

Efficient point-based simulation of four-way coupled particles in turbulence at high number density

Xander M. de Wit,¹ Rudie P. J. Kunnen,¹ Herman J. H. Clercx,¹ and Federico Toschi^{1,2}

¹*Fluids and Flows group and J.M. Burgers Center for Fluid Mechanics,
Department of Applied Physics and Science Education,
Eindhoven University of Technology, 5600 MB Eindhoven, Netherlands*

²*CNR-IAC, I-00185 Rome, Italy*

(Dated: April 17, 2024)

In many natural and industrial applications, turbulent flows encompass some form of dispersed particles. Although this type of multiphase turbulent flow is omnipresent, its numerical modeling has proven to be a remarkably challenging problem. Models that fully resolve the particle phase are computationally very expensive, strongly limiting the number of particles that can be considered in practice. This warrants the need for efficient reduced order modeling of the complex system of particles in turbulence that can handle high number densities of particles. Here, we present an efficient method for point-based simulation of particles in turbulence that are four-way coupled. In contrast with traditional one-way coupled simulations, where only the effect of the fluid phase on the particle phase is modeled, this method additionally captures the back-reaction of the particle phase on the fluid phase, as well as the interactions between particles themselves. We focus on the most challenging case of very light particles or bubbles, which show strong clustering in the high-vorticity regions of the fluid. This strong clustering poses numerical difficulties which are systematically treated in our work. Our method is valid in the limit of small particles with respect to the Kolmogorov scales of the flow and is able to handle very large number densities of particles. This method paves the way for comprehensive studies of the collective effect of small particles in fluid turbulence for a multitude of applications.

I. INTRODUCTION

Understanding the behavior of particles in turbulent flows is paramount in many fields, including engineering, meteorology, oceanography, and astrophysics. The delicate interplay between particle dynamics and the background turbulence encompasses a rich phenomenology [1–5]. To accurately describe this complex system, in general, one needs to consider the reciprocal interactions between fluid and particles, as well as the mutual interactions between particles. From the computational point of view, modeling this complete interplay turns out to be challenging and often computationally prohibitive.

Therefore, most numerical studies historically have focused on the so-called one-way coupling paradigm, where only the effect of turbulence on the particle dynamics is modeled but not vice-versa. This is a reasonable approximation for small particles in the very dilute limit and was shown to recover much of the phenomenology observed in experiments, such as preferential concentration, particle dispersion and clustering [6–9]. In some cases, however, one may be particularly interested in the back-reaction that the particle phase has on the underlying fluid phase. This requires a so-called two-way coupled simulation, where also this back-reaction is explicitly modeled. The most accurate way to model this back-reaction is by fully resolving the geometry of the particles and enforcing the no-slip condition at the particle-fluid interface, e.g. via approaches such as the Immersed Boundary Method [10–14]. Though possibly being very accurate, this method is also computationally very expensive and in practice, this limits the number of particles that can be considered to

around $\mathcal{O}(10 - 100)$ [14]. To study the collective effect of a large number of small particles dispersed in the turbulent flow, this is insufficient, warranting the need for more computationally efficient reduced order modeling.

In this work, we therefore consider instead a point-particle based model which includes the back-reaction force on the fluid and that can easily be extended to millions of particles [15–17]. To accurately handle substantial volume fractions of particles and/or to handle cases of strong particle clustering, however, we need to go one step further and model also the interactions between particles, bringing us in the realm of the so-called four-way coupled simulations.

We perform Direct Numerical Simulations (DNS) of the fluid phase with spatio-temporal velocity field $\mathbf{u}(\mathbf{x}, t)$ using a standard pseudospectral method that solves the incompressible Navier-Stokes equations

$$\frac{D\mathbf{u}}{Dt} = -\nabla p + \nu \Delta \mathbf{u} + \mathbf{f} + \mathbf{f}_p, \quad (1)$$

$$\nabla \cdot \mathbf{u} = 0, \quad (2)$$

with $D\mathbf{u}/Dt \equiv \partial\mathbf{u}/\partial t + \mathbf{u} \cdot \nabla \mathbf{u}$ denoting the material derivative. Here, p is the local pressure, ν denotes the kinematic viscosity of the fluid and \mathbf{f} is the forcing of the fluid, which we choose to be $\hat{\mathbf{f}}(\mathbf{k}) = \epsilon \hat{\mathbf{u}}(\mathbf{k}) / \sum_{k_f \leq |\mathbf{k}| < k_f + 1} |\hat{\mathbf{u}}(\mathbf{k})|^2$, acting on the low wavenumbers $k_f \leq |\mathbf{k}| < k_f + 1$, ensuring a constant energy injection rate ϵ [18]. The feedback force of the particle phase on the fluid phase is captured in \mathbf{f}_p , which is derived in Section II A. For time integration, we use a 2nd order Adam-Bashfort scheme, while the viscous term

is integrated exactly using integrating factors. For details on the Eulerian scheme, see Appendix A

To describe the particle phase, we resort to the Maxey-Riley equations of motion for dispersed particles [19]. This point-based model requires that the particle has a sufficiently small size to see a smooth velocity field, which in practice amounts to an $\mathcal{O}(1)$ multiple of the Kolmogorov scale η . We consider the limit where we can neglect gravity and for simplicity we also disregard the Basset history force that corrects for the unsteady boundary layer around the particle [19, 20]. We then retain only the Stokes drag term and the combined pressure gradient and added mass term, yielding

$$\frac{d\mathbf{V}}{dt} = \beta \frac{D\mathbf{u}(\mathbf{X}, t)}{Dt} - \frac{1}{\tau_p}(\mathbf{V} - \mathbf{u}(\mathbf{X}, t)), \quad (3)$$

$$\frac{d\mathbf{X}}{dt} = \mathbf{V} + (\text{collisions}). \quad (4)$$

Here \mathbf{V} and \mathbf{X} denote the velocity and position of the particle, respectively. Furthermore, the density ratio between the particle and the fluid is parameterized by $\beta = 3/(1 + 2\rho_p/\rho_f)$ with ρ_p and ρ_f the density of the particle and fluid phase, respectively. Finally, τ_p denotes the particle response time, which for spheres is given by $\tau_p = D^2/(12\beta\nu)$ with D the particle diameter. The latter is typically non-dimensionalized with respect to the Kolmogorov timescale of the fluid τ_η , yielding the Stokes number $St = \tau_p/\tau_\eta$. The fluid velocity and its material derivative is evaluated at the particle position using spline interpolation [21]. Note that in the integration of the particle position, we also take into account the particle-particle collisions, which is treated in Section III A, completing our four-way coupling method.

Arguably among the most striking phenomena of the behavior of particles in turbulence is the preferential concentration of particles that are heavier or lighter than the fluid, observed at moderate $St \sim \mathcal{O}(1)$. While heavy particles ($\beta \ll 1$) tend to preferentially concentrate in high-straining regions, light particles ($\beta \gg 1$) show the opposite behavior, strongly concentrating in high-vorticity regions of the turbulent flow [9]. While the method presented here is valid for small particles with any density ratio and/or St , we focus on bubbles $\beta = 3$ at moderate $St = 1$, which show the strongest clustering due to their preferential concentration in the vortex filaments of the turbulent flow [2]. As we shall show, this strong clustering poses additional difficulties to the four-way coupling scheme and thus acts as a *worst-case-scenario* benchmark to our method.

Section II treats the implementation of the particle-fluid interaction, while section III details the scheme for the particle-particle interaction. Finally, section IV covers an example the modulation of the turbulent energy spectrum under the influence of four-way coupled bubbles and conclusions are drawn in section V.

II. PARTICLE-FLUID INTERACTION

A. Momentum conservation: the delta-coupling

The fluid phase is resolved on an Eulerian grid using a pseudospectral method as laid out in Appendix A. To account for the back-reaction of the particle on the fluid, we need to consider the particle as a source of momentum. The particle phase then poses an additional term in the Navier-Stokes equations that follows from the conservation of momentum between the particle and fluid phase. The open question that remains is how the total momentum that is transferred from the particle to the fluid is distributed in space. In general, this depends on the size and shape of the particle and requires resolving the full particle as is done e.g. in the Immersed Boundary Method. In the limit where the particle is small with respect to the Kolmogorov length of the underlying turbulent flow, however, the spatial distribution of the back-coupling force conveniently reduces to a δ -distribution [15, 22]. This can be treated efficiently in the DNS.

To derive the expression for the two-way coupling, we start from the integral conservation of momentum in the fluid phase with volume \mathcal{V}_f with one submerged particle with volume \mathcal{V}_p and surface S_p [23]

$$\int_{\mathcal{V}_f} \rho_f \frac{D\mathbf{u}}{Dt} d\mathcal{V} = \int_{\mathcal{V}_f} \rho_f [-\nabla p + \nu \Delta \mathbf{u} + \mathbf{f}] d\mathcal{V} + \int_{S_p} \boldsymbol{\sigma} \cdot \mathbf{n} dS, \quad (5)$$

with \mathbf{n} the surface normal of the particle and $\boldsymbol{\sigma}$ the fluid stress. Then, for the particle phase, we can write momentum conservation as

$$\rho_p \mathcal{V}_p \frac{d\mathbf{V}}{dt} = - \int_{S_p} \boldsymbol{\sigma} \cdot \mathbf{n} dS. \quad (6)$$

When evaluated for the particle, this results in the Maxey-Riley equations of motion Eq. (3). For the fluid phase, however, in general, substituting Eq. (6) into Eq. (5) and extending the integral over the full domain $\mathcal{V} = \mathcal{V}_p + \mathcal{V}_f$ gives

$$\begin{aligned} \int_{\mathcal{V}} \rho_f \frac{D\mathbf{u}}{Dt} d\mathcal{V} &= \int_{\mathcal{V}} \rho_f [-\nabla p + \nu \Delta \mathbf{u} + \mathbf{f}] d\mathcal{V} \\ &+ \int_{\mathcal{V}} \left[\rho_f \frac{D\mathbf{u}}{Dt} - \rho_p \frac{d\mathbf{V}}{dt} \right] \mathcal{V}_p \delta(\mathbf{x} - \mathbf{X}(t)) d\mathcal{V}. \end{aligned} \quad (7)$$

The total two-way coupling force is then given by a sum over all particles i [23, 24]

$$\mathbf{f}_p = \sum_i^{N_p} \left[\frac{D\mathbf{u}}{Dt} - \frac{\rho_p}{\rho_f} \frac{d\mathbf{V}_i}{dt} \right] \mathcal{V}_p \delta(\mathbf{x} - \mathbf{X}_i(t)). \quad (8)$$

Note that the two-way coupling term not only reflects the momentum exchange (second term in the brackets), but also has an added mass term (first term in the brackets).

The two-way coupling force Eq. (8) can be computed by evaluating the Lagrangian particle acceleration and




FIG. 1. Panel (a): Energy spectra of the one-phase turbulent flow and the two-way coupled bubbly turbulent flow with uniform volume fraction of bubbles $\alpha_c = 0.1$. Panel (b): the two-way coupled energy spectrum compensated by the one-way energy spectrum reveals that the two-way coupling enhances the energy spectrum in the forcing and inertial range, while it attenuates the spectrum in the dissipative range. It furthermore shows that the two-way coupled system is in close agreement with a one-way system where the energy injection rate ϵ and viscosity ν are enhanced by a factor $1/(1 - \alpha_c)$, consistent with the theoretical prediction Eq. (11). The resolution is $N^3 = 128^3$ and the number of particles is $N_p = 2\,097\,152 (= 128^3)$. The shaded regions denote statistical errorbars.

the Eulerian material derivative, interpolated at the position of the particle. To evaluate the δ -function, practically, it is smoothed by extrapolating to the eight nearest grid cells around the particle center using linear volume-weighting, summing to $1/\mathcal{V}_c$ with \mathcal{V}_c the volume of one Eulerian computational cell [15–17]. This ensures that the δ -function integrates to unity. By summing over all particles, we can then obtain \mathbf{f}_p evaluated on the full Eulerian grid, after taking into account the communication at the boundaries of the computational processes. As a final step, \mathbf{f}_p is projected onto its divergence-free part to retain incompressibility. Semantically, this entails a redefinition of the pressure that compensates the divergent part.

B. Two-way coupling validation: homogeneous distribution of bubbles

As a point of validation, we can consider a homogeneous distribution of bubbles. For bubbles $\rho_p/\rho_f = 0$, so that only the added mass term in Eq. (8) remains. Note that in this case, the particle equation of motion vanishes from the back-reaction \mathbf{f}_p and only the spatial distribution of the bubbles becomes important. Indeed, in a continuum description, we obtain

$$\mathbf{f}_p = \alpha(\mathbf{x}, t) \frac{D\mathbf{u}}{Dt}, \quad (9)$$

with $\alpha(\mathbf{x}, t) \equiv \sum_i^{N_p} \mathcal{V}_p \delta(\mathbf{x} - \mathbf{X}_i(t))$ the local volume fraction of bubbles. This yields

$$\frac{D\mathbf{u}}{Dt} = -\nabla p + \nu \Delta \mathbf{u} + \mathbf{f} + \alpha(\mathbf{x}, t) \frac{D\mathbf{u}}{Dt}. \quad (10)$$

We can then analytically compute the effect of the two-way forcing on the underlying fluid if we enforce a homogeneous distribution of particles $\alpha(\mathbf{x}, t) = \alpha_c$. Then Eq. (10) reduces to

$$\frac{D\mathbf{u}}{Dt} = -\frac{1}{1 - \alpha_c} \nabla p + \frac{1}{1 - \alpha_c} \nu \Delta \mathbf{u} + \frac{1}{1 - \alpha_c} \mathbf{f}. \quad (11)$$

That is, solving the two-way coupled system should become equivalent to solving a one-phase system, where both the viscosity ν as well as the energy injection rate ϵ are enhanced by a factor $1/(1 - \alpha_c)$. The prefactor to the pressure gradient can be absorbed in a redefinition of the effective pressure and thus has no dynamical influence.

This correspondence is confirmed numerically as laid out in Fig. 1 by considering the turbulent kinetic energy spectrum, validating the implementation of our two-way coupling algorithm for the case where we enforce a homogeneous distribution of immobilized bubbles. Only in the highest wavenumber range, some quantitative differences between the two-way coupled system and the enhanced one-way coupled system can be observed. We attribute these discrepancies to discretization errors, since the two-way coupling involves interpolation and discrete time-stepping, while the viscous term is integrated exactly using integrating factors. Indeed, we confirmed that upon refining the spatiotemporal resolution, this remaining discrepancy at high wavenumbers vanishes (not shown).

We thus understand that the added mass effect of bubbles enhances the turbulence in the forcing and inertial range, while it simultaneously also enhances the viscous dissipation, attenuating the turbulence in the dissipative range. However, while we have shown this for the case of a uniform density of bubbles, in reality, due to prefer-

ential sampling of bubbles in vortex filaments, we should expect the back-reaction of bubbles to be highly localized in space, which can change the global influence of the bubbles on the underlying turbulence.

From the numerical point of view, this poses two additional problems in the pursuit to obtain reliable simulations of bubbly turbulence. For one, since bubbles sample the Kolmogorov-like scales of vortex filaments, the resulting high spatial dishomogeneity emphasizes the action of the back-reaction force in the high wavenumber range. This warrants more caution in treating the dealiasing of the modified velocity field, as will be discussed in the next section. The second problem concerns the stability of the two-way coupled system, as discussed after.

C. Dealiasing

After that the particle feedback force $\mathbf{f}_p(\mathbf{x})$ is computed in real space on the N^3 grid, it is transformed into Fourier space, yielding $\hat{\mathbf{f}}_p(\mathbf{k})$ on the N^3 Fourier grid with $\mathbf{k} \in [-(N/2 - 1), \dots, N/2]^3$. In a conventional pseudospectral algorithm with a 2/3-dealiasing rule, all Fourier space fields would be zeroed for all $|\mathbf{k}| > (2/3)N/2$ to avoid spurious aliasing contributions from the non-linear term.

However, in the case of light particles, due to their strong clustering, the particle feedback force becomes strongly localized and rough (large gradients), such that $\hat{\mathbf{f}}_p(\mathbf{k})$ has significant contributions in the high wavenumber range. Hence, truncating $\hat{\mathbf{f}}_p(\mathbf{k})$ using a 2/3-dealiasing rule would be a strong approximation, significantly decreasing the intensity of the two-way coupling, and it should thus be avoided. See Fig. 2 for an example of the consequences of 2/3-dealiasing on a delta-function.




FIG. 2. A 2/3-dealiasing rule applied to a delta-function in real space reduces the magnitude of the delta-peak and produces spurious oscillations in its vicinity.

Therefore, we propose a slightly different dealiasing approach where the non-linear term is computed on a larger

grid than the one used for other real space quantities such as the feedback force, avoiding the need for truncation of the latter. The scheme is laid out in Fig. 3.

The zero-padded Fourier transform from Fourier space to the real space convolution grid and vice-versa can be computed efficiently as a pruned FFT and by skipping 1D transforms that are identically zero. This is implemented using the P3DFFT package [25]. Note that the dealiasing approach suggested here raises the effective resolution of the simulation to $(3N/2)^3$ as compared to a N^3 resolution with a conventional 2/3-dealiasing rule.

D. Stability

The strong clustering of light particles not only poses a challenge due to the roughness that it introduces in the feedback force, but it can also compromise the stability of the integration of the two-way coupled system. This becomes apparent by considering Eq. (10). Recall that the explicit time derivative $\partial \mathbf{u} / \partial t$ with respect to which the integration is performed is contained in the material derivative. Now, Eq. (10) shows that due to the two-way coupling, the time derivative is fed back into the evolution equation itself. Practically, since the material derivative as sensed by the particles is evaluated on one timestep Δt prior to the current integration timestep, this effectively reduces Eq. (10) to

$$\frac{\partial \mathbf{u}(t)}{\partial t} \approx \alpha \frac{\partial \mathbf{u}(t - \Delta t)}{\partial t} + \mathcal{F}[\mathbf{u}], \quad (12)$$

with $\mathcal{F}[\mathbf{u}]$ capturing the other forces (pressure gradient, advective acceleration, viscosity and external force). This is a dynamical system of the type $b_n = \alpha b_{n-1}$, which becomes unstable if $\alpha > 1$, yielding exponentially growing solutions. This indicates that we should expect the two-way coupled system of bubbles to become unstable if the local volume fraction α exceeds unity. Physically, this can never occur if one enforces excluded volume interactions. This warrants the need for explicit treatment of the particle-particle interaction, indicating that in order to obtain reliable two-way coupled simulations of bubbly turbulence, one in fact needs to resort to four-way coupled simulations. This is treated in the next section III.

The particle-particle interaction that will be considered here, however, is an approximate method in order to be able to computationally handle large number densities of bubbles. As a consequence, not all excluded volume interactions are strictly enforced at all times. To ensure stability, we therefore need to explicitly enforce that the local volume fraction does not exceed unity by carefully clipping the particle feedback force in those cases. To that extent, we multiply the particle feedback force by a local clipping factor $c(\mathbf{x}, t)$ as

$$c(\mathbf{x}, t) = \begin{cases} 1 & \text{for } \alpha(\mathbf{x}, t) \leq \alpha_0, \\ \frac{\alpha_0}{\alpha(\mathbf{x}, t)} & \text{for } \alpha(\mathbf{x}, t) > \alpha_0. \end{cases} \quad (13)$$

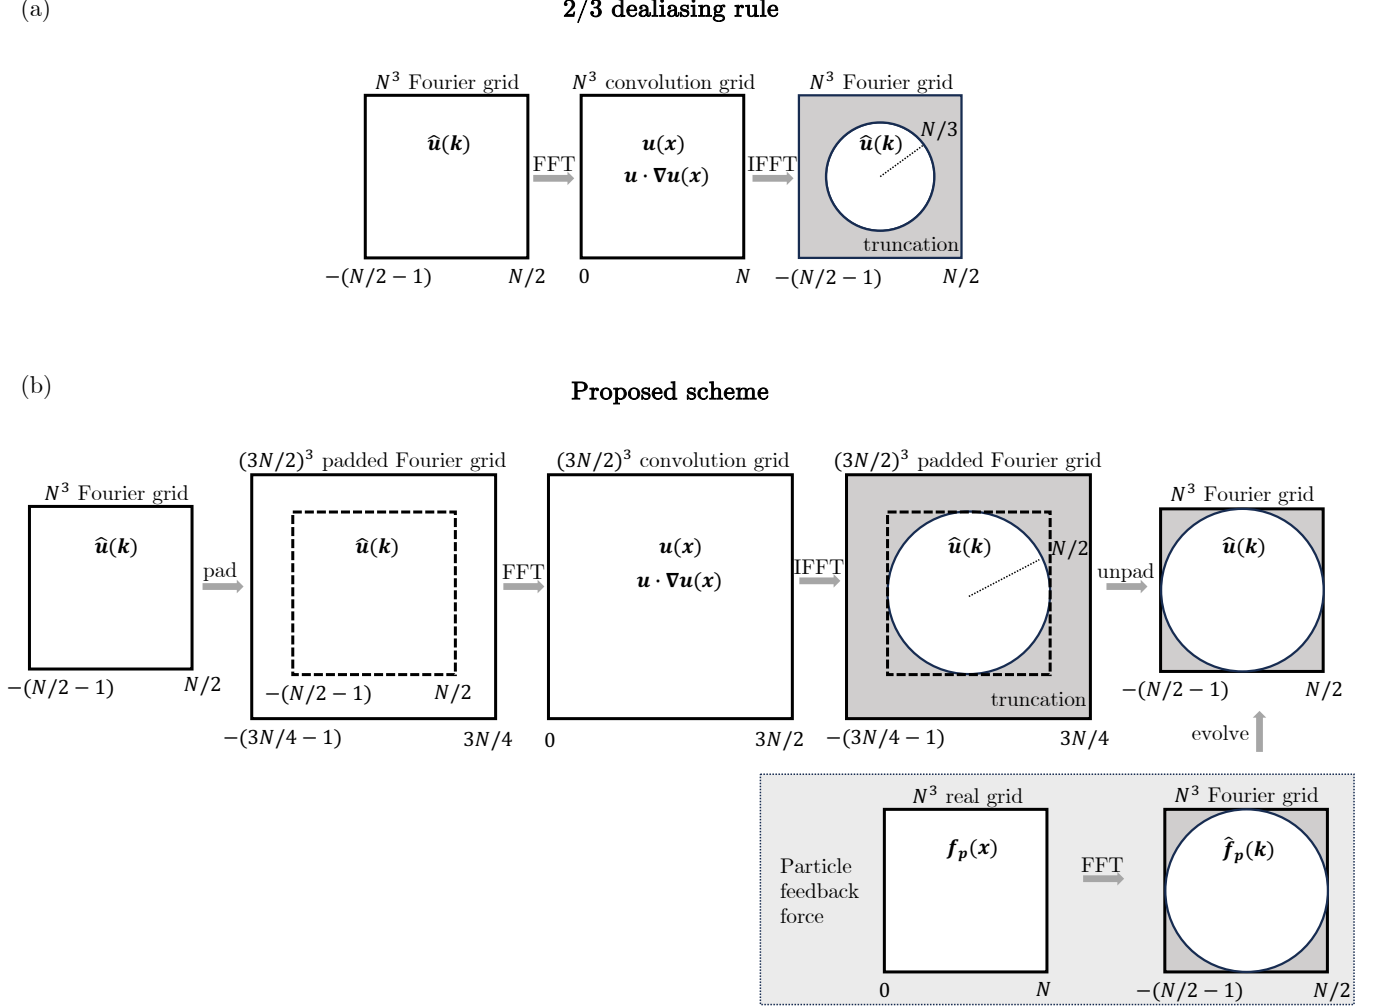


FIG. 3. A schematic of the proposed dealiasing routine. The schematic depicts the analog 2D approach, but a 3D setting is implied. Panel (a): the conventional 2/3 dealiasing rule, where all quantities are computed on a N^3 grid and the Fourier space quantities are truncated at $|\mathbf{k}| > N/3$. Panel (b): our suggested routine. The flow field is evolved on a N^3 Fourier grid. For computation of the non-linear term, the flow field is inflated and padded with zeros to make a $(3N/2)^3$ grid, which is transformed to real space, where the convolution is carried out on the $(3N/2)^3$ convolution grid. The non-linear term is then taken back to Fourier space, where all contributions on wavenumbers $|\mathbf{k}| > N/2$ are removed for dealiasing. This yields a fully dealiased flow field on the N^3 Fourier grid. Other contributions such as the particle feedback force can be computed on a N^3 real space and N^3 Fourier space grid, needing only truncation of the wavenumbers $|\mathbf{k}| > N/2$, rather than $|\mathbf{k}| > N/3$ as in the conventional 2/3-dealiasing routine.

By carefully clipping the particle feedback force in just the right places and times, this prefactor enforces that the effective local volume fraction never exceeds α_0 , ensuring the stability of the integration. While any $\alpha_0 < 1$ works, we set it to the close packing fraction of spheres $\alpha_0 = \pi/(3\sqrt{2}) \approx 0.74$.

For our benchmarking simulations of four-way coupled bubbly turbulence (bubbles with $D = 0.8\eta$ at varying volume fraction $\bar{\alpha}$), we find that the total integrated magnitude of \mathbf{f}_p (L1 norm) is clipped by less than 1% as the volume fraction remains $\bar{\alpha} \lesssim 15\%$, which we deem acceptable.

E. Performance

The two-way coupling routine involves one loop over all Lagrangian particles as well as one FFT over the Eulerian grid. This routine can thus always be run within a runtime that is less than the combined runtime of the Eulerian part (involving two FFTs) and the integration of Lagrangian trajectories (involving also one loop over the Lagrangian particles, but with a heavier workload). Hence, the two-way coupling has a minimal impact on the overall performance.

III. PARTICLE-PARTICLE INTERACTION

A. Efficient excluded volume interaction

To enforce the excluded volume hard-sphere interactions between particles, one needs to solve overlaps between particles at every time step. Solving all excluded volume interactions between all particles submerged in the fluid at large number densities is potentially a daunting problem from the computational point of view. However, instead of checking for all $N_p^2/2$ possible collisions between particles, we can use a boxing approach. This is a classical approach borrowed from Molecular Dynamics that only checks locally for possible collisions [26]. Another computational difficulty is posed by particles that are involved in multiple collisions. Solving all overlaps recursively can in turn also become computationally intractable for sizable clusters of overlapping particles. We therefore propose an approximate algorithm for the excluded volume interaction, with efficiency in mind, that at every timestep solves only the strongest overlap for every particle. The second most severe overlap can then be solved in the next timestep, and so on. This allows us to have an excluded volume interaction algorithm that is efficient enough to allow for large number densities (with the number of particles on the same order as the number of Eulerian grid points), but also accurate enough to ensure a minimal total overlap volume of the particles. We coin this algorithm *YOCO* “You Only Collide Once”. The algorithm is laid out below.




FIG. 4. The boxing approach boxes all particles contained in process p_n (solid boxes), while it sends (shaded solid boxes) and receives (shaded dashed boxes) the outer boxes to its nearest neighboring processes. Overlaps are then only checked between the particles contained within each box itself and its nearest neighboring boxes.

At every timestep:

1. **Boxing.** The bubbles are collected in cubic boxes with sides that are maximally one radius of a particle. This requires the creation of an array of boxes that tiles the full 3D space. By looping over all particles, this array is filled with references to the particles contained in that box. Each parallel process needs to communicate one extra layer of boxes in each direction to its neighboring processes to facilitate the overlap checking in the next step. See Fig. 4.
2. **Overlap checking.** By looping over each box, we check for overlaps between particles within the box as well as between particles in the box and in the directly neighboring boxes. To avoid double counting, we check only for overlaps between particle i and j when the globally unique identifier of particle i is larger than that of particle j . All overlaps are recorded in a local list, storing the references of the overlapping particles and the linear size of the overlap $Q = D - |\mathbf{X}_i - \mathbf{X}_j|$.
3. **Sorting.** The list of overlaps is sorted from largest to smallest Q .
4. **Purging.** We eliminate every second occurrence of a particle, leaving only the strongest overlap for each particle in the list. This can be done efficiently using a hash map as a look-up table to keep track of whether particles have already occurred in the list.
5. **Solving collisions.** All remaining overlaps in the list are solved by moving each pair of overlapping particles outwards over their mutual centerline such that the new $Q = 0$ for this pair as depicted in Fig 5.




FIG. 5. The excluded volume hard-sphere interactions are enforced by moving overlapping pairs of particles radially outwards over their mutual centerline from $\mathbf{X}_i, \mathbf{X}_j$ to $\mathbf{X}'_i, \mathbf{X}'_j$.

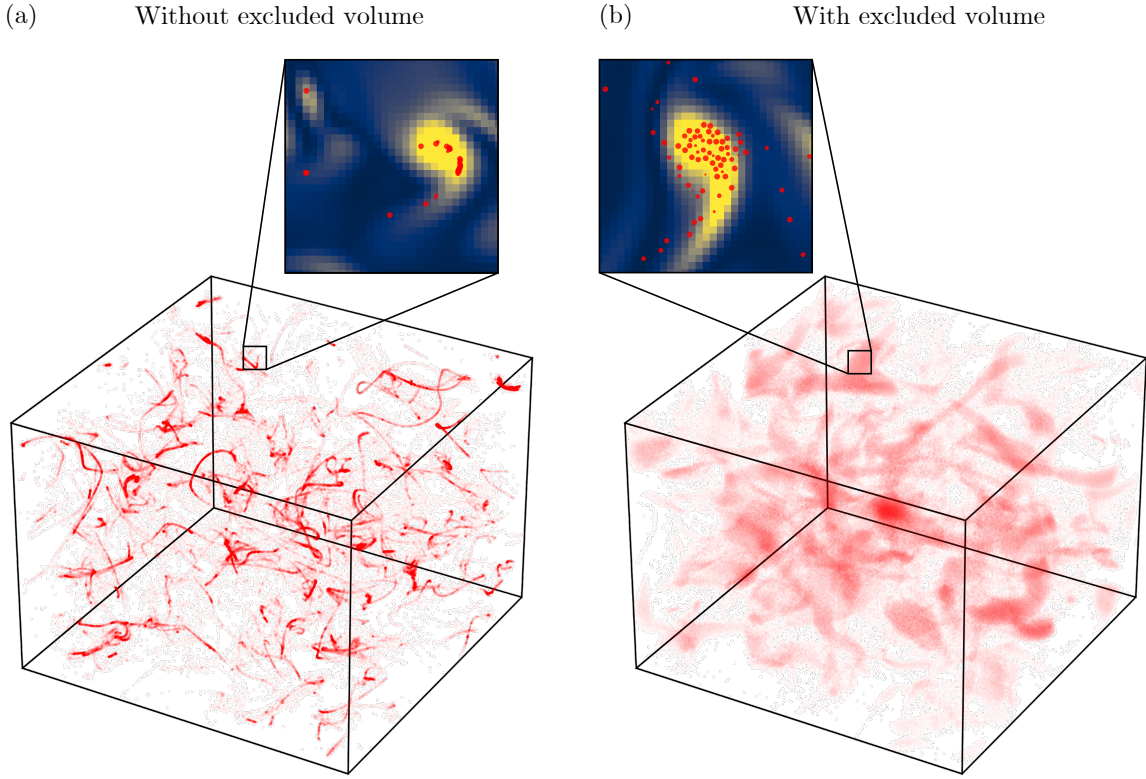


FIG. 6. A snapshot of particle positions for a simulation without excluded volume interactions (a) and with excluded volume interactions (b) with bubbles of size $D = 0.8\eta$ and average volume fraction $\bar{\alpha} = 3.5\%$. The resolution is $N^3 = 128^3$ and number of particles $N_p = 262\,144$ and Taylor scale Reynolds number $\text{Re}_\lambda \approx 35$. The inset zooms show a small cross section of the full domain, where the background color indicates the local enstrophy $\Omega^2 = |\nabla \times \mathbf{u}|^2$.

To assess the effectiveness of this *YOCO* algorithm in enforcing the most important excluded volume interactions, we perform a simulation of bubbly turbulence with and without excluded volume interaction. A resulting exemplary snapshot of particle positions is provided in Fig. 6. This shows that, indeed, due to excluded volume interactions, particles become notably more spread out around the vicinity of vortex filaments, rather than collapsing into the centers of the filaments.

At the more quantitative level, we can compute the overlap volume by Monte-Carlo sampling probes in the whole space and checking whether they are inside or outside at least one particle. The fraction of Monte-Carlo samples that are inside at least one particle will then converge to V_u/V with V_u the total union volume of all particles and V the total volume of the simulation domain. In the case of perfect excluded volume interactions, one should obtain $V_u \rightarrow \sum_i^{N_p} \mathcal{V}_p$, the summed volume of all particles. We can thus define the overlap ratio as

$$\gamma = \frac{\sum_i^{N_p} \mathcal{V}_p - V_u}{\sum_i^{N_p} \mathcal{V}_p}, \quad (14)$$

which yields $\gamma \rightarrow 0$ in the case of perfect excluded volume interactions and $\gamma \rightarrow 1$ in the case of total overlap (vanishing union volume). The results of this Monte-Carlo

sampling are presented in Fig. 7, which shows that the overlap ratio for the simulation without excluded volume interaction yields $\gamma = (78.3 \pm 0.1)\%$ while in the simulation with excluded volume interactions an overlap ratio $\gamma = (0.5 \pm 0.1)\%$ remains. This shows that the *YOCO* algorithm is very effective in reducing the overlap between particles, with an almost negligible remaining overlap ratio. The histogram of the found overlaps for the cases with and without excluded volume interactions is provided in Fig. 8. This shows that for the case with excluded volume interactions, most remaining overlaps are small, while for the case without excluded volume interactions, the most probable overlap is close to full overlap $Q/D \approx 1$. This further assures us of the effectiveness of the proposed *YOCO* algorithm. We find that when pushing the volume fraction $\bar{\alpha}$ even further into the very dense regime, the overlap ratio remains $\gamma \lesssim 3\%$ while $\bar{\alpha} \lesssim 15\%$.

B. Evaluation of the material derivative

A final point of attention concerns the evaluation of the material derivative at the position of the particle as needed for the particle equation of motion Eq. (3). Recall that the material derivative at the position of the particle

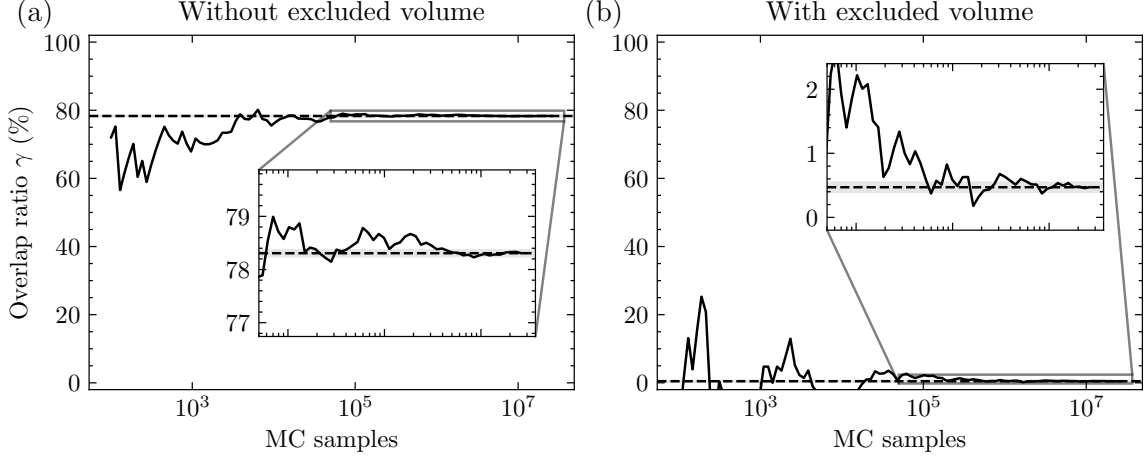


FIG. 7. Monte-Carlo sampling of the overlap volume for a simulation with bubbles of size $D = 0.8\eta$ and average volume fraction $\bar{\alpha} = 3.5\%$ for the case without excluded volume interaction (a) and with excluded volume interaction (b) between bubbles. The figures show the convergence over the obtained overlap ratio as a function of the number of Monte-Carlo samples taken.

is given as

$$\frac{D\mathbf{u}(\mathbf{X}, t)}{Dt} = \frac{\partial \mathbf{u}(\mathbf{X}, t)}{\partial t} + \mathbf{u}(\mathbf{X}, t) \cdot \nabla \mathbf{u}(\mathbf{X}, t). \quad (15)$$

A commonly used trick in Lagrangian tracking is to evaluate the material derivative at the particle position from the full derivative of the fluid velocity at the position of the particle $d\mathbf{u}(\mathbf{X}, t)/dt$. This quantity is numerically easily accessible, as it involves only the result of the interpolation of the velocity at the previous timestep. Then, considering that

$$\frac{d\mathbf{u}(\mathbf{X}, t)}{dt} = \frac{\partial \mathbf{u}(\mathbf{X}, t)}{\partial t} + \frac{d\mathbf{X}(t)}{dt} \cdot \nabla \mathbf{u}(\mathbf{X}, t), \quad (16)$$




FIG. 8. Histogram of the relative overlap Q/D of the found overlapping pairs in a snapshot of a simulation without excluded volume interaction (dashed line) and with excluded volume interaction (solid line) for a simulation with bubbles of size $D = 0.8\eta$ and average volume fraction $\bar{\alpha} = 3.5\%$.

the material derivative at the position of the particle can be obtained as

$$\frac{D\mathbf{u}(\mathbf{X}, t)}{Dt} = \frac{d\mathbf{u}(\mathbf{X}, t)}{dt} + [\mathbf{u}(\mathbf{X}, t) - \mathbf{V}(t)] \cdot \nabla \mathbf{u}(\mathbf{X}, t). \quad (17)$$

This conveniently avoids the evaluation of the partial time derivative of the fluid velocity at the current position of the particle, which would require knowledge of the previous state of the full flow field. However, note that here we have tacitly assumed that $d\mathbf{X}(t)/dt = \mathbf{V}(t)$. While this is usually true, this is not the case with our treatment of particle-particle collisions, which can alter the particle position due to collisions in a way that does not follow the integration of the velocity alone.

Therefore, when applying particle-particle interactions, we do not evaluate the material derivative at the particle position according to Eq. (17), but we instead evaluate Eq. (15) directly. This involves keeping one previous flow snapshot in memory and interpolating both the current velocity as well as the previous velocity at the current position of the particle in order to evaluate the partial time derivative. This can be done with a minimal impact on the overall performance of the code.

C. Performance

Treating pairwise particle collision is a problem that can in principle be of quadratic complexity in the number of particles. While the algorithm proposed here effectively reduces the complexity, for example through the boxing approach that employs the locality of collisions, the complexity of the algorithm is still expected to exceed linearity. For an increasingly large number of particles, the collision routine can thus become the computationally heaviest part of the full four-way coupling code. We have therefore carefully assessed the performance of the




FIG. 9. Performance of the four-way coupling method as a function of the number of particles N_p . Solid lines with symbols denote the runtime of the proposed *YOCO* algorithm for the particle-particle interaction for particles with different buoyancy ratios. The horizontal dotted line/curve coined Eulerian and Lagrangian denote the runtime of routines for the Eulerian part and of the integration of the Lagrangian trajectories, respectively. The vertical dashed-dotted line denotes the number of Eulerian grid points $N^3 = 128^3$ for reference. The particle size $D = 0.65\eta$, such that the volume fraction varies from $\bar{\alpha} = 0.01\%$ to $\bar{\alpha} = 29\%$.

four-way coupling method as laid out in Fig. 9. Since the performance of the particle-particle routine depends on the number of realized collisions, the slowest runtime is obtained for bubbles ($\beta = 3$), as expected. However, as can be appreciated from the figure, thanks to the efficient implementation of the *YOCO* algorithm, the collision algorithm only starts to bottleneck the overall performance as the number of particles approaches the number of Eulerian grid points (N^3). Indeed, here we obtain a scaling of the runtime with the number of particles with a power between 1 and 2.

IV. EXAMPLE OF FOUR-WAY COUPLING

Finally, we present an example simulation of the full four-way coupled bubbly turbulence in Fig. 10. At contrast with the artificial test case where bubbles were kept fixed to enforce uniform density as presented in Fig. 1, here we see that the amplification of the large and inertial scales vanishes. Instead, when the bubbles are free to move and preferentially concentrate, we observe a slight attenuation of the intermediate to dissipative scales and a strong amplification of the smallest scales. We argue that the attenuation is an effect of increased effective viscosity, while the amplification of the smallest scales can be understood as a redistribution of energy towards the smallest scales where the particles are active. These findings are qualitatively consistent with earlier experimental and numerical findings [23, 27–29]. In future

work, we plan to assess the effect of four-way coupling further using the method presented here, studying e.g. the effect on the preferential concentration, pair dispersion, and energy spectra under the influence of different concentrations of various particles.

V. CONCLUSIONS AND OUTLOOK

In this work, we have presented a four-way coupling approach for simulating small particles in a turbulent flow, tailored towards supporting high number densities of particles and/or cases of intense particle clustering. While more accurate methods exist, which extend the support of particle sizes up to the integral scales of turbulence, these more complex methods compromise on computational efficiency, limiting the number of particles that can be resolved. Our point-based method, on the other hand, has a computational cost that is of a similar order as the one for the evolution of the Eulerian flow, at least when the number of particles remains of the same order of the number of Eulerian grid points, which in practice amounts to millions of particles. This method is thus well suited for studying the collective effect of large volume fractions of small submerged particles in turbulent flows as encountered in plenty of industrial and natural applications [4].

Moreover, our method can be straightforwardly extended to more complicated systems, for example when there are other forces acting on the particle, such as an external driving force. By extending the equation of motion of the particle, the feedback mechanism presented here automatically couples its dynamics to the dynamics of the background turbulence through the mutual momentum exchange. Another extension that one can make is to include rotational dynamics of the particle as well as its feedback onto the flow, as is laid out in Ref. [30]. The versatility of this method thus opens up many opportunities to study a large variety of complex particle-laden fluid systems in which resolving a large number of particles is important to capture their collective physical effect.

ACKNOWLEDGEMENTS

We thank Dmitry Pekurovsky for the development of the P3DFFT library. This publication is part of the project “Shaping turbulence with smart particles” with project number OCENW.GROOT.2019.031 of the research programme Open Competitie ENW XL which is (partly) financed by the Dutch Research Council (NWO).




FIG. 10. Panel (a): Energy spectra of the one-phase turbulent flow and the four-way coupled bubbly turbulent flow with average volume fraction of bubbles $\bar{\alpha} = 3.5\%$ with size $D = 0.8\eta$. Panel (b): Four-way coupled energy spectrum compensated by the one-way energy spectrum. This shows a slight attenuation of the turbulence at intermediate to dissipative scales and a relatively strong enhancement of the smallest scales. The shaded regions denote statistical errorbars. The resolution is $N^3 = 128^3$ and number of particles $N_p = 262\,144$ and Taylor scale Reynolds number $Re_\lambda \approx 35$. Compare with Fig. 1 that is run at enforced uniform particle density.

Appendix A: Scheme for Eulerian pseudospectral solver

For the Eulerian part, a classical pseudospectral approach is employed. To solve the Navier-Stokes equations

$$\frac{\partial \mathbf{u}}{\partial t} + \mathbf{u} \cdot \nabla \mathbf{u} = -\nabla p + \nu \Delta \mathbf{u} + \mathbf{f}, \quad (\text{A1})$$

under the condition of incompressibility $\nabla \cdot \mathbf{u} = 0$, we write the velocity field \mathbf{u} in terms of its vector potential \mathbf{b} as

$$\mathbf{u} = \nabla \times \mathbf{b}, \quad (\text{A2})$$

which inherently ensures incompressibility since $\nabla \cdot \nabla \times \mathbf{b} = 0$.

Taking the curl of Eq. (A1) and using the definition of vorticity $\boldsymbol{\omega} = \nabla \times \mathbf{u}$, then yields the Navier-Stokes equations in vorticity-velocity formulation

$$\frac{\partial \boldsymbol{\omega}}{\partial t} = \nabla \times (\mathbf{u} \times \boldsymbol{\omega}) + \nu \Delta \boldsymbol{\omega} + \nabla \times \mathbf{f}. \quad (\text{A3})$$

This can be rewritten as a dynamical equation for the vector potential by noting that $\boldsymbol{\omega} = -\Delta \mathbf{b}$ as

$$-\frac{\partial \Delta \mathbf{b}}{\partial t} = \nabla \times (\mathbf{u} \times \boldsymbol{\omega}) - \nu \Delta^2 \mathbf{b} + \nabla \times \mathbf{f}. \quad (\text{A4})$$

This equation is solved in spectral \mathbf{k} -space, where we define the Fourier transform $\hat{\mathbf{b}}(\mathbf{k}) = \mathcal{F}[\mathbf{b}(\mathbf{x})]_{\mathbf{k}}$. Then we can write Eq. (A4) in the final pseudospectral form

$$|\mathbf{k}|^2 \frac{\partial \hat{\mathbf{b}}}{\partial t} = i\mathbf{k} \times \mathcal{F}[\mathbf{u} \times \boldsymbol{\omega}]_{\mathbf{k}} - \nu |\mathbf{k}|^4 \hat{\mathbf{b}} + i\mathbf{k} \times \hat{\mathbf{f}}. \quad (\text{A5})$$

In practice, this equation is evaluated on a finite \mathbf{k} -space of size N^3 . This should be large enough to resolve the smallest active scale, being the dissipative Kolmogorov scale η , such that the largest wavenumber $k_{\max} \approx 2\pi/\eta$. The non-linear term $\mathcal{F}[\mathbf{u} \times \boldsymbol{\omega}]_{\mathbf{k}}$ is evaluated in real space by first computing the velocity and vorticity in real space as

$$\mathbf{u} = \mathcal{F}^{-1} [i\mathbf{k} \times \hat{\mathbf{b}}]_{\mathbf{x}}, \quad \boldsymbol{\omega} = \mathcal{F}^{-1} [|\mathbf{k}|^2 \hat{\mathbf{b}}]_{\mathbf{x}}. \quad (\text{A6})$$

and then carrying out the cross product between them using pointwise multiplication and transforming this product back to the Fourier space. This is the essence of the pseudospectral method, as evaluating the non-linear term in spectral space would involve a convolution of quadratic complexity, while the Fourier transform can be performed with lin-log complexity, owing to the Fast Fourier Transform (FFT) algorithm [31]. In 3D on an N^3 Eulerian grid, the complexity of the pseudospectral method thus becomes $\mathcal{O}(N^3 \log N)$. For an efficient parallel implementation of the FFT in 3D, we use the P3DFFT package [25].

The pseudospectral method produces higher harmonics in the non-linear term that need to be dealiased. This dealiasing needs to be done after the non-linear term is transformed back into the Fourier space. Our dealiasing approach is laid out in the main text in section II C.

To solve Eq. (A5), we need to integrate it in time. Let us rewrite it as

$$\frac{\partial \hat{\mathbf{b}}(\mathbf{k}, t)}{\partial t} = c\hat{\mathbf{b}}(\mathbf{k}, t) + K(\hat{\mathbf{b}}(\mathbf{k}, t), t), \quad (\text{A7})$$

where $c \equiv -\nu |\mathbf{k}|^2$ and $K(\hat{\mathbf{b}}(\mathbf{k}, t), t) \equiv |\mathbf{k}|^{-2}(i\mathbf{k} \times \mathcal{F}[\mathbf{u} \times \boldsymbol{\omega}]_{\mathbf{k}} + i\mathbf{k} \times \hat{\mathbf{f}})$. Then, we can solve for the vis-

cosity exactly by treating it as an integrating factor

$$\frac{\partial}{\partial t} [\hat{\mathbf{b}}(\mathbf{k}, t) \exp(-ct)] = K(\hat{\mathbf{b}}(\mathbf{k}, t), t) \exp(-ct). \quad (\text{A8})$$

Integrating for a small timestep dt from t_n to $t_{n+1} = t_n + dt$ then gives

$$\begin{aligned} \hat{\mathbf{b}}(\mathbf{k}, t_{n+1}) &= \hat{\mathbf{b}}(\mathbf{k}, t_n) \exp(cdt) \\ &+ \exp(cdt) \int_0^{dt} \exp(-c\tau) K(\hat{\mathbf{b}}(\mathbf{k}, t_n + \tau), t_n + \tau) d\tau. \end{aligned} \quad (\text{A9})$$

To evaluate the temporal integral numerically, we employ a second-order accurate Adam-Bashfort scheme, yielding

$$\begin{aligned} \hat{\mathbf{b}}(\mathbf{k}, t_{n+1}) &\approx \hat{\mathbf{b}}(\mathbf{k}, t_n) \exp(cdt) + \left[\frac{3dt}{2} K_n \exp(cdt) \right] \\ &- \left[\frac{dt}{2} K_{n-1} \exp(2cdt) \right], \end{aligned} \quad (\text{A10})$$

where we abbreviated $K_n \equiv K(\hat{\mathbf{b}}(\mathbf{k}, t_n), t_n)$ and $K_{n-1} \equiv K(\hat{\mathbf{b}}(\mathbf{k}, t_{n-1}), t_{n-1})$. This completes the implementation of the Eulerian part of the numerical method presented here.

-
- [1] S. Balachandar and J. K. Eaton, Turbulent Dispersed Multiphase Flow, *Annual Review of Fluid Mechanics* **42**, 111 (2009).
 - [2] F. Toschi and E. Bodenschatz, Lagrangian properties of particles in turbulence, *Annual Review of Fluid Mechanics* **41**, 375 (2009).
 - [3] V. Mathai, D. Lohse, and C. Sun, Bubbly and Buoyant Particle-Laden Turbulent Flows, *Annual Review of Condensed Matter Physics* **11**, 529 (2020).
 - [4] L. Brandt and F. Coletti, Particle-laden turbulence: Progress and perspectives, *Annual Review of Fluid Mechanics* **54**, 159 (2022).
 - [5] R. Benzi and F. Toschi, Lectures on turbulence, *Physics Reports* **1021**, 1 (2023), lectures on turbulence.
 - [6] M. R. Maxey, The motion of small spherical particles in a cellular flow field, *The Physics of Fluids* **30**, 1915 (1987).
 - [7] A. Crisanti, M. Falcioni, A. Provenzale, P. Tanga, and A. Vulpiani, Dynamics of passively advected impurities in simple two-dimensional flow models, *Physics of Fluids A: Fluid Dynamics* **4**, 1805 (1992).
 - [8] E. Balkovsky, G. Falkovich, and A. Fouxon, Intermittent distribution of inertial particles in turbulent flows, *Physics Review Letters* **86**, 2790 (2001).
 - [9] E. Calzavarini, M. Kerscher, D. Lohse, and F. Toschi, Dimensionality and morphology of particle and bubble clusters in turbulent flow, *Journal of Fluid Mechanics* **607**, 13 (2008).
 - [10] C. S. Peskin, Flow patterns around heart valves: A numerical method, *Journal of Computational Physics* **10**, 252 (1972).
 - [11] E. A. Fadlun, R. Verzicco, P. Orlandi, and J. Mohd-Yusof, Combined immersed-boundary finite-difference methods for three-dimensional complex flow simulations, *Journal of Computational Physics* **161**, 35 (2000).
 - [12] P. Orlandi and S. Leonardi, Dns of turbulent channel flows with two- and three-dimensional roughness, *Journal of Turbulence* **7**, N73 (2006).
 - [13] J. H. Seo and R. Mittal, A sharp-interface immersed boundary method with improved mass conservation and reduced spurious pressure oscillations, *Journal of Computational Physics* **230**, 7347 (2011).
 - [14] R. Verzicco, Immersed boundary methods: Historical perspective and future outlook, *Annual Review of Fluid Mechanics* **55**, 129 (2023).
 - [15] J. K. Eaton, Two-way coupled turbulence simulations of gas-particle flows using point-particle tracking, *International Journal of Multiphase Flow* **35**, 792 (2009).
 - [16] T. H. Van den Berg, S. Luther, I. M. Mazzitelli, J. M. Rensen, F. Toschi, and D. Lohse, Turbulent bubbly flow, *Journal of Turbulence* **7**, 1 (2009).
 - [17] R. Monchaux and A. Dejoan, Settling velocity and preferential concentration of heavy particles under two-way coupling effects in homogeneous turbulence, *Physical Review Fluids* **2**, 104302 (2017).
 - [18] A. G. Lamorgese, D. A. Caughey, and S. B. Pope, Direct numerical simulation of homogeneous turbulence with hyperviscosity, *Physics of Fluids* **17**, 015106 (2004).
 - [19] M. R. Maxey and J. J. Riley, Equation of motion for a small rigid sphere in a nonuniform flow, *The Physics of Fluids* **26**, 883 (1983).
 - [20] M. A. T. van Hinsberg, H. J. H. Clercx, and F. Toschi, Enhanced settling of nonheavy inertial particles in homogeneous isotropic turbulence: The role of the pressure gradient and the basset history force, *Physical Review E* **95**, 023106 (2017).
 - [21] M. A. T. van Hinsberg, J. H. M. t. T. Boonkamp, F. Toschi, and H. J. H. Clercx, Optimal interpolation schemes for particle tracking in turbulence, *Physical Review E* **87**, 043307 (2013).
 - [22] M. Boivin, O. Simonin, and K. D. Squires, Direct numerical simulation of turbulence modulation by particles in isotropic turbulence, *Journal of Fluid Mechanics* **375**, 235 (1998).
 - [23] I. M. Mazzitelli, D. Lohse, and F. Toschi, On the relevance of the lift force in bubbly turbulence, *Journal of Fluid Mechanics* **488**, 283–313 (2003).
 - [24] M. R. Maxey, E. J. Chang, and L. P. Wang, Simulation of interactions between microbubbles and turbulent flows, *Applied Mechanics Reviews* **47**, S70 (1994).
 - [25] D. Pekurovsky, P3DFFT: A framework for parallel computations of fourier transforms in three dimensions, *SIAM Journal on Scientific Computing* **34**, C192 (2012).
 - [26] M. P. Allen and D. J. Tildesley, *Computer simulation of liquids* (Oxford University Press, 1987).
 - [27] M. Lance and J. Bataille, Turbulence in the liquid phase of a uniform bubbly air–water flow, *Journal of Fluid Mechanics* **222**, 95–118 (1991).
 - [28] G. Riboux, F. Risso, and D. Legendre, Experimental characterization of the agitation generated by bubbles rising at high reynolds number, *Journal of Fluid Mechanics* **643**, 509–539 (2010).
 - [29] E. Alm  ras, V. Mathai, D. Lohse, and C. Sun, Experi-

- mental investigation of the turbulence induced by a bubble swarm rising within incident turbulence, *Journal of Fluid Mechanics* **825**, 1091–1112 (2017).
- [30] H. I. Andersson, L. Zhao, and M. Barri, Torque-coupling and particle-turbulence interactions, *Journal of Fluid Mechanics* **696**, 319 (2012).
- [31] M. Frigo and S. Johnson, The design and implementation of FFTW3, *Proceedings of the IEEE* **93**, 216 (2005).