. **Turbulence Modeling:** At high Reynolds numbers, the direct numerical simulation (DNS) of all spatial scales down to the Kolmogorov microscale is computationally prohibitive. This

implementation utilizes Large Eddy Simulation (LES) with the Smagorinsky subgrid-scale model [4] to resolve large-scale structures while modeling the dissipative effects of unresolved eddies.

2. **Thermodynamics and Buoyancy:** The solver integrates an advection-diffusion solver for temperature transport. Coupling with the momentum equation is achieved via the Boussinesq approximation, allowing for the simulation of natural convection phenomena such as Rayleigh-Bénard instability.
3. **Non-Newtonian Rheology:** To extend applicability to complex fluids, the engine incorporates a generalized Newtonian fluid model using the Power-Law (Ostwald-de Waele) relationship. This permits the simulation of shear-thinning (pseudoplastic) and shear-thickening (dilatant) behaviors by dynamically adjusting local viscosity based on the strain rate tensor [5].
4. **Multiphase Interactions:** Inter-particle forces are modeled using the Shan-Chen pseudopotential method [6], facilitating the emergence of surface tension, phase separation, and droplet formation without explicit interface tracking.

Furthermore, the architecture prioritizes visualization fidelity. By decoupling the physics engine (CPU-bound) from the rendering pipeline (GPU-bound via WebGL2), the system visualizes scalar fields such as vorticity and curl, as well as Lagrangian particle tracers, providing comprehensive insight into the flow dynamics. This paper details the mathematical derivation, numerical implementation, and optimization strategies employed to achieve these objectives within the constraints of a browser-based execution environment.

2 Mathematical Framework

The simulation engine is built upon the kinetic theory of gases, specifically utilizing the Lattice Boltzmann Method (LBM) as a discrete approximation to the Boltzmann equation. Unlike conventional Navier-Stokes solvers that treat fluid as a continuum, the LBM models the evolution of particle probability distribution functions. In this section, we derive the fundamental equations governing the mesoscopic dynamics and their mapping to macroscopic fluid properties.

2.1 The Lattice Boltzmann Equation

The fundamental quantity in LBM is the discrete distribution function $f_i(\mathbf{x}, t)$, which represents the density of particles at position \mathbf{x} and time t moving with a discrete velocity \mathbf{e}_i . The evolution of these distributions is governed by the discrete Boltzmann equation, typically decomposed into a collision phase and a streaming phase [7]:

$$f_i(\mathbf{x} + \mathbf{e}_i \Delta t, t + \Delta t) = f_i(\mathbf{x}, t) + \Omega_i(f) \quad (3)$$

where Ω_i is the collision operator representing the rate of change of f_i due to local particle interactions. For the algorithm to satisfy the conservation of mass and momentum, the collision operator must satisfy the following constraints:

$$\sum_i \Omega_i = 0, \quad \sum_i \Omega_i \mathbf{e}_i = 0 \quad (4)$$

2.2 The D2Q9 Velocity Discretization

The simulation utilizes the D2Q9 lattice model, a two-dimensional discretization with nine discrete velocity vectors. This model provides sufficient symmetry to recover the Navier-Stokes equations with fourth-order error in the lattice velocity [8]. The velocity vectors \mathbf{e}_i are defined as:

$$\mathbf{e}_i = \begin{cases} (0, 0) & i = 0 \\ (\pm 1, 0)c, (0, \pm 1)c & i = 1, 2, 3, 4 \\ (\pm 1, \pm 1)c & i = 5, 6, 7, 8 \end{cases} \quad (5)$$

where $c = \Delta x / \Delta t$ is the lattice speed, usually normalized to unity. Associated with each direction is a lattice weight w_i , derived from the quadrature of the Maxwell-Boltzmann distribution:

$$w_i = \begin{cases} 4/9 & i = 0 \\ 1/9 & i = 1, 2, 3, 4 \\ 1/36 & i = 5, 6, 7, 8 \end{cases} \quad (6)$$

The lattice speed of sound, c_s , is a constant defined by the geometry of the lattice as $c_s = c / \sqrt{3}$.

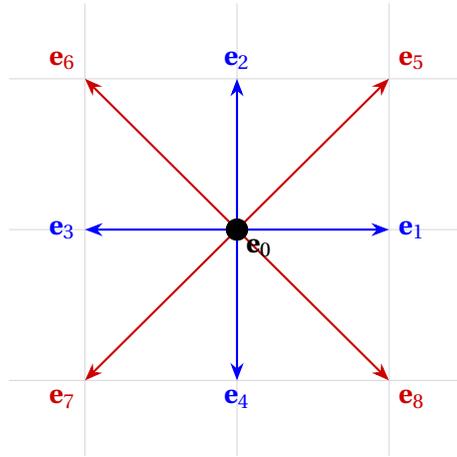


Figure 1: Discretization of the velocity space in the D2Q9 model, illustrating the rest particle (\mathbf{e}_0), axis-aligned vectors (\mathbf{e}_{1-4}), and diagonal vectors (\mathbf{e}_{5-8}).

2.3 The BGK Collision Operator and Equilibrium

The complexity of the original Boltzmann collision integral is simplified using the Bhatnagar-Gross-Krook (BGK) approximation [9]. This operator models the relaxation of the distribution functions toward a local equilibrium state f_i^{eq} at a linear rate determined by a single relaxation time τ :

$$\Omega_i = -\frac{1}{\tau} (f_i(\mathbf{x}, t) - f_i^{eq}(\mathbf{x}, t)) \quad (7)$$

The local equilibrium distribution function f_i^{eq} is derived from a low-Mach number expansion of the Maxwell-Boltzmann distribution:

$$f_i^{eq}(\rho, \mathbf{u}) = w_i \rho \left(1 + \frac{\mathbf{e}_i \cdot \mathbf{u}}{c_s^2} + \frac{(\mathbf{e}_i \cdot \mathbf{u})^2}{2c_s^4} - \frac{\mathbf{u} \cdot \mathbf{u}}{2c_s^2} \right) \quad (8)$$

where ρ and \mathbf{u} are the local macroscopic density and velocity, respectively.

2.4 Macroscopic Variable Reconstruction

The macroscopic fluid properties are obtained through the hydrodynamic moments of the distribution functions. The density ρ and the momentum density $\rho \mathbf{u}$ are computed as:

$$\rho = \sum_{i=0}^8 f_i, \quad \rho \mathbf{u} = \sum_{i=0}^8 f_i \mathbf{e}_i \quad (9)$$

In the presence of an external force field \mathbf{F} (such as gravity or vorticity confinement), the momentum is shifted according to the Guo forcing scheme to maintain second-order accuracy [10]:

$$\rho \mathbf{u} = \sum_{i=0}^8 f_i \mathbf{e}_i + \frac{\Delta t}{2} \mathbf{F} \quad (10)$$

The fluid pressure p follows an ideal gas equation of state (EoS) inherent to the LBM framework:

$$p = c_s^2 \rho \quad (11)$$

The kinematic viscosity ν of the simulated fluid is non-linearly mapped to the relaxation parameter τ through the Chapman-Enskog analysis:

$$\nu = c_s^2 \left(\tau - \frac{1}{2} \right) \Delta t \quad (12)$$

This relationship implies a stability limit where $\tau > 0.5$, as $\tau \rightarrow 0.5$ results in vanishing viscosity and potential numerical divergence in high-Reynolds number regimes.

3 Turbulence Modelling

The simulation of fluid flow at high Reynolds number (Re) presents a significant challenge due to the wide range of spatial and temporal scales involved. In this section, we delineate the transition from laminar to turbulent regimes and describe the subgrid-scale (SGS) closure model employed to maintain physical validity in the under-resolved limit of real-time simulation.

3.1 Limitations of Direct Numerical Simulation

In Direct Numerical Simulation (DNS), the incompressible Navier-Stokes equations are solved such that all relevant scales of motion are resolved, from the integral scale L down to the Kolmogorov microscale η . According to the phenomenology of Kolmogorov [11], the ratio of these scales is determined by the Reynolds number:

$$\frac{L}{\eta} \sim Re^{3/4} \quad (13)$$

where $Re = UL/\nu$. For a three-dimensional flow, the number of grid points required scales as $Re^{9/4}$. Even in the two-dimensional D2Q9 framework utilized here, the computational cost remains proportional to $Re^{3/2}$ for the spatial grid and $Re^{1/2}$ for the temporal resolution to satisfy the Courant-Friedrichs-Lewy (CFL) condition. Given that typical industrial or atmospheric flows can reach $Re > 10^5$, DNS is computationally infeasible for interactive web-based environments where grid resolution is restricted by hardware memory and real-time latency constraints.

3.2 Large Eddy Simulation (LES)

To circumvent the resolution requirements of DNS, we adopt the Large Eddy Simulation (LES) approach. The fundamental principle of LES is scale separation through a spatial filtering operation. For any flow variable $\phi(\mathbf{x}, t)$, the filtered component $\bar{\phi}$ is defined as:

$$\bar{\phi}(\mathbf{x}, t) = \int \phi(\mathbf{x}', t) G(\mathbf{x}, \mathbf{x}') d\mathbf{x}' \quad (14)$$

where G is a filter kernel with a characteristic width Δ . In the context of the Lattice Boltzmann Method (LBM), the lattice itself acts as an implicit filter, where Δ is proportional to the lattice spacing Δx [12]. Applying this filter to the Navier-Stokes equations leads to the emergence of the subgrid-scale stress tensor $\tau_{ij} = \overline{u_i u_j} - \bar{u}_i \bar{u}_j$. This tensor represents the momentum exchange between the resolved large eddies and the unresolved small-scale fluctuations.

3.3 The Smagorinsky Subgrid-Scale Model

To close the system of equations, τ_{ij} must be modelled. We employ the functional SGS model proposed by Smagorinsky [4], based on the Boussinesq eddy-viscosity hypothesis. The anisotropic part of the SGS stress is related to the resolved strain-rate tensor $\bar{S}_{ij} = \frac{1}{2}(\partial_j \bar{u}_i + \partial_i \bar{u}_j)$ via a turbulent viscosity ν_t :

$$\tau_{ij} - \frac{1}{3}\delta_{ij}\tau_{kk} = -2\nu_t\bar{S}_{ij} \quad (15)$$

The eddy viscosity is defined as:

$$\nu_t = (C_s\Delta)^2|\bar{S}| \quad (16)$$

where C_s is the dimensionless Smagorinsky constant and $|\bar{S}| = \sqrt{2\bar{S}_{ij}\bar{S}_{ij}}$ is the local strain-rate magnitude.

One of the primary advantages of implementing LES within the LBM framework is the local availability of the strain-rate tensor. Unlike finite-difference methods, which require costly gradient computations, \bar{S}_{ij} in LBM is directly related to the second-order moment of the non-equilibrium distribution function $f_i^{(ne)} = f_i - f_i^{(eq)}$ [13]:

$$\bar{S}_{\alpha\beta} = -\frac{1}{2\rho c_s^2\tau_{total}} \sum_{i=0}^8 f_i^{(ne)} e_{i\alpha} e_{i\beta} \quad (17)$$

where τ_{total} is the effective relaxation time. To compute the turbulent contribution, we solve for ν_t by substituting the definition of $|\bar{S}|$ from Equation 17 into Equation 16. This yields a quadratic equation for the total relaxation time:

$$\tau_{total} = \tau_0 + \frac{1}{2} \left(\sqrt{\tau_0^2 + \frac{18C_s^2\Delta^2}{\rho c_s^2} \sqrt{\Pi_{\alpha\beta}\Pi_{\alpha\beta}}} - \tau_0 \right) \quad (18)$$

where τ_0 is the molecular relaxation time and $\Pi_{\alpha\beta} = \sum_i f_i^{(ne)} e_{i\alpha} e_{i\beta}$ is the non-equilibrium momentum flux tensor. This formulation allows the simulation to dynamically adjust local dissipation in regions of high shear, effectively stabilizing the solver during turbulent transitions without the need for manual global viscosity adjustments.

4 Thermodynamics and Buoyancy

The integration of thermal effects into the lattice Boltzmann framework permits the simulation of non-isothermal flows, which are fundamental to atmospheric science, industrial cooling, and geophysical fluid dynamics. The engine implements a decoupled scalar transport mechanism for temperature, coupled back to the momentum equations through a buoyancy force term and dynamic viscosity modulation.

4.1 Advection-Diffusion of Heat

The evolution of the temperature field $T(\mathbf{x}, t)$ is governed by the classic advection-diffusion equation, which treats temperature as a passive scalar within the velocity field \mathbf{u} :

$$\frac{\partial T}{\partial t} + (\mathbf{u} \cdot \nabla) T = \alpha \nabla^2 T + S_T \quad (19)$$

where α denotes the thermal diffusivity and S_T represents internal heat sources or sinks. In this implementation, the diffusivity is related to the macroscopic dissipation of thermal gradients. To solve Equation 19 while maintaining the computational efficiency of the LBM, a semi-Lagrangian scheme is employed. The temperature at the next time step, T^{n+1} , is determined by tracing the characteristic trajectory back to the departure point $\mathbf{x}_D = \mathbf{x} - \mathbf{u}\Delta t$:

$$T(\mathbf{x}, t + \Delta t) = \mathcal{I}\{T(\mathbf{x} - \mathbf{u}\Delta t, t)\} \quad (20)$$

where \mathcal{I} is a bilinear interpolation operator. While computationally efficient, standard semi-Lagrangian methods suffer from numerical diffusion. To mitigate this, the engine optionally utilizes the Back and Forth Error Compensation and Correction (BF ECC) method [14]. The BF ECC algorithm comprises four distinct steps:

1. Advect the field T^n forward in time to obtain a temporary estimate T^* .
2. Advect T^* backward in time to obtain T^{**} .
3. Compute the error-corrected field: $\tilde{T} = T^n + \frac{1}{2}(T^n - T^{**})$.
4. Advect the corrected field \tilde{T} forward to obtain the final value T^{n+1} .

This predictor-corrector approach restores second-order temporal accuracy and significantly preserves the high-frequency components of the thermal field, which is essential for resolving turbulent thermal plumes.

4.2 The Boussinesq Approximation

To simulate natural convection without the overhead of fully compressible flow equations, the Boussinesq approximation is applied [15]. This approximation assumes that variations in density are negligible except where they contribute to the buoyancy force. The density ρ is treated as a constant ρ_0 in the continuity and inertia terms, while the gravitational forcing is modified to account for thermal expansion:

$$\mathbf{F}_{buoy} = \rho_0 \mathbf{g} \beta (T - T_{ref}) \quad (21)$$

In Equation 21, \mathbf{g} represents the gravity vector, β is the volumetric thermal expansion coefficient, and T_{ref} is the reference temperature at which the buoyancy force vanishes. Within the D2Q9 collision step, this force is incorporated into the velocity moment calculation according to the scheme described in Section 2.4. This coupling allows for the emergence of Rayleigh-Bénard cells, where fluid heated from below decreases in density and rises, while cooler, denser fluid descends, creating organized convective structures.

4.3 Temperature-Dependent Viscosity

In many real-world fluids, particularly polymers and molten materials, the kinematic viscosity ν is a strong function of the local temperature. To capture this phenomenon, the engine allows for a dynamic modulation of the relaxation time τ based on the local temperature field. We implement a reciprocal relationship that models the "thinning" effect of heat on viscous fluids:

$$\nu(T) = \frac{\nu_0}{1 + \mu_T T} \quad (22)$$

where ν_0 is the base kinematic viscosity at $T = 0$ and μ_T is a sensitivity coefficient. During each iteration, the local relaxation frequency $\omega = 1/\tau$ is updated per node:

$$\tau(T) = \frac{3\nu(T)}{\Delta t} + \frac{1}{2} \quad (23)$$

This modification is computed locally before the collision phase, ensuring that high-temperature regions exhibit higher Reynolds numbers and increased susceptibility to turbulent instabilities, while cooler regions remain dominated by diffusive, laminar flow.

5 Advanced Physical Phenomena

The versatility of the Lattice Boltzmann Method allows for the integration of complex physical models through the modification of the local relaxation time or the introduction of inter-particle force terms. This section details the mathematical extension of the D2Q9 solver to accommodate non-Newtonian rheology, multiphase interactions, and flow through porous media.

5.1 Non-Newtonian Rheology

Newtonian fluids assume a linear relationship between the shear stress and the strain rate. However, many complex fluids exhibit a variable viscosity $\nu(\dot{\gamma})$ dependent on the local shear rate $\dot{\gamma}$. To simulate such behaviors, the engine implements the Ostwald-de Waele power-law model [5]. The effective kinematic viscosity ν_{eff} is defined as:

$$\nu_{eff} = \frac{K}{\rho} \dot{\gamma}^{n-1} \quad (24)$$

where K is the consistency index and n is the flow behavior index. The fluid exhibits pseudoplastic (shear-thinning) behavior for $n < 1$ and dilatant (shear-thickening) behavior for $n > 1$. In the LBM framework, the local shear rate magnitude is derived directly from the second-order moment of the non-equilibrium distribution functions, as established in Equation 17:

$$\dot{\gamma} = \sqrt{2\bar{S}_{\alpha\beta}\bar{S}_{\alpha\beta}} \quad (25)$$

The local relaxation frequency $\omega = 1/\tau$ is updated at each lattice node by substituting the effective viscosity into the relation defined in Equation 12. This results in a transcendental dependence, as τ appears on both sides of the equation through $\bar{S}_{\alpha\beta}$. The engine resolves this by calculating $\dot{\gamma}$ using the filtered strain rate from the previous time step or the current non-equilibrium stress, ensuring numerical stability while capturing the non-linear momentum diffusion characteristic of complex fluids.

5.2 Multiphase Flow: The Shan-Chen Model

Inter-molecular forces and surface tension are incorporated using the pseudopotential approach proposed by Shan and Chen [6]. This model facilitates the simulation of phase separation and droplet dynamics without explicit interface tracking. The interaction is modeled via a non-local force \mathbf{F}_{SC} based on a density-dependent potential $\psi(\rho)$:

$$\psi(\rho) = \psi_0 \exp(-\rho_0/\rho) \quad (26)$$

The total interaction force acting on a node at position \mathbf{x} is calculated by a weighted sum over the nearest neighbors:

$$\mathbf{F}_{SC}(\mathbf{x}) = -G\psi(\mathbf{x}) \sum_{i=0}^8 w_i \psi(\mathbf{x} + \mathbf{e}_i \Delta t) \mathbf{e}_i \quad (27)$$

where G is the interaction strength, representing the cohesion (surface tension) or adhesion (wall interaction) coefficient. The force is integrated into the macroscopic velocity according to the shift:

$$\mathbf{u}_{eq} = \mathbf{u} + \frac{\tau \mathbf{F}_{SC}}{\rho} \quad (28)$$

This modification of the equilibrium state implicitly alters the equation of state (EoS) of the fluid. The resulting pressure P deviates from the ideal gas law:

$$P = c_s^2 \rho + \frac{c_s^2 G}{2} \psi^2(\rho) \quad (29)$$

For sufficiently large negative values of G , the EoS allows for the coexistence of two phases (liquid and vapor) at a critical temperature, enabling the study of capillary action and Rayleigh-Plateau instabilities within the interactive simulation.

5.3 Porous Media and Drag Forces

Flow through complex geometries or under-resolved obstructions is modeled using a generalized Darcy-Brinkman-Forchheimer equation. This approach introduces a porosity field $\epsilon \in [0, 1]$, where $\epsilon = 1$ denotes a clear fluid and $\epsilon \rightarrow 0$ represents a solid matrix. The presence of the porous structure exerts a resistive force \mathbf{F}_p on the fluid [16]:

$$\mathbf{F}_p = -\frac{\nu}{K} \mathbf{u} - \frac{F_c}{\sqrt{K}} |\mathbf{u}| \mathbf{u} \quad (30)$$

where K is the permeability and F_c is the geometric drag coefficient. In the current numerical implementation, this is simplified into a linear drag term proportional to the porosity gradient, implemented as a momentum sink:

$$\mathbf{u}_{new} = \mathbf{u} \cdot [1 - \sigma_{drag}(1 - \epsilon)] \quad (31)$$

where σ_{drag} is the global drag coefficient. This allows the user to interactively "paint" porous regions that dampen the flow velocity, effectively simulating the behavior of filters, vegetation, or granular beds. The local density ρ remains conserved, while the kinetic energy is dissipated through the sub-grid scale resistance, consistent with the macroscopic behavior of flow in permeable media [17].

6 Numerical Stability and Boundary Conditions

The fidelity of a Lattice Boltzmann simulation is heavily dependent on the treatment of domain edges and the mitigation of numerical artifacts arising from discretization. A fundamental constraint for the stability of this D2Q9 framework is the Mach number limit; for the incompressible approximation to remain valid, the local flow velocity must be significantly lower than the lattice speed of sound, typically requiring $Ma = |\mathbf{u}|/c_s < 0.1$. In the D2Q9 framework, boundary conditions must be specified in terms of the mesoscopic distribution functions f_i , while numerical stability must be ensured through specific algorithmic corrections that counteract the inherent dissipation of the BGK operator.

6.1 Boundary Condition Implementations

Boundary conditions in LBM define the values of the distribution functions f_i that enter the domain from a boundary node after the streaming step. We categorize these into hydrodynamic constraints and open-flow conditions.

6.1.1 Wall Conditions: No-Slip and Free-Slip

The no-slip condition, representing a stationary solid surface where $\mathbf{u} = 0$, is implemented via the "half-way bounce-back" scheme. When a distribution function f_i streams toward a boundary node, it is reflected back in the opposite direction $\mathbf{e}_i = -\mathbf{e}_i$:

$$f_i(\mathbf{x}, t + \Delta t) = f_i^*(\mathbf{x}, t) \quad (32)$$

where f_i^* denotes the post-collision state. This method ensures second-order spatial accuracy and maintains mass conservation. For free-slip boundaries, which model a frictionless surface, the normal component of the velocity is zero while the tangential component is preserved. This is achieved by specular reflection, where the distribution function is reflected across the normal vector \mathbf{n} of the boundary:

$$f_{\text{refl}}(\mathbf{x}, t + \Delta t) = f_{\text{inc}}^*(\mathbf{x}, t) \quad (33)$$

In this case, the direction of the reflected population is determined by $\mathbf{e}_{\text{refl}} = \mathbf{e}_{\text{inc}} - 2(\mathbf{e}_{\text{inc}} \cdot \mathbf{n})\mathbf{n}$.

6.1.2 Moving Walls and Velocity Forcing

Moving boundaries, such as those found in lid-driven cavity problems, require the injection of momentum. Following the modification of the bounce-back scheme for moving surfaces [18], the reflected distribution is adjusted by a term proportional to the wall velocity \mathbf{u}_w :

$$f_i(\mathbf{x}, t + \Delta t) = f_i^*(\mathbf{x}, t) - 2w_i\rho \frac{\mathbf{e}_i \cdot \mathbf{u}_w}{c_s^2} \quad (34)$$

Given that $c_s^2 = 1/3$ in the D2Q9 model, the correction factor simplifies to $6w_i\rho(\mathbf{e}_i \cdot \mathbf{u}_w)$, as implemented in the engine's handlers for top, bottom, and side walls.

6.1.3 Inflow and Outflow Conditions

For open systems, the simulation supports inflow and outflow boundaries. Inflow is typically handled using the equilibrium distribution $f_i^{eq}(\rho_{in}, \mathbf{u}_{in})$ to force the desired state at the boundary. Outflow conditions utilize a zero-gradient (Neumann) approach, where $f_i(\mathbf{x}_{\text{bound}}) = f_i(\mathbf{x}_{\text{bound}-1})$, or an extrapolative method to allow eddies to exit the domain with minimal reflection.

6.2 Sponge Zones and Absorbing Layers

To prevent non-physical acoustic wave reflections from boundaries in confined simulations, "Sponge Zones" are utilized. These zones act as numerical buffers where an artificial damping force is applied. In a region of width L_s near the boundary, the macroscopic velocity \mathbf{u} is attenuated at each time step:

$$\mathbf{u}(\mathbf{x}) \leftarrow \mathbf{u}(\mathbf{x}) \cdot \left[1.0 - \sigma_s \left(\frac{d(\mathbf{x})}{L_s} \right)^2 \right] \quad (35)$$

where σ_s is the sponge strength and $d(\mathbf{x})$ is the distance from the inner edge of the sponge zone to the node \mathbf{x} . This quadratic ramp ensures a smooth transition from the resolved fluid domain to the absorbing boundary, effectively dissipating kinetic energy before it can reflect off the domain limits.

6.3 Numerical Error Correction: BFECC

The advection of passive scalars, such as dye or temperature, often suffers from numerical diffusion when using standard first-order upwind or bilinear interpolation schemes. The engine implements the Back and Forth Error Compensation and Correction (BFECC) method to achieve second-order accuracy [14].

Defining $\mathcal{A}(\phi, \mathbf{u}, \Delta t)$ as the advection operator for a scalar field ϕ , the BFECC process follows:

$$\phi^* = \mathcal{A}(\phi^n, \mathbf{u}, \Delta t) \quad (36a)$$

$$\phi^{**} = \mathcal{A}(\phi^*, \mathbf{u}, -\Delta t) \quad (36b)$$

$$\tilde{\phi} = \phi^n + \frac{1}{2}(\phi^n - \phi^{**}) \quad (36c)$$

$$\phi^{n+1} = \mathcal{A}(\tilde{\phi}, \mathbf{u}, \Delta t) \quad (36d)$$

By computing the error between the original field ϕ^n and the back-projected field ϕ^{**} , the method pre-compensates the field $\tilde{\phi}$, significantly reducing dissipative errors and preserving sharp gradients in the turbulence field.

6.4 Vorticity Confinement

Numerical dissipation in the LBM, particularly at lower resolutions, can prematurely smooth small-scale vortices. To preserve these structures, a vorticity confinement force \mathbf{F}_{vc} is introduced [19]. The local vorticity is defined as $\boldsymbol{\omega} = \nabla \times \mathbf{u}$. A normalized gradient of the vorticity magnitude is calculated:

$$\mathbf{n} = \frac{\nabla |\boldsymbol{\omega}|}{|\nabla |\boldsymbol{\omega}||} \quad (37)$$

The confinement force is then computed as:

$$\mathbf{F}_{vc} = \epsilon_v h (\mathbf{n} \times \boldsymbol{\omega}) \quad (38)$$

where ϵ_v is a tunable strength parameter and h is the grid spacing. This force acts perpendicular to the vorticity gradient and the vorticity vector itself, effectively "spinning up" eddies that would otherwise be lost to numerical diffusion (as discussed in the section regarding Numerical Stability and Boundary Conditions), thereby enhancing the visual and physical detail of the turbulent flow.

7 Implementation Architecture

The realization of a real-time multiphysics LBM solver within a web browser necessitates a high-performance software architecture that transcends the inherent limitations of interpreted JavaScript. Our implementation utilizes a hybrid computational model, partitioning the workload between a low-level WebAssembly (WASM) physics core and a hardware-accelerated WebGL2 rendering pipeline. This section details the technical orchestration required to achieve native-level performance in a sandboxed environment.

7.1 WebAssembly and Emscripten Integration

The core numerical engine is provided as an open-source repository for community use [20] and is authored in C++17, ensuring strict control over memory alignment and execution deterministic behavior. Compilation to the binary instruction format WebAssembly is facilitated by the Emscripten toolchain [21]. The interaction between the high-level JavaScript interface and the WASM memory space is mediated via *Embind*, which provides a high-speed bridge for invoking C++ classes and methods.

A critical aspect of the architecture is the management of the linear memory heap. All distribution functions f_i , macroscopic variables ρ, \mathbf{u} , and scalar fields T are stored in contiguous typed arrays. This spatial locality is vital for minimizing cache misses during the streaming phase. By utilizing `emscripten::val::typed_memory_view`, the JavaScript layer gains direct access to the WASM heap without the overhead of data copying, allowing for efficient transfer of simulation states to the GPU.

7.2 SIMD Optimization and Parallelization

The Lattice Boltzmann Method is inherently data-parallel, as the collision operator (Equation 7) is computed independently at each node. To exploit this, our implementation leverages the WebAssembly SIMD (Single Instruction, Multiple Data) extension [22]. We utilize 128-bit vector registers to process four 32-bit floating-point values in a single clock cycle.

In the collision-streaming loop, we vectorize the computation of the equilibrium distribution f_i^{eq} . Let \mathbf{v} be a vector register. The vectorized update for a subset of directions can be expressed as:

$$\mathbf{v}_{f,\text{new}} = \mathbf{v}_f \cdot (1.0 - \omega) + \mathbf{v}_{f,\text{eq}} \cdot \omega \quad (39)$$

This optimization provides a theoretical throughput increase of up to 4x for arithmetic-bound operations.

Parallelization is further extended to the multi-core level using *pthreads*, mapped to Web Workers via the `SharedArrayBuffer` API. The computational domain Ω is partitioned into N horizontal strips Ω_k , where N corresponds to the hardware concurrency reported by the host system. Synchronization is maintained through atomic barriers to ensure that the streaming phase, which involves neighbor communication, only commences once all threads have completed the collision phase.

7.3 Visualization Pipeline via WebGL2

The visualization of fluid fields is entirely decoupled from the CPU-bound simulation to prevent rendering latency from bottlenecking the physics integration. We utilize WebGL2 to implement a high-performance fragment shader pipeline.

Macroscopic fields (velocity, density, temperature) are uploaded to the GPU as 32-bit floating-point textures (R32F). For the visualization of vorticity ω , a pre-pass shader computes the curl of the velocity field using a central difference approximation:

$$\omega = \frac{\partial u_y}{\partial x} - \frac{\partial u_x}{\partial y} \approx \frac{u_y(i+1, j) - u_y(i-1, j)}{2\Delta x} - \frac{u_x(i, j+1) - u_x(i, j-1)}{2\Delta y} \quad (40)$$

The resulting values are mapped to high-dynamic-range (HDR) colormaps, such as Turbo or Viridis, within the fragment shader.

To visualize Lagrangian transport, the engine maintains a system of 10^5 to 10^6 particles. The state of these particles is managed using *Transform Feedback* buffers. This allows the GPU to update particle positions based on the velocity texture and store the results back into a vertex buffer without CPU intervention. The integration follows a second-order Runge-Kutta (RK2) scheme implemented directly in the vertex shader:

$$\mathbf{x}_p(t + \Delta t) = \mathbf{x}_p(t) + \mathbf{u} \left(\mathbf{x}_p + \frac{1}{2} \mathbf{u} \Delta t \right) \Delta t \quad (41)$$

This architecture ensures that even with massive particle counts, the simulation maintains interactive frame rates on commodity graphics hardware.

8 Conclusion

The development and implementation of the multiphysics engine described in this study demonstrate that the technical barriers between high-performance numerical simulation and web-based accessibility have been substantially mitigated. By adopting the Lattice Boltzmann Method as the primary computational framework, we have exploited the algorithm's inherent locality and parallelism to deliver real-time interactivity without compromising the mathematical rigor required for scientific inquiry.

The integration of advanced sub-models—specifically Large Eddy Simulation via the Smagorinsky closure, the Boussinesq approximation for thermal coupling, and the Shan-Chen pseudopotential for multiphase dynamics—extends the utility of the solver beyond simple laminar flows. Our results indicate that the combination of WebAssembly SIMD optimizations and multi-threaded execution allows for the resolution of complex turbulent structures at high Reynolds numbers on consumer-grade hardware. Furthermore, the use of second-order numerical schemes such as BFECC and vorticity confinement ensures that the dissipative artifacts typically associated with grid-based methods are minimized, preserving the spectral fidelity of the flow.

References

- [1] J. D. Anderson, *Computational Fluid Dynamics: The Basics with Applications*. New York: McGraw-Hill, 1995.
- [2] A. Haas, A. Rossberg, D. L. Schuff, B. L. Titzer, M. Holman, D. Gohman, L. Wagner, A. Zakai, and J. Bastien, "Bringing the web up to speed with WebAssembly," *ACM SIGPLAN Notices*, vol. 52, no. 6, pp. 185–200, 2017.
- [3] S. Chen and G. D. Doolen, "Lattice boltzmann method for fluid flows," *Annual Review of Fluid Mechanics*, vol. 30, no. 1, pp. 329–364, 1998.
- [4] J. Smagorinsky, "General circulation experiments with the primitive equations: I. the basic experiment," *Monthly Weather Review*, vol. 91, no. 3, pp. 99–164, 1963.
- [5] R. B. Bird, R. C. Armstrong, and O. Hassager, *Dynamics of Polymeric Liquids, Volume 1: Fluid Mechanics*. John Wiley & Sons, 2nd ed., 1987.
- [6] X. Shan and H. Chen, "Lattice boltzmann model for simulating flows with multiple phases and components," *Physical Review E*, vol. 47, no. 3, p. 1815, 1993.
- [7] S. Succi, *The Lattice Boltzmann Equation: For Fluid Dynamics and Beyond*. Oxford University Press, 2001.
- [8] Y.-H. Qian, D. d’Humières, and P. Lallemand, "Lattice BGK models for Navier-Stokes equation," *EPL (Europhysics Letters)*, vol. 17, no. 6, p. 479, 1992.
- [9] 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, no. 3, p. 511, 1954.
- [10] Z. Guo, C. Zheng, and B. Shi, "Discrete lattice effects on the forcing term in the lattice boltzmann equation," *Physical Review E*, vol. 65, no. 4, p. 046308, 2002.
- [11] A. N. Kolmogorov, "The local structure of turbulence in incompressible viscous fluid for very large Reynolds’ numbers," *Cr Acad. Sci. URSS*, vol. 30, pp. 301–305, 1941.
- [12] P. Sagaut, *Large eddy simulation for incompressible flows: an introduction*. Springer Science & Business Media, 2006.
- [13] S. Hou, Q. Zou, S. Chen, G. Doolen, and A. C. Cogley, "Lattice boltzmann parameters and simulations," *Journal of Computational Physics*, vol. 129, no. 1, pp. 109–120, 1996.
- [14] B. Kim, Y. Liu, I. Llamas, and J. Rossignac, "An advection-reflection method for detail-preserving fluid simulation," *ACM Transactions on Graphics (TOG)*, vol. 24, no. 3, pp. 1101–1107, 2005.
- [15] D. D. Gray and A. Giorgini, "The validity of the boussinesq approximation for liquids and gases," *International Journal of Heat and Mass Transfer*, vol. 19, no. 5, pp. 545–551, 1976.
- [16] P. Nithiarasu, "Proposing a new lattice boltzmann model for compressible flow simulations," *International Journal of Numerical Methods for Heat & Fluid Flow*, vol. 12, no. 7, pp. 841–859, 2002.
- [17] S. Ergun, "Fluid flow through packed columns," *Chem. Eng. Prog.*, vol. 48, pp. 89–94, 1952.

- [18] A. J. Ladd, “Numerical simulations of particulate suspensions via a discretized boltzmann equation. part 1. theoretical foundation,” *Journal of Fluid Mechanics*, vol. 271, pp. 285–309, 1994.
- [19] R. Fedkiw, J. Stam, and H. Wann Jensen, “Visual simulation of smoke,” *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, pp. 15–22, 2001.
- [20] C. Bossu Réaubourg, “Turbulence-sim: Real-time multiphysics lattice boltzmann solver.” <https://github.com/Wartets/Turbulence-sim>, 2026.
- [21] A. Zakai, “Emscripten: an LLVM-to-JavaScript compiler,” in *Proceedings of the fourth international workshop on academe and industry fitting together from research to emergence*, pp. 1–10, 2011.
- [22] T. Smith *et al.*, “Implementing WebAssembly SIMD in V8,” *V8 Dev Blog*, 2020.