

Computer Programs in Physics

HTMPC: A heavily templated C++ library for large scale particle-based mesoscale hydrodynamics simulations using multiparticle collision dynamics [☆]



Elmar Westphal ^{a, ID}, Segun Goh ^{b, ID}, Roland G. Winkler ^{b, ID}, Gerhard Gompper ^{b, ID, *}

^a Peter Grünberg Institute and Jülich Centre for Neutron Science, Forschungszentrum Jülich, 52425 Jülich, Germany

^b Theoretical Physics of Living Matter, Institute for Advanced Simulation, Forschungszentrum Jülich, 52425 Jülich, Germany

ARTICLE INFO

<https://go.fzj.de/HTMPC>

Keywords:

Mesoscale hydrodynamic simulation
Multiparticle collision dynamics (MPC)

ABSTRACT

We present HTMPC, a Heavily Templated C++ library for large-scale simulations implementing multi-particle collision dynamics (MPC), a particle-based mesoscale hydrodynamic simulation method. The implementation is plugin-based, and designed for distributed computing over an arbitrary number of MPI ranks. By abstracting the hardware-dependent parts of the implementation, we provide an identical application-code base for various architectures, currently supporting CPUs and CUDA-capable GPUs. We have examined the code for a system of more than a trillion MPC particles distributed over a few thousand MPI ranks (GPUs), demonstrating the scalability of the implementation and its applicability to large-scale hydrodynamic simulations. As showcases, we examine passive and active suspension of colloids, which confirms the extensibility and versatility of our plugin-based implementation.

Program summary

Program Title: HTMPC

CPC Library link to program files: <https://doi.org/10.17632/xnxh68zhbt.1>

Licensing provisions: MIT

Programming language: C++17, CUDA C++ (optional), MPI (optional)

Supplementary material: Supplementary Information, User Manual

Nature of problem: Complex fluids in soft, active, and living matter are characterized by a wide range of relevant length- and time-scales, from nanometers to millimeters, and from sub-microseconds to seconds. Their dynamics is often governed by the hydrodynamics of the embedding aqueous medium. Thus, it is essential for the numerical study of such systems to develop efficient simulation techniques and highly parallel computer codes, especially when large system sizes and emergent collective behavior are considered. Several mesoscale simulation techniques have been developed in the last decades for this purpose. Multi-particle collision dynamics (MPC), a particle-based hydrodynamics simulation technique, is a promising ansatz for such an endeavor. It is also important to develop an easy-to-extend implementation, so that the code can be adapted to various soft and living matter systems as desired.

Solution method: We develop an implementation of MPC that can exploit large-scale high-performance computing resources for hydrodynamic simulations of complex fluids. The code provides a C++ template library, which is plugin-based and can be extended by user-written plugins, implementing particles or objects interacting with the surrounding fluid. Calculations can be distributed over an arbitrary number of MPI ranks and accelerated with the current implementation supporting CUDA-capable GPUs. The code includes essential features of state-of-the-art MPC algorithms, e.g., thermostat, local angular-momentum conservation, and a variety of boundary conditions, such as periodic, no-slip (both also supporting shear flow) and slip. Simulation data can be written to and read from disk.

Additional comments including restrictions and unusual features: The code provides an option to produce perfectly reproducible particle trajectories, independent of the MPI setup of a simulation system, as long as the underlying

[☆] The review of this paper was arranged by Prof. David W. Walker.

* Corresponding author.

E-mail address: g.gompper@fz-juelich.de (G. Gompper).

<https://doi.org/10.1016/j.cpc.2024.109494>

Received 25 July 2024; Received in revised form 13 December 2024; Accepted 26 December 2024

1. Introduction

Solid bodies immersed in fluids experience hydrodynamic forces by local fluid motion, which can be of thermal origin on micro- and sub-micrometer scales or initiated by the motion of other immersed objects. Such hydrodynamic interactions are long range in nature and therefore determine the dynamical and structural properties of suspensions. They are particularly relevant for soft matter systems, which encompass traditional complex fluids such as amphiphilic mixtures, colloidal suspensions, polymer solutions, but also biological matter [1]. Soft matter systems are characterized by a wide range of relevant length and time scales – from the atomic scales of solvent particles to the mesoscale of colloidal particles, and the macroscopic scale of phase transitions – and they are sensitive to thermal fluctuations and to weak perturbations [2]. Theoretical and simulation studies reveal the influence of hydrodynamic interaction on the dynamics of polymers [3–5] and colloids [6,7] in solution in absence and presence of flow fields [4,8]. Similarly, biological fluids [9] share many rheological properties with soft-matter systems. For example, in blood flow, not only the collisions between the deformable red blood cells but also the hydrodynamic interactions due to flow of the embedding blood plasma are essential [10,11]. Cellular flows may also involve active components, so that their behavior can then be modeled as active fluids. Nature provides a plethora of autonomous microswimmer [12,13], particular examples are flagellated bacteria [14,15] and sperm cells [16], which propel themselves via hydrodynamic interactions.

The presence of the disparate time, length, and energy scales in soft matter systems renders the application of conventional (all atom) simulation techniques challenging. To mimic the behavior of atomistic systems on the length scale of the mesoscopic objects requires the application of “coarse-grained” or mesoscopic approaches [1]. Accordingly, mesoscale simulation methods have been developed such as Dissipative Particle Dynamics (DPD) [17,18], Lattice Boltzmann (LB) [19–21], Direct Simulation Monte Carlo (DSMC) [22,23], and multiparticle collision dynamics (MPC) [1,24].

Among the mesoscale simulation techniques, MPC is particularly suitable for simulations on massively parallel GPU-based supercomputers [25,26]. The main advantage is that the interactions and momentum exchange between the fluid particles occur locally in collision cells, which renders the calculation of large inter-particle distances obsolete, but still gives rise to long-range hydrodynamic interactions [5,27]. In addition, the streaming of every fluid particles is carried out independently. Moreover, MPC incorporates thermal fluctuations by construction. Therefore, MPC is suitable for massive parallel computing and can be utilized to investigate large thermal systems beyond scales that a single processor can handle.

Since the original formulation in 1999 [28], several variants and revisions of MPC have been proposed [1], and it has been applied to a wide range of physical problems. Particularly, the transport properties of MPC fluids have been obtained theoretically [29–33] and its hydrodynamic correlation function has been derived [27]. The applications cover such diverse aspects as sedimenting colloidal suspensions [34,35] and polyelectrolyte electrophoresis [36], liquid crystals [37], polymeric systems and protein solutions [4,38–47], ferro- [48] and nanofluids [49], and even self-propelled particles [50–54], active nematics [55] and active polymers [56,57], as well as biological systems such as suspensions of bacteria [14,15,58–60], sperm [61–63], and parasites, e.g., trypanosomes [64] and *Plasmodium falciparum* [65].

In this article, we present a highly efficient parallel implementation of MPC, which is suitable for both CPU- and GPU-based computers, and has been evolved from the previous GPU-based implementation pre-

sented in Ref. [25]. Our new code has been tested up to 1.5 trillion MPC fluid particles on 3,540 NVIDIA A100 GPUs.

2. Multiparticle collision dynamics

In MPC, a fluid is represented by N_s point particles of mass m with the positions \mathbf{r}_i and the velocities \mathbf{v}_i ($i = 1, \dots, N_s$). Their time evolution is described by a two-step dynamics of alternating streaming and collision steps. In the streaming step, the MPC particles move ballistically according to

$$\mathbf{r}_i(t+h) = \mathbf{r}_i(t) + h\mathbf{v}_i, \quad (1)$$

where h is the MPC time step. The inter-particle interactions are modeled via a momentum conserving stochastic process in the collision step [1,24,28]. Here, all MPC particles are sorted into the cells with side length a of a simple cubic lattice, each of which contains on average $\langle N_c \rangle$ particles. Then, in the stochastic rotation dynamics (SRD) variant of MPC [1], the relative velocities $\mathbf{v}_{i,c} \equiv \mathbf{v}_i - \mathbf{v}_{cm}$ of the particles within a cell with respect to the center-of-mass velocity $\mathbf{v}_{cm} \equiv (\sum_{i=1}^{N_c} \mathbf{v}_i)/N_c$ of the cell are rotated by a fixed angle α around a randomly oriented axis according to

$$\mathbf{v}_i(t+h) = \mathbf{v}_{cm}(t) + \mathbf{D}(\alpha)\mathbf{v}_{i,c}(t), \quad (2)$$

where \mathbf{D} denotes the rotation matrix. The orientation of the rotation axis is independently chosen for every cell and collision step. Notably, the entire collision lattice is randomly shifted before every collision step to insure Galilean invariance [29,66].

Several theoretical and technical advances have been achieved in the course of the development of MPC. Here we summarize some of the important issues, which have also been taken into account in our implementation.

Angular-momentum conservation Violation of angular-momentum conservation may lead to non-physical torques due to non-symmetric components in the stress tensor [67–69]. To avoid such artifacts, cell-level angular-momentum conservation can be enforced (MPC-SRD+a) by adding an additional constraint to the velocity in Eq. (2) at each collision step [68,70],

$$\mathbf{v}_i(t+h) = \mathbf{v}_{cm} + \mathbf{D}(\alpha)\mathbf{v}_{i,c} + \boldsymbol{\omega}_c(t) \times \mathbf{r}_{i,c}, \quad (3)$$

where $\mathbf{r}_{i,c} \equiv \mathbf{r}_i - \mathbf{r}_{cm}$ is the relative position with respect to the center-of-mass position \mathbf{r}_{cm} , and the angular velocity is given as

$$\boldsymbol{\omega}_c = m\mathbf{I}_c^{-1} \sum_{i=1}^{N_c} (\mathbf{r}_{i,c} \times (\mathbf{v}_{i,c} - \mathbf{D}(\alpha)\mathbf{v}_{i,c})). \quad (4)$$

Here, \mathbf{I}_c is the moment-of-inertia tensor of the particles in the collision cell with respect to the center of mass.

Thermostat As the MPC-SRD algorithm given by Eq. (2) conserves energy, the fluid shows thermodynamic behavior of a microcanonical ensemble. Whenever a canonical ensemble is appropriate, i.e., a constant temperature is desired, the system has to be thermostatted. This is achieved by a Maxwell-Boltzmann scaling (MBS) approach [71], which rescales the relative velocities in the cells after each collision step via the scaling factor κ

$$\mathbf{v}'_i = \mathbf{v}_{cm} + \kappa\mathbf{v}_{i,c}. \quad (5)$$

The scaling factor itself is determined from the kinetic energy of the particles within a cell, i.e.,

Schematic overview for an example use case:

(c) Plugin Call Sequence

(a) Partial Steps: INIT, START, REORDER, MOVE, PROCESS, PROCESS_CELLS, ROTATE_VELOCITIES, FINALIZE

(b) Implemented Specializations

Plugin	Performing at steps
Properties	INIT, FINALIZE
Boundaries	(none)
Domain_decomposition	INIT
Cells	INIT, START, PROCESS_CELLS
Fluid	INIT, START, REORDER, MOVE, PROCESS, ROTATE_VELOCITIES

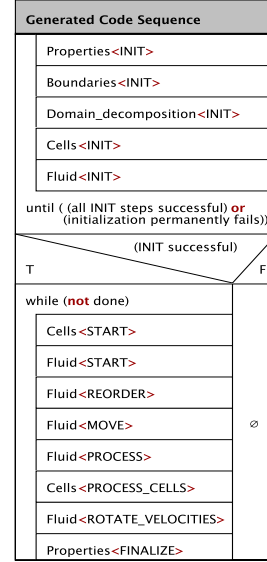


Fig. 1. Schematic overview for an example use case. (a) Partial steps are defined as an enumeration in the configuration class. (b) Plugins containing specialized functor templates to be called at partial steps listed in (a). (c) The plugin call sequence ultimately generated at compile time.

$$\kappa = \sqrt{\frac{2E_k}{\sum_i m v_{i,c}^2}}, \quad (6)$$

where the random kinetic energy E_k suffices the distribution function

$$P(E_k) = \frac{1}{E_k \Gamma(f/2)} \left(\frac{E_k}{k_B T} \right)^{f/2} \exp \left(-\frac{E_k}{k_B T} \right), \quad (7)$$

with $f = 3(N_c - 1)$ the number of degrees of freedom and Γ the gamma function.

Boundary conditions In addition to periodic boundary conditions, the fluid can be in contact with confining walls and surfaces of embedded particles such as colloids. Slip as well as no-slip boundary conditions can be employed in MPC. During the streaming step, collisions between MPC particles and boundaries result in a momentum change. Specifically, all the MPC particles which have crossed boundaries are translated back onto the boundaries. For no-slip boundary conditions, the bounce-back rule is applied, where the velocity is simply reverted, i.e., $\mathbf{v}_i \rightarrow -\mathbf{v}_i$. For slip-boundary conditions, specular reflection is employed. In case of, e.g., interactions with solid bodies such as colloidal particles, momentum conservation implies a change of the linear momentum of MPC particles as well as the linear and angular momenta of the colloid. In the collision step, collision cells are added either in walls or larger immersed particles. The part of these collision cells in walls or immersed particles are filled with uniformly distributed phantom particles, which are identical to MPC particles, such that the whole collision cell contains the same average number $\langle N_c \rangle$ of fluid particles [33,72]. This enhances wall-fluid and colloid-fluid interactions and reduces boundary slip [33,72]. Specifically, it is required due to the random shift of the collision lattice, otherwise the particle number varies in cells cut by boundaries. The velocities of the phantom particles are taken from a Gaussian distribution with zero mean and variance $\sqrt{k_B T/m}$. These phantom particles participate in the collision step as all other fluid particles.

External fields A constant gravitational force along a coordinate axis in combination with two parallel confining walls leads to a parabolic Poiseuille flow [73]. Importantly, in this case, the MPC fluid has to be thermalized [74].

Parameters The simulation of fluid properties requires a suitable choice of the MPC parameters [5,34]. By the nature of the MPC collision process, hydrodynamics in the fluid appears on length scales λ_c larger than the collision cell only. The actual length scale itself depends on the collision time step and increases with increasing collision step [27]. Moreover, in fluids, momentum transport dominates over mass transport. This is expressed by the Schmidt number $S_c = \nu/D$, where ν is the kinematic viscosity and D the diffusion coefficient of the MPC particles [75]. In MPC, a large Schmidt number is achieved for small collision step sizes and large collision angles. Hence, we recommend to choose these values in the range $h \leq 0.1 \sqrt{ma^2/(k_B T)}$ and $\alpha > 90^\circ$ [27,75].

Proper hydrodynamics for objects embedded in MPC such as colloids or polymers requires an appropriate choice of their size. Because of the condition $\lambda_c/a > 1$, a colloid radius should be $R_c/a \geq 3$ (see Sec. 5) [5]. Similarly, in a simulation of a polymer its bond-length and monomer diameter should be larger than the collision cell size [39].

Due to the point-particle nature, a MPC fluid is compressible with an ideal gas equation of state. The compressibility can be controlled by the average number of fluid particles in a collision cell. For systems with passive colloids or polymers, $\langle N_c \rangle = 10$ is recommended. However, in systems with self-propelled particles, e.g., squirmers, a larger average number $\langle N_c \rangle$ may be necessary to avoid depletion effects, particularly at large concentrations; here, we suggest $\langle N_c \rangle = 50$ for Péclet numbers in the range $Pe \sim 10^1 - 10^2$ [53,54,76] (see Sec. 5).

3. Implementation

3.1. Code structure

Basic structure The code is a C++17 template library. It consists of a class template representing a simulation environment using plugins that contain the simulation data and/or the code to manipulate it. This environment is generated using rules and definitions provided by a descriptor class that is passed as a template argument. This class contains a sequence of plugin templates, an enumeration of partial simulation steps to be executed in every iteration (see Fig. 1), an arbitrary number of type-declarations and constexpr (compile time) variable-declarations and a sequence of parsers to interpret these declarations.

Partial simulation steps In each MPC step a number of operations is performed on the simulation data. To coordinate these operations, each iteration step can be divided into an arbitrary number of partial steps, such as the movement of fluid particles, the generation of rotation matrices, etc. As the partial steps have to be performed in a given order, they are defined as an enumeration. The code allows the user to extend this basic enumeration by inserting an arbitrary number of additional steps into it, provided all existing steps remain in their original order.

Plugins The data and code for the simulation are stored in plugin class templates. These are listed in a sequence given in the configuration class and are instantiated according to rules also given there. In the generated simulation class, instantiations of all plugin templates specialized by the configuration class are inherited and can be accessed in different, configurable ways. The plugins may contain functor templates specialized for each partial iteration step, containing the code to be performed at that particular step. Besides data and code to be performed regularly at certain times in the iteration step, plugins can also represent boundary conditions by adding additional functor templates that can be applied to particle data in certain situations such as particle movement or processing. The included plugins implement the MPC fluid, MPC collision cells, different boundary conditions, etc. Additional plugins can be added following the given interface.

Generating iteration steps As described, each iteration step is divided into a number of partial steps. During each iteration step the partial steps will be performed in the order given in their enumeration. During each partial step, all plugins will be traversed in the order given in the plugin sequence and if they contain a functor template specialized for that particular step, the functor will be instantiated to be called. The only exception to this is the initialization step, the functor of which is only called once after the simulation environment is constructed. Only during initialization, it is possible to change the order of execution of the plugins, making it possible to process dependencies between them. Otherwise, all dependencies must be implicitly resolved by the given order of plugins. Regular operation will commence as soon as all plugins are initialized. If the initialization of one or more plugins fails permanently, the program will be aborted.

3.2. Distributed simulation

The code is designed to optionally distribute the workload over an arbitrary number of MPI ranks, which is achieved by splitting the simulation system into spatial subdomains. In addition, particles are reassigned to the subdomains covering their current position at regular intervals. Between these reordering operations, particles that leave their subdomains are temporarily sent to the appropriate subdomain for processing and their updated velocities are imported back.

3.2.1. Domain decomposition

The amount of necessary communication depends on the number of particles exchanged, and therefore, the surface area of the subdomains. For optimal performance, the domain decomposition algorithm minimizes the surface area of subdomains, by generating subdomains close to cubic shapes. To this end, the total number of subdomains given by the number of MPI ranks is factorized at first. The currently longest side of the system will then be subsequently divided by the largest remaining factor in a loop over all factors from the factorization. If a shear flow is considered, the system is not split along the corresponding discontinuous axis to avoid difficulties during particle exchanges resulting from the fact that we can not define a fixed neighborhood relation between subdomains due to the continuous shifts in shear-offset across the boundaries. Lastly, the number of remaining collision cells in each direction, given as the remainder of the system size divided by the number of domains, will be distributed evenly over the subdomain sizes, i.e. a size of 20 splits into 3 will result in partial sizes of (7,7,6).

3.2.2. Temporary exchange of particles for processing

During every iteration step, the respective collision cell of each MPC particle is calculated from the position of particles, the random shift of the current iteration, and applicable boundary conditions. If the collision cell of a MPC particle is outside the subdomain the particle is currently assigned to, then the particle must temporarily be transferred to the respective subdomain and processed there. This implementation was chosen over the usage of overlapping subdomains for several reasons: while every particle is always a member of exactly one cell, overlapping cells may be shared between 2, 4 or 8 subdomains for surface, edge or corner cells respectively, leading to more complex communication patterns when distributing partial and redistributing combined results. Depending on the configured features, this may happen more than once during each iteration step, because the change in angular momentum of each cell can only be calculated later in the step. Additionally, each fluid particle traveling beyond the overlapping cells of its assigned subdomain would require the whole system to be spatially reordered immediately, because otherwise its properties could not be processed properly. This would require additional monitoring effort and possibly more frequent expensive reordering.

Each subdomain has one export buffer - similar to the fluid particle storage - for preparing sets of particles to be sent to other domains, and one import buffer for receiving particles from other domains. These buffers only contain the data members necessary for the currently activated feature set. Whenever a particle needs to be exported, the counter of particles to be sent from the current subdomain to the destination subdomain is incremented, and a data triplet of the particle index, destination subdomain, and previous export count (as a destination-relative index) is appended to a vector. After finding all particles to be exported, the prefix sum of the export counts of destination subdomains is calculated, which represents the offsets for each destination subdomain in the output buffer. Then the vector of data triplets is traversed and the data of each particle contained in the traversed vector is copied to the export buffer position calculated by adding the destination offset from the prefix sum and the destination-relative index of the particle. In this way, a buffer containing particle data is ordered by their respective destination subdomains, which is efficient especially on GPUs, as each step can be performed in parallel and the export buffer for all domains can be built with a single kernel call.

In almost all cases fluid particles are received from direct neighbors and the exchange is done using a `MPI_Neighbor_alltoall` collective call. Otherwise, a more expensive fallback path using `MPI_Alltoall` is used. Then a prefix sum is calculated over the import counts, resulting in the import offsets for each subdomain. Such obtained import- and export-counts and -offsets are then used in neighbor-collective or discrete MPI-calls to exchange the necessary data members, which can be done without explicitly transferring data to/from the GPUs by using a CUDA-aware MPI implementation. Now the code is ready to perform all remaining MPC steps, after which the processed particle velocities are transferred back to their original subdomains and copied back from the exchange buffer to their respective particles.

3.2.3. Permanent exchange of particles during reordering

During the course of the simulation, a growing number of particles move away from the subdomain they were originally assigned to, resulting in a higher communication load as well as a growth in buffer size. To avoid such penalty in performance, the code reassigns (or reorders) particles to the domains where they are most likely processed. Reasonable reorder intervals must balance the time spent for reordering particles and the time gained through faster operation, which also depends on the time step length and system architecture. In most setups, an optimal interval lies between ten and a few hundred iteration steps. The algorithm for transferring the particles is the same as described in Sec. 3.2.2 for temporary particle transfers. After transfer, however, the relocated particles are inserted and deleted to/from the permanent particle sets of the domains.




Fig. 2. Influence of particle reordering on runtime for 200 iteration steps with AMC. Here, $L = 128a$, $\langle N_c \rangle = 10$, $h = 0.05a\sqrt{m/(k_B T)}$, and $\alpha = 130^\circ$. The test runs were performed on a GTX2070 GPU. The definitions of labels are given in Table S1 (SI). (For interpretation of the colors in the figure(s), the reader is referred to the web version of this article.)

During the exchange step, the particles in each subdomain are also spatially reordered, which is beneficial for performance in general, as shown in Fig. 2, and therefore, is also applied to single-rank systems. Reordering the particles in each subdomain proceeds similarly to the filling of export buffers described above, but particles are ordered by collision cells instead of destination subdomains. Particle counts per cell are determined, using the particles still residing in their current subdomain and, in case of distributed systems, particles received from other subdomains in the import buffer. The current particle counts per cell are again used as cell-relative indices and the prefix sum of the final particle counts is used as cell offsets for the reordered particles.

To save memory, reordering is done property-wise, which necessitates only a temporary reorder buffer with the size of the largest data member of the particles (usually the velocity). Specifically, each data member of the currently resident particles is copied to the temporary buffer. The original data storage is resized to fit the data of all remaining resident plus imported particles. Then the data of particles in the temporary buffer that are not exported and those from the newly imported ones in the import buffer are copied to their new positions in the resized main data location. It is also possible to perform this step component-wise for each particle property, to reduce the reorder buffer size by a factor of three, which in turn, corresponds to roughly 15-25% of the total memory footprint. Due to inefficient memory access patterns and repeated access, however, it comes with a cost of increasing the runtime by a few percent, the exact value of which depends on the reorder intervals.

3.3. Reproducibility

Additions are not always associative for floating point values, i.e., $(a + b) + c$ is not necessarily equal to $a + (b + c)$ [77,78]. In a multi-threaded or distributed code, such non-associativity causes diverging results between runs with identical initial conditions, because the order of operations is undefined and also depends on the distribution of data. Even small fluctuations will result in significantly different particle trajectories after a sufficient number of iterations. While each of these runs is still valid, perfectly matching trajectories may be desirable, for justification or debug purposes.

To achieve this, the code optionally provides bit-perfect reproducibility of particle trajectories by replacing the floating point summations with fixed point summations, where fixed point numbers assign a fixed count of bits to the integral part [79]. If the resulting range is

not defined appropriately, overflows may occur leading to undefined behavior (see e.g., Ref. [80]) and invalidating the results, which implies an inherent tradeoff between range and precision. For typical simulations the range of values of the cell parameters is limited and can be well estimated, allowing us to determine the necessary precision for fixed point numbers. The generation of initial particles is also adapted accordingly to achieve perfect reproducibility not only on a single GPU, but also on distributed runs, independent of rank counts or domain decompositions, as long as the same system architecture, compiler, and compiler settings are used. Even though the summations themselves should not be influenced by associated parameters, different hardware components, numerical algorithms or optimizations may introduce subtle changes in the results of other operations. Additionally, the order of particles in memory or in files saved to disk is undefined due to their dependence on the domain decomposition layout and the undefined order of atomic operations.

Fluid particle ID In the given implementation, unless particle reordering is switched off in an undistributed system, it is not possible to follow certain particles by their offset alone due to the regular exchange of particles between subdomains and/or particle reordering for performance purposes. The code provides an option that assigns an integer tag to each fluid particle. If the respective option is enabled, this tag is, by default, generated as a consecutive, unique number for each particle, which is independent of the domain decomposition for otherwise identical systems. Opposed to the current storage offset or subdomain, the ID is permanently assigned to a certain particle and can be used to trace that particle over the course of a simulation. In combination with the reproducibility feature, fluid particles with the same unique initial ID are guaranteed to have identical trajectories over different runs for an identical system setup. The ID-value can, however, be changed by the user at any time during the simulation to arbitrary (and also non-unique) numbers and can therefore be adopted to i.e. mark or group certain particles.

3.4. Additional code features

Here additional code features are summarized. More detailed information can be found in Supplementary Information (SI).

Thermostat The implementation of the Maxwell-Boltzmann scaling (MBS) thermostat [71] mainly consists of two components. First, the kinetic energies of all particles are summed up cell-wise during particle processing. Then a scaling factor for each cell is calculated during cell processing, which is applied to all particle velocities in the cell during the particle rotation. Applying the MBS thermostat increases total calculation time by approximately 10%, as shown in Fig. 3.

Local conservation of angular momentum Angular-momentum conservation (AMC) adds new calculation steps and a significant amount of additional calculations to several other steps. Consequently, conserving local angular momentum in combination with the thermostat option will increase calculation times by a factor larger than 2, see Fig. 3. Due to the significantly higher number of operations and their respective error propagation, the residual momentum is roughly two orders of magnitude higher than for plain MPC.

Boundary conditions Boundary conditions are implemented as special plugins with additional functionality. Like all other plugins, they may have their own state, data, and code. Additionally, they have a functor template that can be specialized for certain situations, usually when a particle is moved or when it is processed after random shift of the collision-cell lattice was applied. Different boundary conditions can be set for different axes of the simulation system. Using the appropriate interface, boundary condition plugins can be added ad libitum. The included plugins support periodic, slip, and no-slip boundary conditions. Optionally, shear flow can also be added.




Fig. 3. Feature dependence of the time spent in partial steps. Here, 1000 iteration steps with $L = 100a$, $\langle N_c \rangle = 10$, $h = 0.05a\sqrt{m/(k_B T)}$, $\alpha = 130^\circ$, and a reorder interval of 200 were performed on a GTX2070 GPU. The definitions of labels are given in Table S1 (SI).

Phantom particles The code provides specialized particle sets that can be used by plugins for the efficient processing of phantom particles. In contrast to standard fluid or solute particles, phantom particles are not subject to movement or distribution, but are created in the subdomain where they are to be processed while random shifts are taken into account. Plugins generating phantom particles initiate their processing, rotation of velocities, and restoration of their local angular momentum at appropriate times, if applicable.

Gravity It is possible to apply a constant gravitational force parallel to one of the axes. The application can be combined with other options and the implemented boundary conditions.

3.5. Hybrid CPU/GPU implementation

The code is designed for distributing the system across an arbitrary number of domains, each of which is processed by one MPI rank. The implementation of intra-domain operations favors GPU acceleration, but also works on one CPU core per MPI rank. CPU-based shared memory parallelization like OpenMP for intra-domain operations was also tested, but the performance turned out to be lower than that of a more fine-grained domain-decomposition distribution.

3.5.1. Identical code base for different architectures

Most of the code processes either particle- or cell-based operations. Each operation is implemented as a functor which is called once for every cell or particle, and is coded to be hardware agnostic for better maintainability and compatibility between platforms. Hardware dependent operations, possibly leading to atomic operations, are hidden in wrapper classes that are implemented for their designated target architecture. Currently, plain CPU code and CUDA-based GPU acceleration on NVIDIA GPUs are implemented, but the concept is modular and could be extended to other architectures without the need to rewrite any of the actual application code.

3.5.2. GPU memory management

Accessing CPU memory from a GPU or vice versa is usually not possible, or comes with a significant performance penalty for most of the relevant past and current GPU architectures. In the code, we employ the CUDA-based implementation providing an allocator for so-called Managed Memory, which is a mechanism providing automated data migration on older GPUs or memory paging on more recent hardware, avoiding the need to explicitly migrate data. Still, alternating access to the same data elements from the CPU and GPU leads to implicit data movement and can significantly hurt performance. For instance, access

to seemingly independent data residing on the same memory page may cause performance issues related to concurrent access of this page by the application running on the GPU and the MPI library running on the CPU. These issues have been addressed by padding managed memory allocations.

3.5.3. GPU acceleration

Code generation and architecture selection is based on C++17 template mechanisms. The only external dependency is an appropriate CUDA compiler. The architecture dependent functionalities that had to be implemented include, among others, an STL-compatible, GPU-based vector template, reduction routines, prefix sums and accumulator classes wrapping atomic additions. As more and more standard library routines are provided in GPU-parallelized versions, some implementations may be updated accordingly in later releases.

3.5.4. Atomic operations

While early releases of our GPU-based MPC routines used a list-based approach for processing the particles belonging to each collision cell [25], the current version uses atomic operations for all summations of this kind. Atomic operations used to be very expensive in early GPU architectures, especially for double precision additions which had to be emulated by loops instead. Recent architectures, however, provide native implementations for double precision additions and much faster support for atomic operations in general. Additionally, even with most spatially ordered particles, the list-based access to particle data results in large amounts of data that have to be read multiple times due to the need of reading data blocks of certain sizes (cache lines), while an atomic-operation based implementation allows for coalesced reading patterns. On the other hand, a list-based approach may show better writing patterns if particles are distributed evenly among cells. As shown in SI (Fig. S1), the relative performance of these approaches differs significantly by GPU architecture and slightly by fluid density. In the few cases where the performance is comparable, cell-list based implementations carry an additional memory-footprint for the list. Therefore, the list approach was abandoned after “Kepler” became the prevalent GPU architecture.

Certain operations such as binning data for the calculation of optional shear profiles, may result in extraordinary high numbers of concurrent atomic operations on a small amount of data, impairing the performance. In this case, optimized local summation patterns may significantly reduce the number of atomic operations [81]. The code may activate such optimizations, depending on their expected efficiency for certain access patterns and GPU architectures.

3.5.5. Random numbers

For an efficient GPU implementation, we choose the Xorshift algorithm [82] as the default random number generator (RNG), which can be implemented in a branch- and loop-free manner, while providing good randomness. Specifically, a system is initialized with a main random seed, which can be either passed as a parameter or generated from the random source of the operating system. Using the main random seed only might lead to the same or at least dependent random values for every cell or necessitate an atomic- or lock-based mechanism for accessing the random seed causing a significant performance penalty. Therefore, the cell random seeds are necessary for processing certain parts in parallel, e.g. generating the random rotation matrix per cell. We note that the modular concept of the code also allows users to replace it by a user-provided RNG with a matching interface.

3.6. Speed boundaries

The computation speed of operations is limited by certain boundaries like the maximum amount of data transferred from/to memory per second or the maximum number of floating point or atomic operations per second that can be performed by the CPU or GPU. Taking into

Table 1

The amounts of data for each MPC particle. * optional; ¹ mandatory for distributed systems, recommended otherwise; ² depending on maximum system size; ³ depending on maximum particle count.

Particle property	Size (bytes)
Position	12
Velocity	24
Cell index ¹	4-8
Reorder buffer ²	8-24
ID-tag ^{*,3}	4-8
Total	48-76

Table 2

The amounts of data for each MPC cell. * optional; ¹ if using thermostat; ² if conserving angular momentum.

Cell property	Size (bytes)
Center of mass velocity	24
Combined mass of particles	8
Particle count	4
Random seed	8
Energy/thermostat ^{*,1}	8
Angular momentum ^{*,2}	24
Center of mass ^{*,2}	24
Inertia Tensor ^{*,2}	48
Total	44-148

account all particle- and cell-related read/write operations, i.e. fluid particle propagation, at least on GPU, is bound by the memory bandwidth. The combination of the rotation of particle velocities and the calculation of several cell parameters needed for local angular-momentum conservation as implemented in our code is the most expensive step, both calculation- and memory-transfer-wise. Even though it involves most of the algorithm's atomic operations, it still runs at approximately 60% of the theoretical memory bandwidth. Without angular-momentum conservation and the atomic operations involved, the rotation of particles is also memory bound.

3.7. Simulation data

3.7.1. Data types and memory management

The storage and calculation precision for most values and operations can be adjusted in the configuration class. The default settings are single-precision floating-point numbers for the fluid particle positions and double precision for their velocities [25]. All calculations on cell properties are performed and stored in double precision. Precisions for suspended particles, e.g., in Molecular Dynamics (MD) plugins can be set as necessary. The amounts of data stored for fluid particles and collision cells in the default settings are summarized in Tables 1 and 2, respectively.

For distributed systems, each subdomain reserves a slightly larger fluid particle capacity to leave room for variations in density between domains. Additionally, import- and export-buffers are needed for particle exchange, whose size depends mostly on the surface area of the domains. As a general rule, the initial buffer size is $\langle N_c \rangle$ multiplied by the number of surface cells, about half of which is caused by the random shift of up to a collision-cell size a for all axes and the remainder to process particles moving out of their assigned subdomain. The number of the latter is reduced by spatially reordering fluid particles at certain intervals. Because memory allocation is an expensive operation, especially on GPUs, the import and export buffers are reserved slightly above the expected size and may be extended by reallocation if needed, but never

shrunk. Additional memory may be used by other plugins. One example is boundary conditions representing no-slip surfaces, which need one cell layer to be partially (depending on the random shift) filled with phantom particles for fluid-surface coupling. If the top and bottom no-slip boundaries are located in separate domains, both domains will need this one-cell-layer-wide buffer, otherwise one buffer will be shared.

3.7.2. Input/output

It is possible to store and read complete system data to/from disk. The state variables (fluid particle positions and velocities) and parameters (including density, random seed, and simulation step count) of the system are written into an easy-to-parse text file, while all particle data and select cell data is written in a binary format. Before writing data, all particles are ordered by their unshifted cell index. The cell offsets are calculated and also saved to enable reading the saved data back into systems of different rank counts or domain decompositions. All data is written to disk by using collective MPI I/O. The space necessary to store a system can be determined by multiplying the number of fluid particles with the storage sizes of their positions, velocities and IDs (if applicable) from the information in Table 1. Additionally, for each cell the random seed, particle count and offset are saved as listed in Table 2. As the amount of header information is negligible, for system of default precision, the estimated amount of data stored reads

$$\begin{aligned}
 D &= N \times [\text{sizeof(position)} + \text{sizeof(velocity)} + \text{sizeof(ID)}] \\
 &\quad + N_{\text{cell}} \times [\text{sizeof(random seed)} + \text{sizeof(cell offset)} \\
 &\quad \quad + \text{sizeof(particle count)}] \\
 &= N \times (12 + 24 + 8) + N_{\text{cell}} \times (8 + 8 + 4). \tag{8}
 \end{aligned}$$

For systems with 2^{32} or less particles, the ID tags and cell offsets are saved as 32 bit integer values requiring 4 bytes of storage space. It is also possible to generate arbitrary particle IDs for debug purposes when reading data that were saved in runs without particle IDs.

4. Performance/benchmarks

Benchmarking is performed on the JUWELS Booster System at Jülich Supercomputing Center [83]. Each of the nodes provides 2 AMD EPYC CPUs (24 cores each) and 4 Nvidia A100 GPUs. The nodes are connected via $4 \times \text{HDR200}$ Infiniband each. Systematic benchmark tests were performed utilizing up to 3072 GPUs or 12288 CPU cores for a wide range of system sizes from $64 \times 64 \times 64$ to $4096 \times 4096 \times 2048$ collision cells at low and high particle densities, i.e., $\langle N_c \rangle = 10$, and 50, respectively, with and without local angular-momentum conservation. While the CPU- and GPU-based measurements for $\langle N_c \rangle = 10$ (Fig. S2), as well as a benchmark comparison to the HOOMD-blue MPC code (Fig. S3) are presented in SI, we mainly focus here on GPU-based measurements for $\langle N_c \rangle = 50$.

Strong scaling During each step, the MPC algorithm performs a relatively small number of calculations on many particles. Therefore maximizing data throughput, which can be achieved by localizing data, is the key to good performance. Accordingly, the strong scaling presented in Fig. 4(a) and (b) shows good efficiency as long as the workload per GPU is high. The scaling efficiency of setups with computationally more expensive options like local angular-momentum conservation or large densities (see also Fig. S1) decreases more slowly, because of their better GPU utilization. For $\langle N_c \rangle = 50$ with the AMC option, a speed-up by a factor of 10 is achieved roughly at 70% efficiency. With a growing number of ranks, communication times become prevalent, as shown in Fig. 5(a) and (b). On the other hand, communication times can be kept on the order of 10% even for 1024 GPUs for sufficiently large simulations, as shown in Fig. 5(c) and (d).

Weak scaling Weak scaling depends on the ratio of the performance of nodes for a given load and the additional communication effort neces-

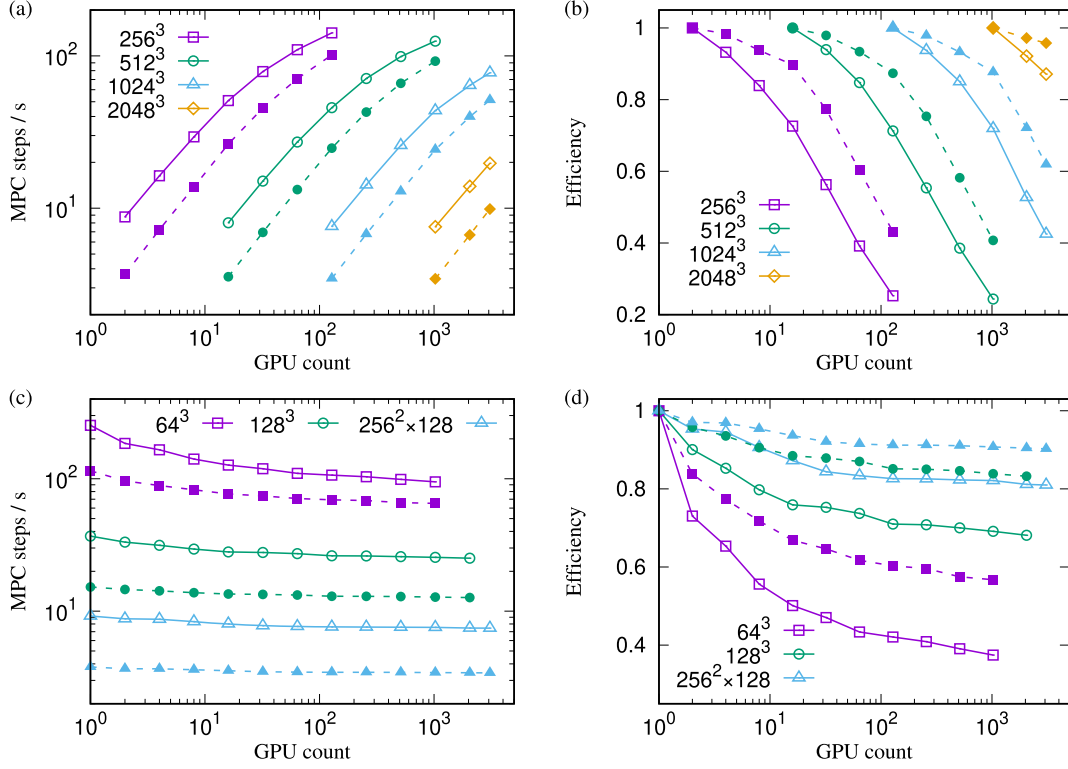


Fig. 4. (a,b) Strong and (c,d) weak scaling behavior for $\langle N_c \rangle = 50$. (a,c) The number of MPC steps per second, as well as (b,d) the corresponding efficiency are shown. Filled symbols (dashed lines) and open symbols (solid lines) represent scaling behaviors of the code with and without AMC, respectively. Here, we use $h = 0.05a\sqrt{m/(k_B T)}$, $\alpha = 130^\circ$, and a reorder interval of 200.

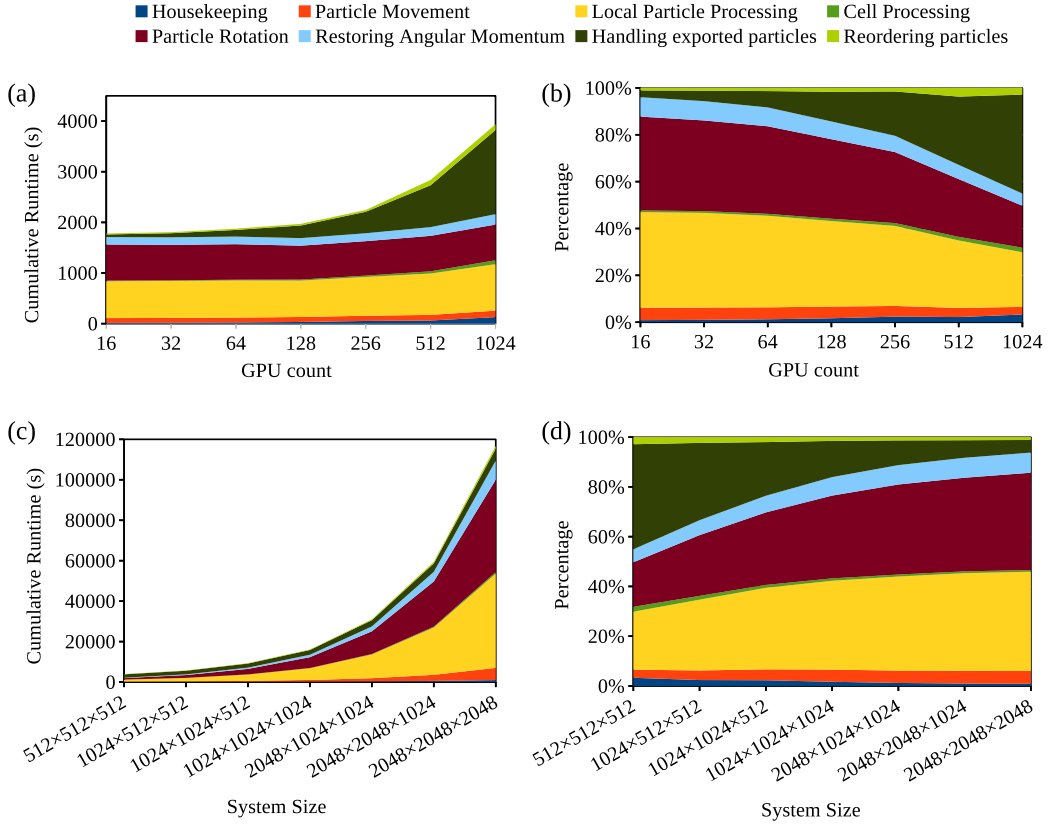


Fig. 5. (a), (c) Cumulative runtime distributions for different tasks, values are summed up runtimes for all ranks, and (b), (d) the percentage of the cumulative runtime occupied by each component. While the results obtained from a fixed system size of $512 \times 512 \times 512$, distributed over different GPU counts as indicated, are presented in (a) and (b), the dependence on system sizes, benchmarked on 1024 GPUs, are shown in (c) and (d). In all cases, $\langle N_c \rangle = 50$, $h = 0.05a\sqrt{m/(k_B T)}$, and $\alpha = 130^\circ$ with a reorder interval of 200, and 400 MPC steps were performed. See Table S1 (SI) for the definitions of labels.

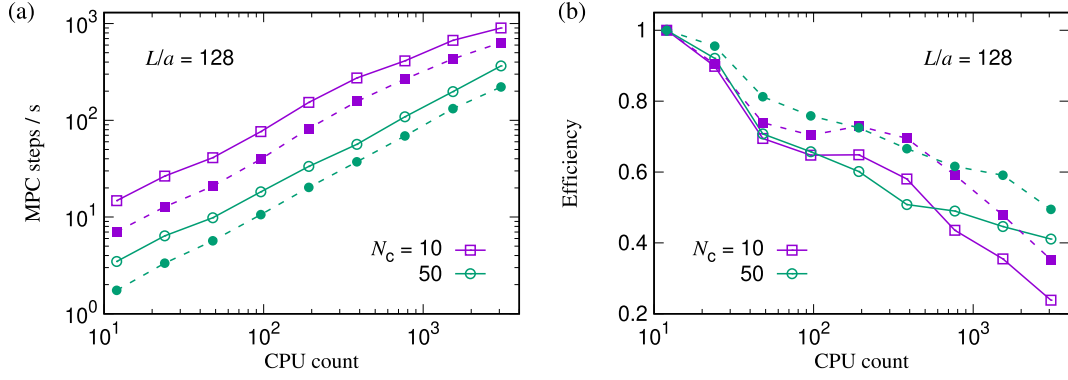


Fig. 6. Strong scalability for a small system size of $L/a = 128$ on CPUs. (a) Number of MPC steps per second and (b) corresponding efficiency for two different MPC densities, as indicated. Filled symbols (dashed lines) and open symbols (solid lines) represent scaling behaviors of the code with and without AMC, respectively. Here, we use $h = 0.05a\sqrt{m/(k_B T)}$, $\alpha = 130^\circ$, and a reorder interval of 200.

sary for adding more nodes with the same load. As shown in Fig. 5, the ratio of calculation versus communication for MPC on GPUs depends heavily on the utilization of the GPUs. This leads to very good weak-scaling behavior already for medium sized workloads, presented in Fig. 4(c) and (d). Weak scaling efficiency for large workloads is almost ideal even for very high GPU counts.

Small system size To demonstrate that HTMPC is also efficiently applicable to inhouse computer clusters with a limited number of nodes, we also perform benchmarks with a small system size of $L/a = 128$, and examine the strong CPU-scalability. As shown in Fig. 6, our code also performs efficiently for both low and high densities of $\langle N_c \rangle = 10$ and $\langle N_c \rangle = 50$. Good scalability with a speedup by a factor of 10 or 100 is easily achievable with a few tens or hundreds of CPUs for this system size, confirming the broad applicability of HTMPC.

5. Showcases

In order to demonstrate the versatility of our code, we present the results of a plugin which is capable of simulating passive and active colloidal particles in solution. The rigid-body translational dynamics of colloidal particles is governed by Newton's equations of motion. For their rotational motion, we employ the quaternion-based formulation, which allows the simulation studies of spherical and non-spherical colloids [51] as well as chiral particles [54] in a unified framework. Fluid-colloid interactions imply a linear and angular momentum transfer by solid-body interactions. We note that this momentum exchange not only affects the MPC fluid particles, but also the colloidal particles. Moreover, for active particles with surface actuation, the slip velocity \mathbf{u}_{slip} at their surface is typically non-zero.

The implementation of colloids requires a plugin that contains multiple features described above – such as angular-momentum conservation for proper hydrodynamics, and additional routines for boundary conditions corresponding to fluid-colloid collisions and for processing of phantom particles, in addition to simulating colloid-colloid interactions. The system can be distributed over an arbitrary number of MPI ranks and all interactions between colloids and fluid or phantom particles are processed locally within each subdomain.

5.1. Colloidal suspensions at thermal equilibrium

The parallel implementation for colloidal particles in a fluid can of course not be as efficient as for pure fluids, as the interactions with the fluid need to be combined and synchronized for all N_{col} colloids in the system. To examine the scalability quantitatively, we again perform benchmark simulations, now including colloidal particles. Here, we consider $N_{\text{col}} = 10^6$ spherical colloids with effective radius $R_{\text{eff}} = 3.25a$. Purely repulsive steric interaction between colloids is modeled via the

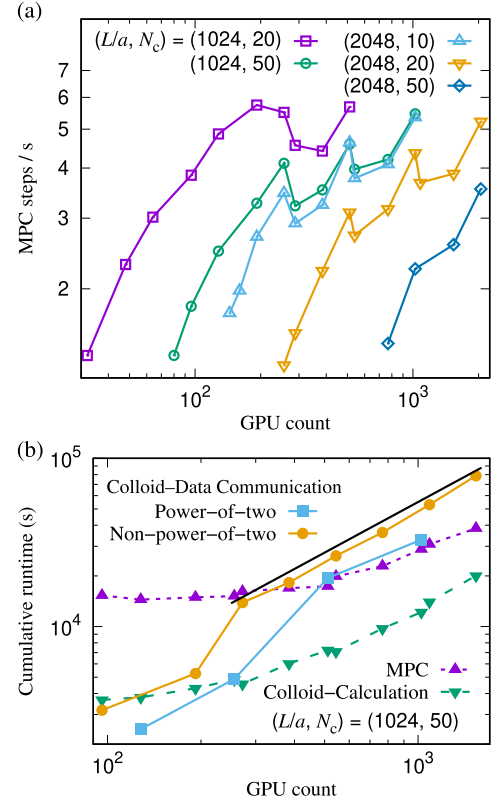


Fig. 7. (a) Scaling behaviors of the code for suspensions of 10^6 colloidal particles with varying system sizes and MPC density. (b) Cumulative runtimes of tasks. The black solid line shows a linear scaling.

separation-shifted Lennard-Jones potential, see SI, Sec. S-IV for details. For the MPC fluid, we use the collision time $h = 0.02a\sqrt{m/(k_B T)}$ and the rotation angle $\alpha = 130^\circ$ for various system sizes and MPC fluid densities.

As shown in Fig. 7(a), the simulated system shows good scaling behavior of run times with the number of GPUs, with two to five times speedup, depending on the system size, beyond which saturation in the speedup is observed. Examination of cumulative runtimes occupied by sub-tasks, shown in Fig. 7(b), indicates that the colloid-data communication load, i.e., the communication effort for colloid-data synchronization, depends linearly on the number of MPI ranks N_{rank} , i.e., roughly proportional to $N_{\text{col}} \times N_{\text{rank}}$, and therefore becomes dominant over the efficiently distributed workload for MPC and the actual colloid calculations, which indeed exhibit sub-linear scaling. The runtime

for colloid-data communication depends sensitively on the number of MPI ranks N_{rank} [84], where powers-of-two are favorable. A significant performance drop for non-power-of-two GPU counts is obtained as illustrated in Fig. 7(a).

5.2. Active systems of self-propelling squirmers

As an example of active (non-equilibrium) systems, we implement a system of spherical squirmers, which model the self-propulsion of microswimmers via active surface actuation [51,85–88]. Specifically, we consider neutral squirmers with the swimming speed u_0 , whose non-zero surface slip velocity is

$$\mathbf{u}_{\text{slip}}^b = \frac{3}{2}u_0 \sin \theta \mathbf{e}_\theta, \quad (9)$$

where θ denotes the polar angle in the body-fixed frame, and \mathbf{e}_θ the unit vector in the θ direction. Squirmer steric repulsion is again modeled via the Lennard-Jones potential of Eq. (S3) in SI. We also note that the amount of additional calculations for the squirmer activity is negligible and the corresponding scaling behavior (results not shown) is almost identical to that of passive colloidal suspensions presented in Fig. 7.

We first consider a dry system (no hydrodynamic interactions) of spherical active Brownian particles (ABPs) as a reference, whose dynamics is characterized by the emergence of a motility-induced phase separation (MIPS) [89,90] for strong self-propulsion, or high Péclet number. Here, we use ABPs of radius $R_c = 3.25a$ and Péclet number $\text{Pe} = u_0/(2R_c D_R) = 128$ with the rotational diffusion coefficient D_R and self-propulsion speed u_0 , and the linear system size of $L = 256a$ with periodic boundary conditions in all three directions, see SI, Sec. S-V for further simulation details. The volume packing fraction is varied from $\rho_{\text{col}} = 0.15$ to 0.50 , which correspond to $N_{\text{col}} = 17,408$ to $N_{\text{col}} = 57,600$. The distribution of local particle densities is obtained via Voronoi tessellation [91] as described in Sec. S-VIA of SI. As shown in Fig. 8(a), the ABP system exhibits MIPS, featuring a metastable liquid-gas coexistence for high packing fractions $\rho_{\text{col}} \gtrsim 0.35$, in line with the previous study [89,92]. The peak densities for the liquid and gas phases are found to be $\rho_{\text{liquid}} \approx 0.6$ and $\rho_{\text{gas}} \approx 0.1$, respectively, also in good agreement with the previous results [92].

We then turn to squirmer systems with the equivalent Péclet number, colloid densities, system sizes, and boundary conditions, to examine hydrodynamic effects on MIPS. Specifically, we use a cubic system of linear size $L = 256a$, and the MPC fluid density (particles per collision cell) $\langle N_c \rangle = 50$, which, together with the collision time $h = 0.02a\sqrt{m/(k_B T)}$ and the rotation angle $\alpha = 130^\circ$, yields the fluid viscosity of $\eta = 111.3\sqrt{mk_B T}/a^2$ [32,54,68]. For squirmers, we consider the radius $R_c = 3a$ or $R_{\text{eff}} = 3.25a$ (see Sec. S-IV, SI), and $u_0 = 0.01152\sqrt{k_B T/m}$. This parameter set results in the rotational diffusion coefficient $D_R = 1.5 \times 10^{-5}\sqrt{k_B T/m}/a$, Reynolds number $\text{Re} = 0.031$, and Péclet number $\text{Pe} = 128$ [54].

As shown in Fig. 8(b), density profiles of hydrodynamic systems consisting of neutral squirmers always exhibit a single peak around the corresponding global packing fraction for the examined range of densities, in sharp contrast to dry ABP systems, indicating that hydrodynamic interactions suppress the motility-induced phase separation, consistent with previous results [76]. We also employ a larger system size of $L = 512a$ with $N_{\text{col}} = 325,424$ and $470,596$, corresponding to $\rho_{\text{col}} \approx 0.35$ and 0.50 , respectively, to check for finite-size effects. Resultant density profiles are essentially identical to those of $L = 256a$ in both cases of ABPs and squirmers, confirming that finite-size effects are not significant. We also present the velocity and vorticity fields of the MPC fluid in Fig. 8(c). We observe the presence of many localized, strong vortical structures evenly distributed across the system, with the typical vorticity $\omega \approx u_0/R_c$. Corresponding squirmer configurations are also provided in Sec. S-VIB of SI.

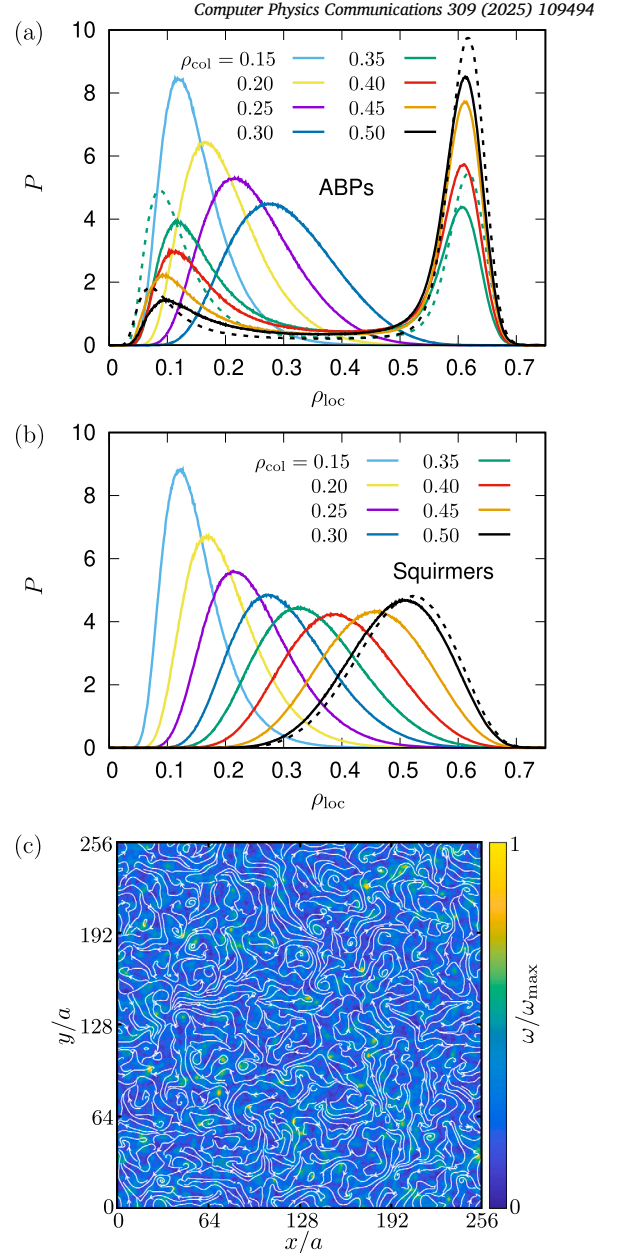


Fig. 8. Local density profiles for (a) ABPs, and (b) squirmers, obtained from Voronoi construction, adjusted by the volume of each Voronoi cells, see SI, Sec. S-VIA for details. In both cases, $\text{Pe} = 128$. We use global densities of $\rho_{\text{col}} = 0.15, 0.2, 0.25, 0.3, 0.35, 0.4, 0.45$, and 0.5 for $L = 256a$ (solid lines), and $\rho_{\text{col}} = 0.35$, and 0.5 for $L = 512a$ (dashed lines), respectively, as indicated. (c) A cross-section of the fluid velocity field (white arrow lines) and the magnitude of vorticity (heat map) are shown for squirmer suspension at $\rho_{\text{col}} = 0.35$. Here, a slice of the thickness $6a = 2R_c$ is used and $\omega_{\text{max}} \approx 1.5 \frac{u_0}{R_c} \gg D_R$.

6. Discussion

We have presented a highly parallelized implementation of multiparticle collision dynamics for hydrodynamic simulations. The MPC algorithm is well suited for parallelization on at least two layers: data access is local and there are no long-range interactions. Therefore, the system can be domain-decomposed and distributed with little communication overhead between subdomains. Additionally, many of the operations performed on the particles are independent. Therefore, the work done in each subdomain can be further parallelized. However, we note that communication is usually much more expensive than calculation, especially for distributed systems utilizing GPUs. Nevertheless, we have

confirmed overall good scaling behaviors up to more than a trillion MPC fluid particles.

Several features of the current code deviate from the previous implementations [25,26]. The major differences can be summarized as follows: (i) Our new implementation is plugin-based. (ii) The code provides a reproducibility option. (iii) We have abandoned the cell-list algorithm. Instead, our new version uses atomic operation directly. (iv) Every collision cell is uniquely assigned to one subdomain to avoid the necessity of partially processing additional cells in halo areas between subdomains and combining the results. (v) Instead of double buffering, we consider different variants of a memory-saving algorithm during reordering, which significantly reduces the memory footprint.

Lastly, as a showcase, we have considered the emergence of motility-induced phase separation (MIPS) in a system of self-propelled squirmers with hydrodynamic interactions. In dry systems of active Brownian particles in three dimensions, MIPS occurs for $Pe \gtrsim 2\text{--}30$ [89,90,92]. Meanwhile, for squirmer suspensions in two dimensions, MPC simulations demonstrate that hydrodynamic interactions significantly affect the collective behavior of active particles, and completely suppress phase separation [76]. However, direct simulation for three-dimensional systems was not achieved so far, most probably due to the computational complexity. Therefore, the results presented here clearly show that our newly implemented code is indeed very efficient, versatile and useful, facilitating investigations of large system sizes that could not be achieved so far.

CRedit authorship contribution statement

Elmar Westphal: Writing – original draft, Visualization, Validation, Software, Methodology, Investigation, Formal analysis, Data curation. **Segun Goh:** Writing – original draft, Visualization, Validation, Investigation, Formal analysis, Data curation. **Roland G. Winkler:** Writing – review & editing, Supervision, Project administration, Methodology, Conceptualization. **Gerhard Gompper:** Writing – review & editing, Supervision, Resources, Project administration, Methodology, Funding acquisition, Conceptualization.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgement

The authors gratefully acknowledge the Gauss Centre for Supercomputing e.V. (www.gauss-centre.eu) for funding this project by providing computing time through the John von Neumann Institute for Computing (NIC) on the GCS Supercomputer JUWELS at Jülich Supercomputing Centre (JSC).

Appendix A. Supplementary material

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.cpc.2024.109494>.

Data availability

Source code and documentation available at <https://go.fzj.de/HTMPC>.

References

- [1] G. Gompper, T. Ihle, D.M. Kroll, R.G. Winkler, Multi-particle collision dynamics: a particle-based mesoscale simulation approach to the hydrodynamics of complex fluids, *Adv. Polym. Sci.* 221 (2009) 1, https://doi.org/10.1007/978-3-540-87706-6_1.
- [2] J.K. Dhont, R.G. Winkler, G. Nägele, D. Richter, G. Gompper, *Soft Matter - From Synthetic to Biological Materials: Lecture Notes of the 39th Spring School 2008*; This Spring School was organized by the Institute of Solid State Research in the Research Centre Jülich on 3-14 March, 2008, No. PreJuSER-510, *Stremethoden*, 2008.
- [3] M. Doi, S.F. Edwards, *The Theory of Polymer Dynamics*, Clarendon Press, Oxford, 1986.
- [4] C.-C. Huang, R.G. Winkler, G. Sutmann, G. Gompper, Semidilute polymer solutions at equilibrium and under shear flow, *Macromolecules* 43 (2010) 10107, <https://doi.org/10.1021/ma101836x>.
- [5] M. Theers, E. Westphal, G. Gompper, R.G. Winkler, From local to hydrodynamic friction in Brownian motion: a multiparticle collision dynamics simulation study, *Phys. Rev. E* 93 (2016) 032604, <https://doi.org/10.1103/PhysRevE.93.032604>.
- [6] J.K.G. Dhont, *An Introduction to Dynamics of Colloids*, Elsevier, Amsterdam, 1996.
- [7] J.T. Padding, A.A. Louis, Hydrodynamic interactions and Brownian forces in colloidal suspensions: coarse-graining over time and length scales, *Phys. Rev. E* 74 (2006) 031402.
- [8] R. Chelakkot, R.G. Winkler, G. Gompper, Flow-induced helical coiling of semiflexible polymers in structured microchannels, *Phys. Rev. Lett.* 109 (2012) 178101.
- [9] K.Y. Wan, Life through the fluid dynamics lens, *Nat. Phys.* 19 (2023) 1744–1745, <https://doi.org/10.1038/s41567-023-02299-7>.
- [10] H. Noguchi, G. Gompper, Shape transitions of fluid vesicles and red blood cells in capillary flow, *Proc. Natl. Acad. Sci. USA* 102 (2005) 14159–14164, <https://doi.org/10.1073/pnas.0504243102>.
- [11] D.A. Fedosov, B. Caswell, G.E. Karniadakis, A multiscale red blood cell model with accurate mechanics, rheology, and dynamics, *Biophys. J.* 98 (10) (2010) 2215–2225, <https://doi.org/10.1016/j.bpj.2010.02.002>.
- [12] E. Lauga, T.R. Powers, The hydrodynamics of swimming microorganisms, *Rep. Prog. Phys.* 72 (2009) 096601, <https://doi.org/10.1088/0034-4885/72/9/096601>.
- [13] G. Gompper, C. Bechinger, S. Herminghaus, R. Isele-Holder, U.B. Kaupp, H. Löwen, H. Stark, R.G. Winkler, Microswimmers—from single particle motion to collective behavior, *Eur. Phys. J. Spec. Top.* 225 (2016) 2061, <https://doi.org/10.1140/epjst/e2016-60095-3>.
- [14] J. Hu, A. Wysocki, R.G. Winkler, G. Gompper, Physical sensing of surface properties by microswimmers – directing bacterial motion via wall slip, *Sci. Rep.* 5 (2015) 9586, <https://doi.org/10.1038/srep09586>.
- [15] S.M. Mousavi, G. Gompper, R.G. Winkler, Wall entrapment of peritrichous bacteria: a mesoscale hydrodynamics simulation study, *Soft Matter* 16 (2020) 4866, <https://doi.org/10.1039/D0SM000571A>.
- [16] J. Elgeti, R.G. Winkler, G. Gompper, Physics of microswimmers—single particle motion and collective behavior: a review, *Rep. Prog. Phys.* 78 (2015) 056601, <https://doi.org/10.1088/0034-4885/78/5/056601>.
- [17] P.J. Hoogerbrugge, J.M.V.A. Koelman, Simulating microscopic hydrodynamics phenomena with dissipative particle dynamics, *Europhys. Lett.* 19 (1992) 155, <https://doi.org/10.1209/0295-5075/19/3/001>.
- [18] P. Español, P.B. Warren, Perspective: dissipative particle dynamics, *J. Chem. Phys.* 146 (15) (2017) 150901, <https://doi.org/10.1063/1.4979514>.
- [19] G.R. McNamara, G. Zanetti, Use of the Boltzmann equation to simulate lattice-gas automata, *Phys. Rev. Lett.* 61 (1988) 2332, <https://doi.org/10.1103/PhysRevLett.61.2332>.
- [20] S. Succi, *The Lattice Boltzmann Equation: for Fluid Dynamics and Beyond*, Oxford University Press, 2001.
- [21] B. Dünweg, A.C. Ladd, Lattice Boltzmann simulations of soft matter systems, *Adv. Polym. Sci.* 221 (2009) 89, https://doi.org/10.1007/978-3-540-87706-6_2.
- [22] G.A. Bird, *Molecular Gas Dynamics and the Direct Simulation of Gas Flows*, Oxford University Press, Oxford, 1994.
- [23] S.J. Plimpton, S.G. Moore, A. Borner, A.K. Stagg, T.P. Koehler, J.R. Torczynski, M.A. Gallis, Direct simulation Monte Carlo on petaflop supercomputers and beyond, *Phys. Fluids* 31 (8) (2019) 086101, <https://doi.org/10.1063/1.5108534>.
- [24] R. Kapral, Multiparticle collision dynamics: simulations of complex systems on mesoscale, *Adv. Chem. Phys.* 140 (2008) 89, <https://doi.org/10.1002/9780470371572.ch2>.
- [25] E. Westphal, S.P. Singh, C.-C. Huang, G. Gompper, R.G. Winkler, Multiparticle collision dynamics: GPU accelerated particle-based mesoscale hydrodynamic simulations, *Comput. Phys. Commun.* 185 (2014) 495, <https://doi.org/10.1016/j.cpc.2013.10.004>.
- [26] M.P. Howard, A.Z. Panagiotopoulos, A. Nikoubashman, Efficient mesoscale hydrodynamics: multiparticle collision dynamics with massively parallel GPU acceleration, *Comput. Phys. Commun.* 230 (2018) 10, <https://doi.org/10.1016/j.cpc.2018.04.009>.
- [27] C.-C. Huang, G. Gompper, R.G. Winkler, Hydrodynamic correlations in multiparticle collision dynamics fluids, *Phys. Rev. E* 86 (2012) 056711, <https://doi.org/10.1103/PhysRevE.86.056711>.
- [28] A. Malevanets, R. Kapral, Mesoscopic model for solvent dynamics, *J. Chem. Phys.* 110 (1999) 8605, <https://doi.org/10.1063/1.478857>.
- [29] T. Ihle, D.M. Kroll, Stochastic rotation dynamics I: formalism, Galilean invariance, Green-Kubo relations, *Phys. Rev. E* 67 (2003) 066705, <https://doi.org/10.1103/PhysRevE.67.066705>.
- [30] T. Ihle, D.M. Kroll, Stochastic rotation dynamics II: transport coefficients, numerics, long time tails, *Phys. Rev. E* 67 (2003) 066706, <https://doi.org/10.1103/PhysRevE.67.066706>.

- [31] N. Kikuchi, C.M. Pooley, J.F. Ryder, J.M. Yeomans, Transport coefficients of a mesoscopic fluid dynamics model, *J. Chem. Phys.* 119 (2003) 6388–6395, <https://doi.org/10.1063/1.1603721>.
- [32] H. Noguchi, G. Gompper, Transport coefficients of off-lattice mesoscale-hydrodynamics simulation techniques, *Phys. Rev. E* 78 (2008) 016706, <https://doi.org/10.1103/PhysRevE.78.016706>.
- [33] R.G. Winkler, C.-C. Huang, Stress tensors of multiparticle collision dynamics fluids, *J. Chem. Phys.* 130 (2009) 074907, <https://doi.org/10.1063/1.3077860>.
- [34] J.T. Padding, A.A. Louis, Hydrodynamic and Brownian fluctuations in sedimenting suspensions, *Phys. Rev. Lett.* 93 (2004) 220601, <https://doi.org/10.1103/PhysRevLett.93.220601>.
- [35] M. Hecht, J. Harting, M. Bier, J. Reinshagen, H.J. Herrmann, Shear viscosity of claylike colloids in computer simulations and experiments, *Phys. Rev. E* 74 (2006) 021403, <https://doi.org/10.1103/PhysRevE.74.021403>.
- [36] S. Frank, R.G. Winkler, Polyelectrolyte electrophoresis: field effects and hydrodynamic interactions, *Europhys. Lett.* 83 (2008) 38004, <https://doi.org/10.1209/0295-5075/83/38004>.
- [37] S. Mandal, M.G. Mazza, Multiparticle collision dynamics for tensorial nematodynamics, *Phys. Rev. E* 99 (2019) 063319, <https://doi.org/10.1103/PhysRevE.99.063319>.
- [38] A. Malevanets, R. Kapral, Solute molecular dynamics in a mesoscopic solvent, *J. Chem. Phys.* 112 (2000) 7260, <https://doi.org/10.1063/1.481289>.
- [39] K. Mussawisade, M. Ripoll, R.G. Winkler, G. Gompper, Dynamics of polymers in a particle based mesoscopic solvent, *J. Chem. Phys.* 123 (2005) 144905, <https://doi.org/10.1063/1.2041527>.
- [40] I. Ali, D. Marenduzzo, J.M. Yeomans, Polymer packaging and ejection in viral capsids: shape matters, *Phys. Rev. Lett.* 96 (2006) 208102, <https://doi.org/10.1103/PhysRevLett.96.208102>.
- [41] A. Nikoubashman, C.N. Likos, Flow-induced polymer translocation through narrow and patterned channels, *J. Chem. Phys.* 133 (2010) 074901, <https://doi.org/10.1063/1.3466918>.
- [42] C.C. Huang, G. Gompper, R.G. Winkler, Effect of hydrodynamic correlations on the dynamics of polymers in dilute solution, *J. Chem. Phys.* 138 (2013) 144902, <https://doi.org/10.1063/1.4799877>.
- [43] Q.-S. Chen, A. Patelli, H. Chaté, Y.-Q. Ma, X.-Q. Shi, Fore-aft asymmetric flocking, *Phys. Rev. E* 96 (2017) 020601, <https://doi.org/10.1103/PhysRevE.96.020601>.
- [44] A. Nikoubashman, M.P. Howard, Equilibrium dynamics and shear rheology of semiflexible polymers in solution, *Macromolecules* 50 (2017) 8279–8289, <https://doi.org/10.1021/acs.macromol.7b01876>.
- [45] M. Liebetreu, C.N. Likos, Hydrodynamic inflation of ring polymers under shear, *Commun. Mater.* 1 (2020) 4, <https://doi.org/10.1038/s43246-019-0006-5>.
- [46] D.S. Devarajan, S. Rekhi, A. Nikoubashman, Y.C. Kim, M.P. Howard, J. Mittal, Effect of charge distribution on the dynamics of polyampholytic disordered proteins, *Macromolecules* 55 (2022) 8987–8997, <https://doi.org/10.1021/acs.macromol.2c01390>.
- [47] Z. Wang, Z.-G. Wang, A.-C. Shi, Y. Lu, L. An, Behaviors of a polymer chain in channels: from zimm to rouse dynamics, *Macromolecules* 56 (6) (2023) 2447–2453, <https://doi.org/10.1021/acs.macromol.3c00013>.
- [48] P. Ilg, Simulating the flow of interacting ferrofluids with multiparticle collision dynamics, *Phys. Rev. E* 106 (2022) 064605, <https://doi.org/10.1103/PhysRevE.106.064605>.
- [49] R. Wang, C. Feng, Z. Zhang, C. Shao, J. Du, What quantity of charge on the nanoparticle can result in a hybrid morphology of the nanofluid and a higher thermal conductivity?, *Powder Technol.* 422 (2023) 118443, <https://doi.org/10.1016/j.powtec.2023.118443>.
- [50] S. Thakur, R. Kapral, Collective dynamics of self-propelled sphere-dimer motors, *Phys. Rev. E* 85 (2012) 026121, <https://doi.org/10.1103/PhysRevE.85.026121>.
- [51] M. Theers, E. Westphal, G. Gompper, R.G. Winkler, Modeling a spheroidal microswimmer and cooperative swimming in a narrow slit, *Soft Matter* 12 (2016) 7372, <https://doi.org/10.1039/C6SM01424K>.
- [52] A.W. Zantop, H. Stark, Emergent collective dynamics of pusher and puller squirmers: swarming, clustering, and turbulence, *Soft Matter* 18 (2022) 6179, <https://doi.org/10.1039/D2SM00449F>.
- [53] K. Qi, E. Westphal, G. Gompper, R.G. Winkler, Emergence of active turbulence in microswimmer suspensions due to active hydrodynamic stress and volume exclusion, *Commun. Phys.* 5 (2022) 49, <https://doi.org/10.1038/s42005-022-00820-7>.
- [54] S. Goh, R.G. Winkler, G. Gompper, Hydrodynamic pursuit by cognitive self-steering microswimmers, *Commun. Phys.* 6 (2023) 310, <https://doi.org/10.1038/s42005-023-01432-5>.
- [55] J. Macías-Durán, V. Duarte-Alaniz, H. Híjar, Active nematic liquid crystals simulated by particle-based mesoscopic methods, *Soft Matter* 19 (2023) 8052–8069, <https://doi.org/10.1039/D3SM00481C>.
- [56] N. Jain, S. Thakur, Collapse dynamics of chemically active flexible polymer, *Macromolecules* 55 (7) (2022) 2375–2382, <https://doi.org/10.1021/acs.macromol.1c02502>.
- [57] J. Clopés Llahí, A. Martín-Gómez, G. Gompper, R.G. Winkler, Simulating wet active polymers by multiparticle collision dynamics, *Phys. Rev. E* 105 (2022) 015310, <https://doi.org/10.1103/PhysRevE.105.015310>.
- [58] J. Hu, M. Yang, G. Gompper, R.G. Winkler, Modelling the mechanics and hydrodynamics of swimming *E. coli*, *Soft Matter* 11 (2015) 7867, <https://doi.org/10.1039/C5SM01678A>.
- [59] T. Eisenstecken, J. Hu, R.G. Winkler, Bacterial swarmer cells in confinement: a mesoscale hydrodynamic simulation study, *Soft Matter* 12 (2016) 8316, <https://doi.org/10.1039/C6SM01532H>.
- [60] L. Ning, X. Lou, Q. Ma, Y. Yang, N. Luo, K. Chen, F. Meng, X. Zhou, M. Yang, Y. Peng, Hydrodynamics-induced long-range attraction between plates in bacterial suspensions, *Phys. Rev. Lett.* 131 (2023) 158301, <https://doi.org/10.1103/PhysRevLett.131.158301>.
- [61] J. Elgeti, U.B. Kaupp, G. Gompper, Hydrodynamics of sperm cells near surfaces, *Biophys. J.* 99 (2010) 1018, <https://doi.org/10.1016/j.bpj.2010.05.015>.
- [62] T. Chinnamasamy, J.L. Kingsley, F. Inci, P.J. Turek, M.P. Rosen, B. Behr, E. Tüzel, U. Demirci, Guidance and self-sorting of active swimmers: 3d periodic arrays increase persistence length of human sperm selecting for the fittest, *Adv. Sci.* 5 (2) (2018) 1700531, <https://doi.org/10.1002/adv.201700531>.
- [63] S. Rode, J. Elgeti, G. Gompper, Sperm motility in modulated microchannels, *New J. Phys.* 21 (2019) 013016, <https://doi.org/10.1088/1367-2630/aaf544>.
- [64] N. Heddergott, T. Krüger, S.B. Babu, A. Wei, E. Stellamanns, S. Uppaluri, T. Pfohl, H. Stark, M. Engstler, Trypanosome motion represents an adaptation to the crowded environment of the vertebrate bloodstream, *PLoS Pathog.* 8 (11) (2012) 1–17, <https://doi.org/10.1371/journal.ppat.1003023>.
- [65] C. Lansche, A.K. Dasanna, K. Quadt, B. Fröhlich, D. Missirlis, M. Tétard, B. Gamain, B. Buchholz, C.P. Sanchez, M. Tanaka, U.S. Schwarz, M. Lanzer, The sickle cell trait affects contact dynamics and endothelial cell activation in Plasmodium falciparum-infected erythrocytes, *Commun. Biol.* 1 (2018) 211, <https://doi.org/10.1038/s42003-018-0223-3>.
- [66] T. Ihle, D.M. Kroll, Stochastic rotation dynamics: a Galilean-invariant mesoscopic model for fluid flow, *Phys. Rev. E* 63 (2001) 020201(R), <https://doi.org/10.1103/PhysRevE.63.020201>.
- [67] I.O. Götz, H. Noguchi, G. Gompper, Relevance of angular momentum conservation in mesoscale hydrodynamics simulations, *Phys. Rev. E* 76 (2007) 046705, <https://doi.org/10.1103/PhysRevE.76.046705>.
- [68] M. Theers, R.G. Winkler, Bulk viscosity of multiparticle collision dynamics fluids, *Phys. Rev. E* 91 (2015) 033309, <https://doi.org/10.1103/PhysRevE.91.033309>.
- [69] M. Yang, M. Theers, J. Hu, G. Gompper, R.G. Winkler, M. Ripoll, Effect of angular momentum conservation on hydrodynamic simulations of colloids, *Phys. Rev. E* 92 (2015) 013301, <https://doi.org/10.1103/PhysRevE.92.013301>.
- [70] H. Noguchi, N. Kikuchi, G. Gompper, Particle-based mesoscale hydrodynamic techniques, *Europhys. Lett.* 78 (2007) 10005.
- [71] C.-C. Huang, A. Chatterji, G. Sutmann, G. Gompper, R.G. Winkler, Cell-level canonical sampling by velocity scaling for multiparticle collision dynamics simulations, *J. Comput. Phys.* 229 (2010) 168, <https://doi.org/10.1016/j.jcp.2009.09.024>.
- [72] A. Lamura, G. Gompper, T. Ihle, D.M. Kroll, Multiparticle collision dynamics: flow around a circular and a square cylinder, *Europhys. Lett.* 56 (2001) 319–325, <https://doi.org/10.1209/epl/i2001-00522-9>.
- [73] L. Cannavacciuolo, R.G. Winkler, G. Gompper, Mesoscale simulation of polymer dynamics in microchannel flows, *Europhys. Lett.* 83 (2008) 34007, <https://doi.org/10.1209/0295-5075/83/34007>.
- [74] C.-C. Huang, A. Varghese, G. Gompper, R.G. Winkler, Thermostat for nonequilibrium multiparticle-collision-dynamics simulations, *Phys. Rev. E* 91 (2015) 013310, <https://doi.org/10.1103/PhysRevE.91.013310>.
- [75] M. Ripoll, K. Mussawisade, R.G. Winkler, G. Gompper, Dynamic regimes of fluids simulated by multi-particle-collision dynamics, *Phys. Rev. E* 72 (2005) 016701.
- [76] M. Theers, E. Westphal, K. Qi, R.G. Winkler, G. Gompper, Clustering of microswimmers: interplay of shape and hydrodynamics, *Soft Matter* 14 (2018) 8590, <https://doi.org/10.1039/C8SM01390J>.
- [77] D. Goldberg, What every computer scientist should know about floating-point arithmetic, *ACM Comput. Surv.* 23 (1) (1991) 5–48, <https://doi.org/10.1145/103162.103163>.
- [78] O. Villa, D. Chavarria-Miranda, V. Gurumoorathi, A. Márquez, S. Krishnamoorthy, Effects of floating-point non-associativity on numerical computations on massively multithreaded systems, in: *Proceedings of Cray User Group Meeting (CUG)*, vol. 3, 2009.
- [79] S. Le Grand, A.W. Götz, R.C. Walker, Spfp: speed without compromise—a mixed precision model for gpu accelerated molecular dynamics simulations, *Comput. Phys. Commun.* 184 (2) (2013) 374–380, <https://doi.org/10.1016/j.cpc.2012.09.022>.
- [80] W. Dietz, P. Li, J. Regehr, V. Adve, Understanding integer overflow in C/C++, *ACM Trans. Softw. Eng. Methodol.* 25 (1) (dec 2015), <https://doi.org/10.1145/2743019>.
- [81] E. Westphal, Voting and shuffling to optimize atomic operations, *Technical Blog (GPU Pro Tip)*, NVIDIA Developer, 2015, <https://developer.nvidia.com/blog/voting-and-shuffling-optimize-atomic-operations>.
- [82] G. Marsaglia, Random number generators, *J. Mod. Appl. Stat. Methods* 2 (2003) 2, <https://doi.org/10.22237/jmasm/1051747320>.
- [83] Jülich Supercomputing Centre, JUWELS cluster and booster: exascale pathfinder with modular supercomputing architecture at Jülich Supercomputing Centre, *J. Large-Scale Res. Facil.* 7 (2021) A138, <https://doi.org/10.17815/jlsrf-7-183>.
- [84] R. Rabenseifner, J.L. Träff, More efficient reduction algorithms for non-power-of-two number of processors in message-passing parallel systems, in: D. Kranzlmüller, P. Kacsuk, J. Dongarra (Eds.), *Recent Advances in Parallel Virtual Machine and Message Passing Interface*, Springer Berlin Heidelberg, Berlin, Heidelberg, 2004, pp. 36–46.
- [85] M.J. Lighthill, On the squirming motion of nearly spherical deformable bodies through liquids at very small Reynolds numbers, *Commun. Pure Appl. Math.* 5 (1952) 109, <https://doi.org/10.1002/cpa.3160050201>.

- [86] J.R. Blake, A spherical envelope approach to ciliary propulsion, *J. Fluid Mech.* 46 (1971) 199, <https://doi.org/10.1017/S002211207100048X>.
- [87] I.O. Götze, G. Gompper, Mesoscale simulations of hydrodynamic squirmer interactions, *Phys. Rev. E* 82 (2010) 041921, <https://doi.org/10.1103/PhysRevE.82.041921>.
- [88] O.S. Pak, E. Lauga, Generalized squirming motion of a sphere, *J. Eng. Math.* 88 (1) (2014) 1, <https://doi.org/10.1007/s10665-014-9690-9>.
- [89] A. Wysocki, R.G. Winkler, G. Gompper, Cooperative motion of active Brownian spheres in three-dimensional dense suspensions, *Europhys. Lett.* 105 (2014) 48004, <https://doi.org/10.1209/0295-5075/105/48004>.
- [90] J. Stenhammar, D. Marenduzzo, R.J. Allen, M.E. Cates, Phase behaviour of active Brownian particles: the role of dimensionality, *Soft Matter* 10 (2014) 1489, <https://doi.org/10.1039/C3SM52813H>.
- [91] C.H. Rycroft, VORO++: a three-dimensional Voronoi cell library in C++, *Chaos* 19 (2009) 041111, <https://doi.org/10.1063/1.3215722>.
- [92] A.K. Omar, K. Klymko, T. GrandPre, P.L. Geissler, Phase diagram of active Brownian spheres: crystallization and the metastability of motility-induced phase separation, *Phys. Rev. Lett.* 126 (2021) 188002, <https://doi.org/10.1103/PhysRevLett.126.188002>.