Recurrence computational fluid dynamics for efficient predictions of long-term particle deposition on a cylinder

RESEARCH ARTICLE | SEPTEMBER 08 2025

Recurrence computational fluid dynamics for efficient predictions of long-term particle deposition on a cylinder

Johannes Hansson ; Thomas Lichtenegger ; Stefan Pirker ; Srdjan Sasic ; Henrik Ström

Physics of Fluids 37, 093320 (2025)
https://doi.org/10.1063/5.0283431

Physics of Fluids ARTICLE pubs.aip.org/aip/pof

Recurrence computational fluid dynamics for efficient predictions of long-term particle deposition on a cylinder

Cite as: Phys. Fluids **37**, 093320 (2025); doi: 10.1063/5.0283431
Submitted: 31 May 2025 · Accepted: 18 August 2025 ·
Published Online: 8 September 2025

Johannes Hansson,[1,a)] Thomas Lichtenegger,[2] Stefan Pirker,[2] Srdjan Sasic,[1] and Henrik Ström[1]

AFFILIATIONS

[1]Division of Fluid Dynamics, Chalmers University of Technology, Gothenburg, Sweden
[2]Department of Particulate Flow Modelling, Johannes Kepler University, Linz, Austria

[a)]Author to whom correspondence should be addressed: johanneh@chalmers.se

ABSTRACT

Particle deposition on an object from a turbulent flow is of considerable interest in many applications. Numerical predictions using conventional computational fluid dynamics (CFD) are challenging due to that a large number of individual deposition events must be observed over a long time for the deposition statistics to converge. Here, we investigate the potential for using recurrence CFD (rCFD) to efficiently and accurately predict particle deposition on a cylinder for Reynolds numbers (Re) in the interval $20 \leq \mathrm{Re} \leq 10\,000$. We quantify the front- and back-side deposition efficiencies independently, analyze the locations and timings of deposition events, and benchmark the computational performance. We find that rCFD predicts deposition efficiencies with similar accuracy as the corresponding CFD simulations, but at a fraction of the computational cost. The most significant deposition occurs on the front side of the cylinder and is very well described for all Reynolds numbers investigated. For $\mathrm{Re} = 10\,000$, we observe a dependence on the rCFD database length in the prediction of the much less effective back-side deposition, as the database only contains a limited subset of the more rare flow behaviors responsible for this deposition. These results can be used to accelerate particle deposition studies by several orders of magnitude, which would bring significant benefits for computationally challenging applications, such as sensor soiling in the car industry, icing on aircraft, and ash build-up in boilers.

I. INTRODUCTION

Particle deposition studies are important for both academic and industrial applications. In academia, there is interest in understanding the underlying mechanisms responsible for particle dispersion and deposition on surfaces.[1–4] In industry, there is often system-specific interest in predicting how a particular design would perform in a given operating environment.[5,6] In these two points, a common view is that high-resolution simulations are needed for properly resolving the behavior of depositing particles close to the surface.

The process of initially suspended particles depositing on the surface of a submerged object in a turbulent fluid flow can be generalized into a few steps.[7] First, the particles are brought to an accumulation region close to the surface by some feature of the flow, known as in-sweeps.[7,8] This near-surface region exhibits high particle concentrations, but an additional process is needed to bring the particles all the way to the surface. This final part is due to either of two contributions: deposition due to particle inertia, or deposition due to diffusion. The comparative strength of the two processes depends mainly on particle inertia, with heavier particles directly impacting the wall, whereas lighter and smaller particles depend on diffusion processes in order to deposit.[9]

To properly capture these near-wall events in numerical simulations, the carrier fluid must be very accurately described close to the surface.[2] Especially, features such as turbulence must be adequately captured in order to properly describe the process of transporting particles from the near-wall accumulation region all the way to the surface.[7] Describing turbulence can be done either by solving for all of the turbulent structures or by modeling them. The former is very computationally demanding, the latter is theoretically challenging and, using currently available methods, not nearly as accurate.[10–13]

Long-term predictions of particle deposition phenomena from turbulent flows are further complicated by the large span of temporal scales involved. An individual deposition event may require the combined action of several coherent structures in the turbulent flow field,

implying that a fine temporal resolution is needed for its description. At the same time, enough deposition events must be simulated for long-term statistics of the process to be accurately determined.

High-fidelity simulation of the carrier fluid can be performed using, for example, direct numerical simulation (DNS) techniques.[1,14] These provide fully resolved flow fields from the Navier–Stokes equations, but suffer from exceptionally high computational costs. Generally, only small-scale problems can be studied using DNS.[15,16] To reduce the computational requirements, several alternative solution techniques have been developed that focus on modeling parts of the small-scale turbulence, rather than resolving it. Two common modeling approaches are known as large eddy simulation (LES)[17,18] and Reynolds-averaged Navier–Stokes (RANS).[19,20] In the former, most of the turbulence is resolved, but the smallest scales are modeled. This simplification reduces computational cost but also sacrifices some accuracy. In the latter method, the underlying Navier–Stokes equations are time-averaged, which reduces computational cost, but unfortunately also accuracy, even further.

To try to overcome some of the limitations in these two approaches, sometimes hybrid LES/RANS methods, such as detached eddy simulation (DES),[21] are used. This method combines the LES and RANS methods by modeling the large-scale structures using the LES approach and the small-scale, near-wall turbulence using RANS. This way, a mix of turbulence-resolving and turbulence-modeling can be used in the same simulation, combining the properties of the two approaches. DES simulations are often used in industrial settings, such as in the automotive[22,23] and aerospace[24,25] industries. In this work, we will use the DES method due to the favorable combination of accuracy and computational performance for external flows at industrially relevant Reynolds numbers.

Unfortunately, even with these simplifications of the solution technique, the computational requirements are still very high, especially for complex, real-world flow systems. Further improvements are therefore needed to reduce computational cost. The most simple and straightforward simplification is to use time-averaged flow fields. This works well for some types of studies, but is inaccurate for investigations that depend on transient flow effects. Particle deposition on a submerged body, such as the cylindrical geometry of interest in this work, is an inherently transient process that depends on the unsteady vortex shedding behind the cylinder.[26,27] It should be noted that these transient effects impact both the front-side deposition (via boundary-layer fluctuations) as well as the backside deposition (via intermittent wall-normal fluid motion in the vortical structures).[1] Therefore, time averaging of the fluid flow fields cannot be expected to represent a successful approach here.

Another strategy for reducing the computational cost is to use reduced-order methods, where the most important degrees of freedom are computed using, for example, proper orthogonal decomposition (POD)[28] or dynamic mode decomposition (DMD).[29] The POD strategy decomposes the flow field evolution into discrete modes, with varying energy levels. The superposition of the most energetic modes can be used to reconstruct an approximation of the original flow fields, reducing the complexity from the full problem description. A disadvantage with this method is that it suffers from stability problems for long time frames. Methods for mitigating these shortcomings have been developed, but they rely on stabilizing the predictions using conventional computational fluid dynamics (CFD), which are computationally expensive. Novel reduced-order approaches based on machine learning, such as β-variational autoencoders, are also being currently researched.[30] While promising, they too produce predictions that diverge over time as small initial fluctuations in the model accumulate. It remains to be investigated whereas the long-term predictions obtained by these approaches are of sufficient quality to accurately predict two-phase flow phenomena such as particle deposition.

Instead, in this work, we use the recurrence computational fluid dynamics (rCFD) framework[31,32] for generating approximate solutions to the flow fields. This method is useful for generating long-term approximations to flow fields that exhibit some degree of periodicity in time. In the rCFD method, a small database of CFD snapshots is used to characterize the flow. Then, these snapshots are rearranged and reused to create a recurrence path that extrapolates the evolution of the flow fields far beyond the range of the database. Since this method only uses flow field snapshots that come directly from a highly resolved CFD simulation, we get very high spatial resolution in the approximation as well. Stability is also guaranteed since (i) the generated recurrence path only consists of previously seen CFD snapshots and (ii) the temporal average and variance of the extrapolated series are the same as those of the database.[33] Furthermore, if an aerodynamics characterization has already been performed using conventional CFD for the system in question, the database contents exist already and no specific training of the rCFD model is required.

When using the rCFD method, a challenge manifests itself in the form of the discrete temporal jumps that are an inherent part of the method. These jumps introduce discontinuities in the flow fields that the particles will experience, potentially leading to unphysical behavior. More specifically, the rCFD approximation may deteriorate the spatio-temporal correlation of the coherent events in the fluid flow description, such that the simulation of the particle deposition process (bulk turbulence, in-sweep, interactions of near-wall fluctuations, and particle inertia) is corrupted. It is therefore not clear *a priori* that rCFD can accurately predict all relevant particle deposition statistics (deposition efficiencies, locations, and temporal correlation).

In this work, we demonstrate the potential for using rCFD to facilitate particle deposition simulations using the example case of particle–laden flow around a cylinder. Both laminar two-dimensional and turbulent three-dimensional flows are considered, with Reynolds numbers ranging between 20 and 10 000. We show that the impact efficiencies computed using rCFD provide accurate approximations of the efficiencies computed using traditional CFD, while deposition locations and temporal correlations exhibit sensitivity to the rCFD database length. Additionally, we show that the rCFD deposition efficiency approximations can be computed at a fraction of the computational cost of conventional methods.

II. THEORETICAL BACKGROUND

In Secs. II A–II E, we present necessary theoretical considerations for modeling the carrier fluid and the suspended particles. The process of particle deposition is also discussed. Finally, a description of how the rCFD method works is provided.

A. Fluid description

The fluid description for the carrier phase is based on the incompressible Navier–Stokes equations,

$$\nabla \cdot \mathbf{u}_f = 0 \tag{1}$$

and

$$\frac{\partial \mathbf{u}_f}{\partial t} + (\mathbf{u}_f \cdot \nabla)\mathbf{u}_f = -\frac{1}{\rho_f}\nabla p + \nabla \cdot \left(\nu_{eff}[\nabla \mathbf{u}_f + (\nabla \mathbf{u}_f)^{\mathrm{T}}]\right), \tag{2}$$

where $\mathbf{u}_f$ is the fluid velocity, ρ_f is fluid density, p is pressure, and $\nu_{eff} = \nu + \nu_t$ is the effective viscosity. Here, $\nu = \mu_f/\rho_f$ is the kinematic viscosity, ν_t is the eddy kinematic viscosity, and μ_f is the fluid dynamic viscosity. In this work, ν_{eff} is obtained using the delayed detached eddy simulation (DDES) method.[34] DDES represents a modification of the original DES method and improves the handling of thick boundary layers and counteracts the issues of grid-induced separation found in traditional DES. The flow in this work is simulated using the OpenFOAM v2313 $k-\omega$ shear-stress transport (SST) DDES solver implementation.[34–36]

Flow around a cylinder is a well-researched problem, with an abundance of studies characterizing the flow behavior.[14,27,37,38] It is known that the external flow is inherently time-dependent for all except the smallest values of the Reynolds number

$$\mathrm{Re} = \frac{U_\infty D}{\nu}, \tag{3}$$

where U_∞ is the fluid free-stream velocity and D is the cylinder diameter. Steady flows are found for $\mathrm{Re} < 40$.[38] Above this limit, the system exhibits an unsteady behavior with regular vortex shedding in the wake. The shedding frequency is often characterized using the Strouhal number, which is defined by

$$\mathrm{Str} = \frac{fD}{U_\infty}, \tag{4}$$

where f is the vortex shedding frequency.

Since vortex shedding is a time-dependent flow phenomenon, a time-dependent mathematical description of the fluid is necessary in order to properly capture the flow behavior. Transient simulations are therefore needed to accurately capture the interaction between the oscillating carrier fluid and the dispersed particles.

B. Particle description

In this work, the Lagrangian particle tracking framework is used. In this description, particles are treated as spherical point particles with mass m, and their trajectory (using the acceleration vector $\mathbf{a}$) is computed using Newton's second law of motion,

$$\mathbf{F} = m\mathbf{a}. \tag{5}$$

The particle force vector $\mathbf{F}$ is, in general, expressed as a sum of several contributing forces, such as drag, lift, and gravity. We are primarily interested in small, non-Brownian particles that are significantly heavier than the carrier fluid; hence, we consider only the effects of the drag force, which is the dominating force at these conditions.[1] The particle Reynolds numbers are consistently lower than unity, so the effects of the lift force can be neglected.[39,40]

The drag force contribution is defined using

$$\mathbf{F}_D = \frac{1}{2}\rho_f \mathbf{u}_{rel}|\mathbf{u}_{rel}|C_D A, \tag{6}$$

where $\mathbf{u}_{rel} = \mathbf{u}_f - \mathbf{u}_p$ represents the relative velocity between the fluid and the particle ($\mathbf{u}_p$), C_D is the drag coefficient, and A is the projected area of the particle into the flow. For the special case of the drag force on a sphere, such as for the particles used in this study, the value of C_D is defined by[41,42]

$$C_D = \frac{24}{\mathrm{Re}_p}\left(1 + \frac{1}{6}\mathrm{Re}_p^{2/3}\right), \tag{7}$$

where

$$\mathrm{Re}_p = \frac{\rho_f |\mathbf{u}_{rel}| d_p}{\mu_f} \tag{8}$$

is the particle Reynolds number.

The point-particle approximation used in the Lagrangian framework simplifies both the mathematical description and the simulation methodology. For example, if the particle is small compared to the fluid spatial scales, then the fluid properties can be assumed constant over the entire particle surface. It is then enough to track the particle center of mass, rather than resolving the entire fluid–particle interface. This simplification then reduces the computational requirements significantly.

However, a downside of this approximation is that it is only valid under certain conditions. Mainly, it is only valid for particles that are small compared to the spatial scales of the fluid. Comparatively large particles cannot be accurately modeled with this approximation since the fluid properties changeover the particle surface. Computationally expensive boundary-resolving methods are then needed for high-accuracy simulations. In this work, we assume that all particles are small enough to justify the point-particle approximation. This is indeed a reasonable assumption for most studies of particle deposition on bluff bodies in industrial settings.

Even with the point-particle approximation, some of the necessary calculations require a corresponding particle diameter, namely the drag force calculation and the particle–wall collision detection mechanism. All particles, therefore, have an associated diameter that is used only by these stages of the simulation.

For the high-Reynolds cases ($\mathrm{Re} \geq 1685$), turbulent fluctuations are present in the carrier fluid, which affects the suspended particles.[1] Since the DDES solution methodology used in this work models some of the turbulence, rather than resolves it, a turbulent dispersion model is needed to account for the effects of unresolved turbulence on the particle trajectories. This work uses the OpenFOAM `StochasticDispersionRAS` model, which is an eddy interaction model that assumes that the turbulence is isotropic. This model works by superposing the effects of a synthetic eddy on each particle by adding a random velocity fluctuation to the fluid velocity seen by the particle.[9] Each particle is affected by exactly one such eddy at a time and the effect is constant within a certain time interval. Once this time has passed, a new synthetic eddy is generated. This time, called the interaction time, τ_{int}, is the minimum of the eddy lifetime τ_e and the transit time τ_R, which represents the time required for the particle to cross the eddy. Therefore,

$$\tau_{int} = \min(\tau_e, \tau_R), \quad \tau_e = \frac{k}{\epsilon}, \quad \tau_R = \frac{\lambda_e}{|\mathbf{u}_{rel}|}, \tag{9}$$

where $\epsilon = C_\mu k\omega$ is the turbulence dissipation. The eddy length scale is defined as

$$\lambda_e = \frac{C_\mu^{3/4} k^{3/2}}{\epsilon}, \tag{10}$$

where $C_\mu = 0.09$. During this time interval, the particles experience random velocity fluctuation

$$\mathbf{u}' = \sigma_f |N_w| \mathbf{d}_r, \tag{11}$$

where $\sigma_f = \sqrt{2k/3}$ is the root mean square of the velocity fluctuations and N_w is a random variable following the standard normal distribution $\mathcal{N}(0,1)$ with mean 0 and standard deviation 1. The random direction vector

$$\mathbf{d}_r = \begin{pmatrix} a\cos\theta \\ a\sin\theta \\ u \end{pmatrix}, \quad \begin{cases} \theta = \omega_{r1} 2\pi \\ u = 2\omega_{r2} - 1 \\ a = \sqrt{1-u^2} \end{cases} \tag{12}$$

represents the direction of the instantaneous velocity fluctuation due to the turbulent dispersion. Here, two additional random variables are defined, ω_{r1} and ω_{r2}, which are sampled from the standard uniform distribution $U(0,1)$.

This work uses a fine mesh, which means that most of the turbulence is resolved, rather than modeled. Thus, the turbulent dispersion model only has a marginal impact on the computed deposition rates. This has been confirmed by simulations, where marginally lower impact efficiencies are found on the back of the cylinder when the turbulent dispersion model is disabled. No differences were found for the front-side deposition rates. In order to account for the small fraction of modeled turbulence, the turbulent dispersion model is enabled in all results presented in this work.

In order to compare results at different levels of particle inertia, we use the Stokes number St, which is defined as

$$\mathrm{St} = \frac{\rho_p U_\infty d_p^2}{9 D \mu_f}, \tag{13}$$

where ρ_p is the particle density. The Stokes number represents the ratio of the particle response time and a corresponding fluid response time. Intuitively, large values of this parameter represent particles with high inertia that travel straight ahead, irrespective of the local flow field behavior. Low values, on the other hand, represent tracer-like particles with low inertia that follow the flow almost perfectly.

In this work, we assume that the particle volume fraction is low, so that the fluid affects the particles, but the particles do not affect the fluid. Additionally, the particles do not interact with each other. This one-way coupling approach also simplifies the interpretation of convergence tests, as the fluid can then be analyzed for convergence independently of the particles.

C. Particle deposition

An important aspect of particle deposition studies is the fraction of particle impacts that lead to deposition. To simplify the description, in this work, we assume that all particle–wall impacts lead to deposition, as is done in similar studies.[1,2,43] A collision is triggered when the center of mass of a spherical particle is closer to the cylinder surface than the particle radius. Once a collision has occurred, the particle deposition location is registered and the particle trajectory is thereafter terminated. A deposited particle becomes completely passive and does not affect any other transport or deposition processes.

The main particle deposition metric used in this work is the impact efficiency, which is defined as

$$\eta = \frac{N_{\text{deposited}}}{N_{\text{injected}}}, \tag{14}$$

where $N_{\text{deposited}}$ is the number of particles that have deposited on the surface and N_{injected} is the number of particles that have been injected upstream in the projected area of the deposition surface. This efficiency is highly dependent on particle inertia, so it is computed individually for each particle size.

D. Recurrence CFD (rCFD)

There are two main versions of the rCFD method: flow-based and transport-based rCFD. In the flow-based version, the fields in the CFD snapshots in the database are played back in an order that approximates the long-term evolution of the flow fields. The evolution of particles and passive scalar fields can then be calculated as if they had been calculated at the same time as the CFD fields. Transport-based rCFD, on the other hand, considers the transport of information between (possibly distant) cells with large time steps, without having to play back the complete field evolution within each step.[44] In this work, we use the flow-based version of rCFD due to the straightforward implementation of particle inertia effects, which play an important role in deposition studies. For the application of the transport-based variant, one would have to describe the particle dynamics from an Eulerian perspective with a passive transport equation. Taking into account the drift relative to the gas flow would be much more challenging and less accurate than for the Lagrangian formulation, which we employ in the present study.

The rCFD method relies on recurrence statistics, which are constructed from pairwise similarity measurements between snapshots in the rCFD database. Several similarity metrics are available, with the choice of measure being up to the user. In this study, we use a distance measure defined by the following equation:[31,32]

$$\mathcal{D}(t,t') = \left(\frac{1}{N} \int \left(\mathbf{u}(\mathbf{r},t) - \mathbf{u}(\mathbf{r},t') \right)^2 \mathrm{d}r^3 \right)^{1/2}, \tag{15}$$

where

$$N = \max_{t,t'} \left(\int \left(\mathbf{u}(\mathbf{r},t) - \mathbf{u}(\mathbf{r},t') \right)^2 \mathrm{d}r^3 \right) \tag{16}$$

is a normalization factor. In these equations, the difference squared of a fluid vector quantity, such as velocity $\mathbf{u}$, is integrated over the volume of the system. Notice that the difference is between fields at different times t and t'. The integration must then be performed for all unique combinations of t and t' in a CFD database of flow field snapshots to yield the complete distance matrix.

The database consists of a short CFD simulation in which the flow fields have been regularly sampled and saved. The sampling should be frequent enough that the fields do not change drastically between sampling points. The length of the sampling should be long enough that the main flow features are captured. This is generally the case once a few oscillation periods have been sampled. A crucial trade-off between storage space, computational requirements, and resulting rCFD approximation accuracy depends on the construction of a high-quality database.

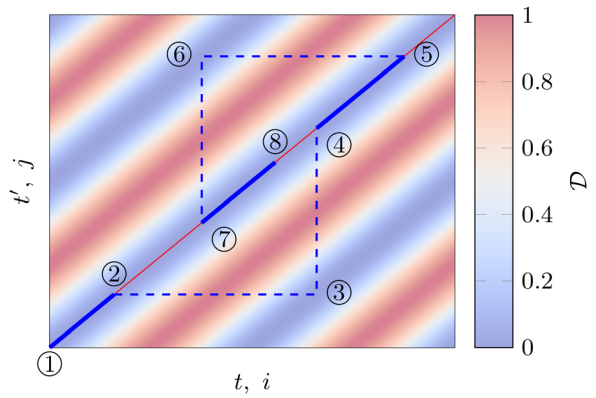

FIG. 1. Schematic illustration of a recurrence matrix. The value of each matrix element $\mathcal{D}_{ji}$ represents the dissimilarity between states i and j in a CFD database according to the chosen difference metric.

Using the distance measure and database above, a recurrence matrix such as the one schematically illustrated in Fig. 1 can be constructed. Here, the i and j indices on the coordinate axes represent the rCFD snapshots in the database at discrete times t and t', respectively. The numerical values between 0 and 1 of each distance matrix element $\mathcal{D}_{ji}$ represent the degree of likeness between the two CFD snapshots i and j. A distance value of $\mathcal{D}_{ji} = 0$ represents the case when the two snapshots are identical, which is true by definition on the main diagonal, where $i = j$. The matrix element with $\mathcal{D}_{ji} = 1$ represents the most dissimilar pair of states i and j in the database. The range of matrix element values is kept between 0 and 1 by the normalization factor in Eq. (16). The presence of clear off diagonal bands in the recurrence matrix indicates that the system exhibits clear recurrent behavior in time. In Fig. 1, we see the main diagonal, one clear off diagonal band, and a small part of a second off diagonal band in the lower right corner of the plot. This structure indicates that this recurrence matrix is associated with a CFD database of almost two complete flow oscillations. Other systems may exhibit similarity between two snapshots without the clear off diagonal structure, which indicates that there are similarities, but not a clear recurring pattern.

E. rCFD recurrence path

Once the recurrence matrix has been constructed, we can generate a recurrence path, which is a time-extrapolated approximation of what the complete solution to the carrier phase flow fields would look like. A flow chart of the path generation algorithm is presented in Fig. 2.

First, the method starts by initializing an empty recurrence path. In terms of implementation, this can be an empty array. Then, a reference to the first element in the CFD database is added to the list, which can be an integer index, with the number 0 indicating the first element in the database. After this, a random segment length is selected. The length represents the number of snapshots from the database that will be added in direct sequence for a particular segment of the rCFD path. This length can at most be until the end of the available database, or to a particular maximum set by the user. Once the current segment length has been chosen, all CFD snapshots from the database in this identified segment are added to the recurrence path. Again, the

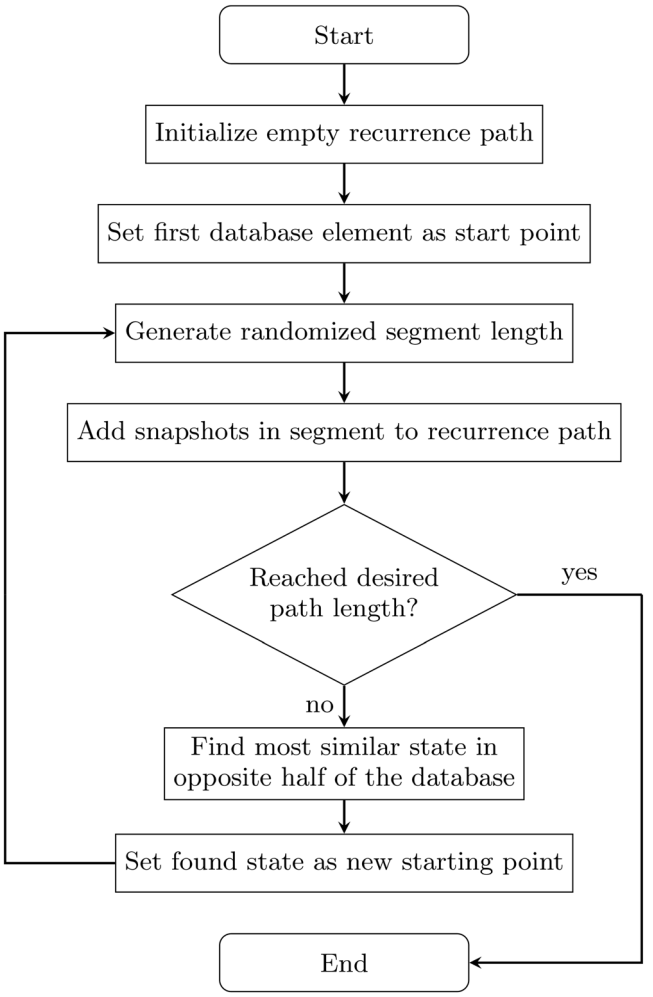

FIG. 2. Flowchart of the rCFD path generation process.

recurrence path can, for example, be an integer array and then the snapshots are added by appending the snapshot database indices to the array. We have now generated a short recurrence sequence. If this sequence is long enough for the studies we want to run, then we end here. If not, then we find the snapshot that is as similar to our current ending snapshot as possible. This new snapshot is selected from the opposite side of the database. In other words, if we end the segment in the second half of the database, we look for similar segments in the first half of the database, and vice versa. This avoids certain database sampling issues, such as the recurrence path only sampling the end of the database. Once the new starting point has been found, we go back to the step "generate randomized segment length" and continue looping until the requested recurrence path length has been reached.

The algorithm can be further illustrated by walking through the recurrence matrix plot (Fig. 1). The path generation process is initialized at the database starting point ①. A random segment length is generated, which puts the segment ending point at ②. We test if we have reached the desired recurrence path length, which we, in this example, have not. Therefore, since we stopped in the first half of the database,

we look for the most similar CFD snapshot in the second half of the database. We find the most similar snapshot ③ by minimizing $\mathcal{D}_{ji}$ along the horizontal dashed line, on the opposite side of the matrix. Taking note of the snapshot time, we set ④ as the starting point for the next segment. The process is then continued until the desired recurrence path length has been reached. In this case, we get the recurrence path

$$① \rightarrow ②,\ ④ \rightarrow ⑤,\ ⑦ \rightarrow ⑧,\ \ldots, \tag{17}$$

where the arrows represent all the intermediate CFD snapshots and the commas indicate jumps in the recurrence path. The jumps might be noticed as small twitches in the resulting approximated flow fields due to minor differences between the ending snapshot in the first segment and the starting snapshot in the second segment. These artifacts appear because the minimized distance metric is a global measure and not necessarily a local one.

Finally, once the path has been generated, it must be translated into a form that is readable by the chosen simulation package. This involves translating the array of snapshot indices into a sequence of simulation data files on computer storage. Further details are software-specific and, therefore, omitted here for brevity.

III. METHOD AND COMPUTATIONAL SETUP

In the following two subsections, we describe problem-specific parameters used in this study, first for the fluid and then for the particles.

A. Fluid setup

In order to properly describe the motion of depositing particles close to the surface, the displacement thickness δ_1 must be properly resolved.[2,45,46] Literature results indicate that using 4–10 grid cells within this thickness should be enough to resolve this part of the boundary layer.[2,47] A previous study suggests a first-cell radial size of at most[2]

$$\Delta \leq \frac{0.3240}{4 \cdot \sqrt{\mathrm{Re}}} D. \tag{18}$$

However, for the simulations in this work, a smaller first-cell radial size of

$$\Delta \approx \frac{d_p}{2} \tag{19}$$

was required, where d_p represents the diameter of the smallest particle size used in any given simulation. This corresponds to approximately

$$\Delta = \frac{0.3240}{32 \cdot \sqrt{\mathrm{Re}}} D. \tag{20}$$

The boundary layer cell size grading is fixed at 10% per cell layer, ensuring that the suggested number of cells in the δ_1 layer is achieved. The computational domain is large enough to prevent domain edge effects from affecting the results.

In terms of discretization schemes, in general, second-order schemes are used for time and space discretization. However, to improve stability, a second-order scheme with a limiter is used for the divergence term in the Navier–Stokes equations. This term then achieves between the first and second order accuracy.

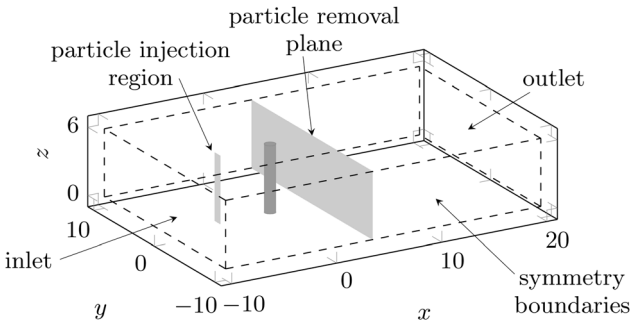

FIG. 3. Simulation domain used in this study. Distances in terms of cylinder diameters *D*.

The simulation domain spans 30, 20, and 6D in the x, y, and z directions, respectively, see Fig. 3. The cylinder is centered $10D$ downstream of the inlet in the x direction, centered in the y direction, and spanning the entire z direction. The inlet has a fixed fluid velocity in the x direction, such that there are no turbulent velocity fluctuations on the inlet boundary. At the outlet, we prescribe a gauge pressure of zero. The side walls act as symmetry planes to minimize boundary effects. The mesh is fine enough that wall functions are not necessary close to the cylinder surface.

This study uses a three-dimensional hexahedral mesh (Fig. 4) with approximately 9.6×10^6 cells. All cases use the same base mesh to simplify comparisons. The cylinder boundary layers are slightly different in each case to account for the condition in Eq. (20). Additionally, the two-dimensional cases use a mesh that is based on a plane cut of the three-dimensional version. The two-dimensional mesh is then very similar in structure to the three-dimensional one, but with a thickness of only a single cell, and consisting of approximately 55 000 cells.

The fluid solver uses the PISO algorithm with a fixed nondimensional time step of $\Delta t U_\infty / D = 3.125 \times 10^{-3}$, which maintains the Courant number at a value less than one. This time step is small enough to accurately capture the particle response time, with at least seven CFD time steps within the shortest particle response time. Additionally, the particle phase solver uses substepping for the time step, which means that the particle equations are solved with a time step that is at most 0.3 of the CFD time step. The particles are then resolved with at least 23 particle time steps within the shortest particle response time. In total, the simulation is run for $tU_\infty/D = 300$ time units, of which the first $tU_\infty/D = 100$ are discarded to eliminate the effects of the startup transients. The rCFD simulations do not have a startup transient, so these simulations skip this initial part. The fluid density is constant ($\rho_p/\rho_f = 1000$), and the fluid kinematic viscosity is adjusted for each case to give the desired Reynolds number, keeping all other properties constant.

In this work, we use very well-resolved DDES, where most of the turbulent kinetic energy is resolved, rather than modeled. This is necessary in order to assure that we have an accurate baseline when compared to literature studies using highly resolved flow simulations. The fraction of resolved turbulent kinetic energy is consistently over 99%.[48] The following Reynolds numbers are simulated in the present work: 20, 100, 1685, 6600, and 10 000.

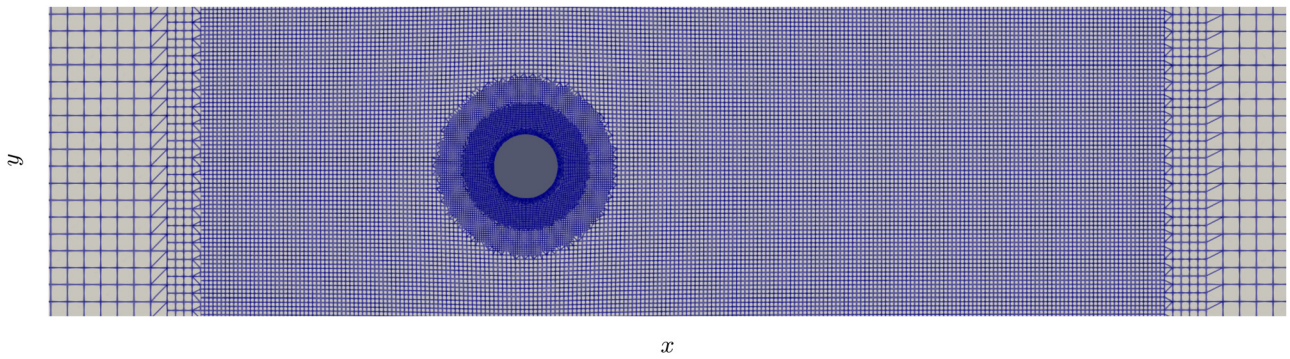

FIG. 4. Cutting plane of the fluid mesh used in this work. Flow is from left to right.

1. Fluid convergence tests

In order to assure proper convergence of the simulations of the carrier fluid phase, mesh convergence tests and comparisons to literature values are performed. The results are assessed using the drag and lift coefficients, as well as the Strouhal number, for the cylinder.

A comparison of the values obtained in this study to values found in the literature is presented in Table I. The present results agree well with the literature data. The values reported for this study were obtained by averaging over 60 vortex shedding periods. Worth noting is that the C_L rms (root mean square) results depend strongly on the aspect ratio of the cylinder, explaining the rather large span of values in previous studies. Given the good agreement between our results and literature values, we consider the carrier fluid simulation setup to be adequate.

B. Particle setup

For the particles, we need to specify several properties such as the number of particles to inject, their diameters, and where to inject them. The particles are injected randomly in the plane of the projected cylinder's frontal surface area, located $5D$ upstream of the cylinder, see Fig. 3. In this region, the fluid flow lines, and hence particle paths, are undisturbed by the presence of the cylinder. Particles are injected over a time span of $tU_\infty/D = 50$ time units, which represent approximately ten vortex shedding periods. Particle injection begins at time $tU_\infty/D = 100$, after the startup transients have subsided. Consequently, particle injection stops at time $tU_\infty/D = 150$, and the existing particles are allowed to propagate freely until the simulation ends at $tU_\infty/D = 300$. The particle density is chosen so that $\rho_p/\rho_f = 1000$, as this density ratio is relevant to both air–water systems and many industrial gas–particle systems. For these reasons, this density ratio is also consistent with most previous studies from the literature.

Particle sizes are determined using the Stokes number in Eq. (13). Numerical values of the Stokes number in particle deposition studies often span in the range of 0.01–10.[1,2,47] Some studies, especially experimental ones, investigate Stokes numbers outside this range, both lower and higher values.[50] In this study, we limit the Stokes number to the range of 0.045–2. This range captures the main flow behavior, while at the same time minimizing the computational cost.

The required number of injected particles is based on the expected number of impacts for a certain particle size. As the expected impact efficiency for a certain particle size depends on the Stokes number, literature values for the expected impact efficiencies[1] may be used to estimate the required number of injected particles. In this study, we aim for at least 100 particle hits on the front side of the cylinder for each Stokes number. Equation (14) can then be used to estimate N_{injected} by using the values $N_{\text{deposited}} = 100$ and η taken from literature. Note that using this estimation procedure does not influence the final computed results for η. The estimation only minimizes the computational effort required to achieve the target $N_{\text{deposited}} \approx 100$ particle impacts per Stokes number needed for good statistics. Based on this estimation procedure, a total of 6×10^6 particles are injected during the course of the simulation. A larger number would lead to a finer sampling of the flow fields with respect to deposition events, but at the cost of increased computational effort. Moreover, to further reduce computational costs, particles are removed from the domain once they pass a fluid cross section plane located behind the cylinder at $x = 4D$, also illustrated in Fig. 3. Here, the fluid is moving strictly away from the cylinder (in the positive x direction), so there are no flow structures that are capable of transporting particles back to the cylinder. The particles can then be removed from the domain without affecting the results.

TABLE I. Comparison of nondimensional fluid flow quantities at Re = 10 000.

Reference	Method	C_D mean.	C_L rms	Strouhal number
This work	DDES	1.091	0.252	0.203
Nguyen and Nguyen (2016)49	DDES	1.133	0.262–0.363	0.196–0.201
Dong and Karniadakis (2005)14	DNS	1.143	0.448	0.203
Norberg (2003)37	Exp.	⋯	0.25–0.45	≈0.2

C. rCFD

The rCFD method has a few parameters that can be adjusted. First, the method is capable of handling coarse-graining of both space and time.[31] In this work, we record every 16th CFD time step, which corresponds to approximately 90–100 sampling points per vortex shedding period. The recording time step is chosen to be fine enough so that the flow fields do not change significantly between steps, while at the same time coarse enough to minimize database storage size. A coarser recording frequency than every 16th time step is probably possible; however, further coarsening of this recording frequency has not been tested in this work. Coarse-graining in space can be performed by mapping the CFD fields to a coarser rCFD mesh with the option to retain sub-grid scale variances. This method is, however, not used in this study.

For the selection of the distance metric between database snapshots, the metric defined in Eq. (15) is used here. Other definitions are possible, depending on which fields are most important in a certain flow system. It is also possible to limit the size of the considered domain, both for the distance calculation and when constructing the CFD database. For example, if only the region immediately around the cylinder is of interest, then the field information far away can be discarded. This saves storage space, which can be a major factor in deciding the CFD database size.

The CFD database need to contain enough flow information to be able to generate good approximations of the flow fields. In this work, we use a fixed database size of $tU/D = 15$ for all cases, which for the transient cases $\mathrm{Re} \geq 100$ correspond to about three vortex shedding periods. Note that the $\mathrm{Re} = 20$ case represents a steady-state flow and does not have a well-defined vortex shedding period. This value is chosen based on the saturation of the mean nearest-neighbor distance in the recurrence matrix,[32] and it is also in line with the expected time duration required to accurately assess the aerodynamic behavior of a complex bluff body using a scale-resolving turbulence model.[51]

The recurrence path length required for the predicted deposition efficiency to converge can be determined by observing the evolution of the signal over time. In the current simulations, the deposition efficiency is within $\pm 10\,\%$ of its final value within approximately two vortex shedding periods, indicating that this signal converges relatively quickly as a large enough portion of the database has been sampled. Nevertheless, all results presented here are obtained from injecting particles over ten vortex shedding periods, to ensure that there is no significant dependence on the chosen recurrence path length in the reported values.

IV. RESULTS

A. Recurrence statistics

An illustration of the recurrence matrix for the carrier phase of the case $\mathrm{Re} = 100$ is illustrated in Fig. 5. The main diagonal is supplemented by several off diagonal bands, indicating that there is a clear recurrence pattern in the flow system. The recurrence is very good, with the off diagonal bands going to almost zero, which indicates close agreement between snapshots at different times t and t'.

Similarly, the recurrence matrix for the case $\mathrm{Re} = 10\,000$ is illustrated in Fig. 6. Here, we also see good agreement and a clear off diagonal band structure. However, the recurrence is not exact, as evidenced by the non-zero minima on the off diagonal bands. This can be explained by the turbulence present in higher-Reynolds-number cases. For low Reynolds numbers, the flow is laminar and random velocity fluctuations due to turbulence do not affect the overall flow fields. We then get essentially perfect symmetry in time. As the flow becomes more turbulent at higher Reynolds numbers, random turbulence effects start to influence the flow fields, introducing variations that impair symmetry in time. This gives rise to the non-zero off diagonal bands in the recurrence matrix. Note that since each case has a different normalization factor N, the element values of the two matrices cannot be directly compared. General trends can, on the other hand, be safely compared.

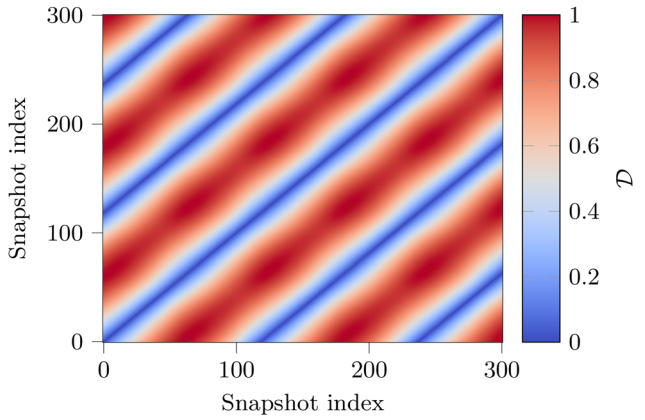

FIG. 5. Recurrence matrix for the CFD database at $\mathrm{Re} = 100$. Plotted are the matrix elements $\mathcal{D}_{ji}$ for the distance metric as defined in Eq. (15).

B. Front-side particle deposition

A validation of the front-side deposition efficiency η_{front} as a function of the particle Stokes number is illustrated in Fig. 7 for the case $\mathrm{Re} = 1685$. The figure presents literature values[1] as the dotted line, and a corresponding DDES simulation performed in this work as the dashed line. It is clear that the DDES simulation corresponds very closely to the simulation results found in the literature.

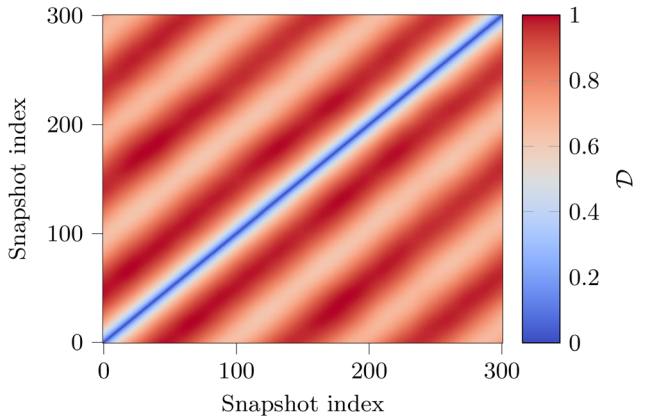

FIG. 6. Recurrence matrix for the CFD database at $\mathrm{Re} = 10\,000$.

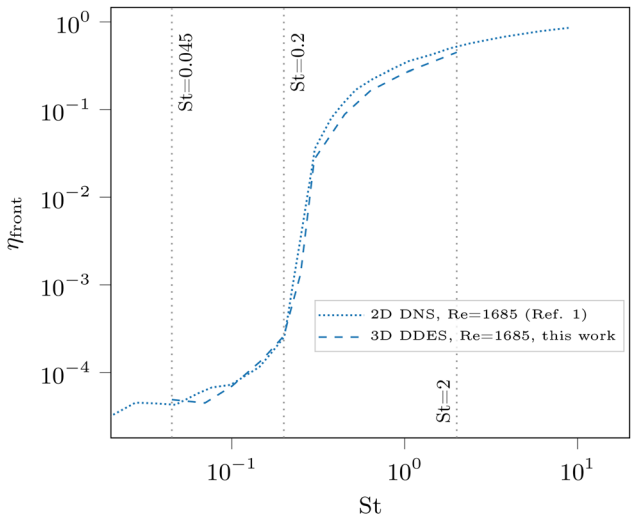

FIG. 7. Front-side particle deposition efficiencies at Re = 1685. The dotted line represents literature values,[1] and the dashed line represents a DDES simulation produced in this work.

As can be seen in the figure, the impact efficiency is heavily dependent on the Stokes number, especially in the transition regime found at intermediate Stokes numbers of $0.2 < \text{St} < 1$. In the high-particle-inertia regime, $\text{St} > 2$, impact efficiencies change slowly with the Stokes number. For the low-particle-inertia regime, $\text{St} < 0.045$, impact efficiencies flatten out.

Beyond $\text{St} > 2$, particle deposition is mainly due to particle inertia. The sudden changes in the carrier phase flow speed and direction are not strong enough to significantly deflect the particle from a collision course with the cylinder, and an impact occurs. In the low-inertia regime, $\text{St} < 0.045$, the suspended particles behave similarly to perfect tracer particles. In this regime, particle-cylinder impacts are mainly due to interception from the finite size of the particles. In other words, the center of mass of the suspended particles follows the fluid streamlines very closely, but the finite radius of the particle means that the surface of the particle impacts the surface of the cylinder when the distance between the particle center of mass and the cylinder surface is within one particle radius. Therefore, the main deposition mechanism in this regime is interception.

With the DDES simulation at Re = 1685 validated against literature in Fig. 7, we present in Fig. 8 a comparison of DDES deposition efficiencies for the Reynolds numbers 1685, 6600, and 10 000 as color-coded dashed lines. The subfigures also illustrate the corresponding deposition efficiencies computed using rCFD flow field extrapolations (solid lines). The DDES and rCFD results match very closely for all Reynolds numbers, indicating that the particle deposition process can be accurately modeled using the rCFD method.

Although there are no disturbances in the upstream flow, the front-side deposition on a cylinder shows strong transient features beyond $\text{Re} > 20$ as the flow is periodically shifted to either side by the vortex shedding, pushing the particles in the boundary layer sideways. The particle interception is therefore enhanced as compared to a stationary boundary layer, and it is expected that this mechanism plays a more prominent role at higher Reynolds numbers.[1] For comparison, a steady-state RANS simulation of the $\text{Re} > 10\,000$ case showed that the front-side deposition efficiencies were accurately captured for $\text{St} > 0.2$, but underpredicted by up to a factor of two for $\text{St} < 0.2$ (not shown). The results in Fig. 7 show that the rCFD approximation of the front-side particle deposition process is very accurate, despite the jumps in the database for the underlying fluid flow fields required by the method. The CFD-to-rCFD comparison is equally excellent for the low-Re cases (not shown), as expected due to the lack of turbulent velocity fluctuations and thus near-perfect recurrence processes at laminar flow conditions.

C. Back-side particle deposition

Backside particle deposition efficiencies as functions of the particle Stokes number are presented in Figs. 9 and 10 for the high-Reynolds

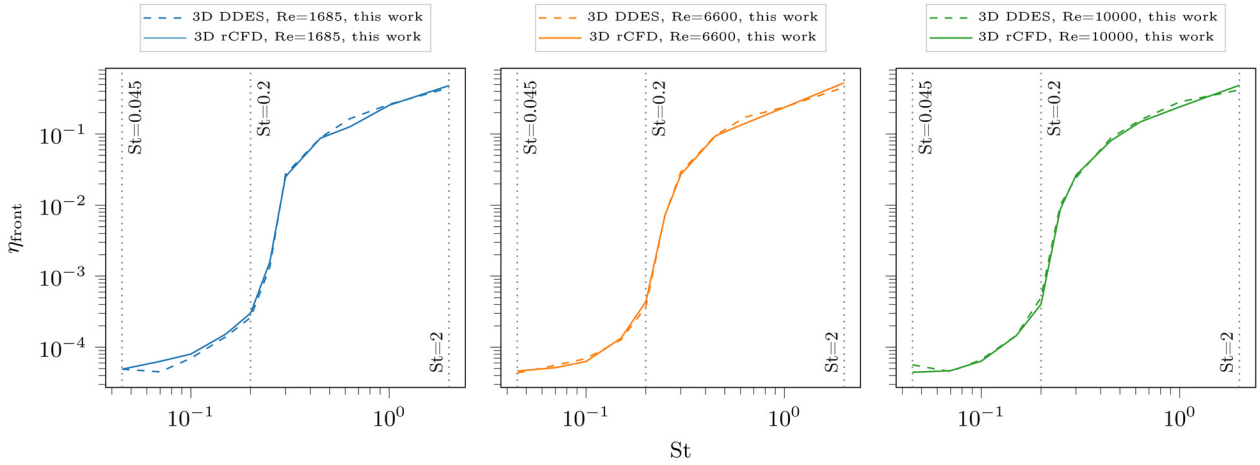

FIG. 8. Deposition efficiencies computed using flow fields from DDES (dashed lines) and rCFD (solid lines) for Re = 1685, 6600, and 10 000. The results are color-coded by Reynolds number. Deposition efficiencies computed using extrapolated flow fields from rCFD correspond closely to the deposition efficiencies computed using conventional DDES flow fields.

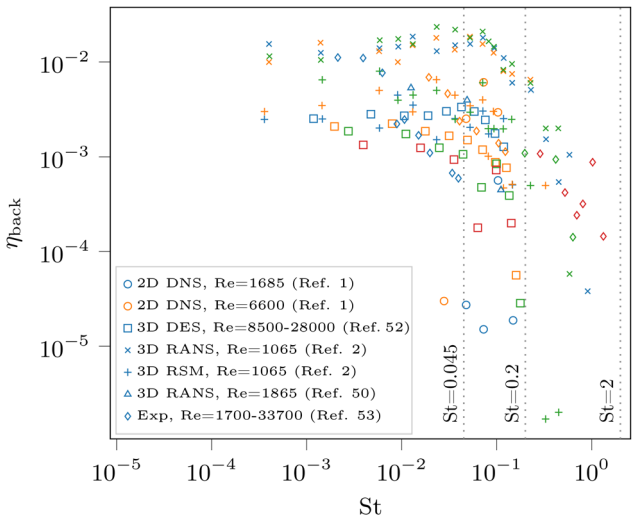

FIG. 9. Literature values of back-side particle deposition efficiencies as a function of the particle Stokes number.[1,2,50,52,53] The graph includes literature values for a square block[50] and a circular disk.[53]

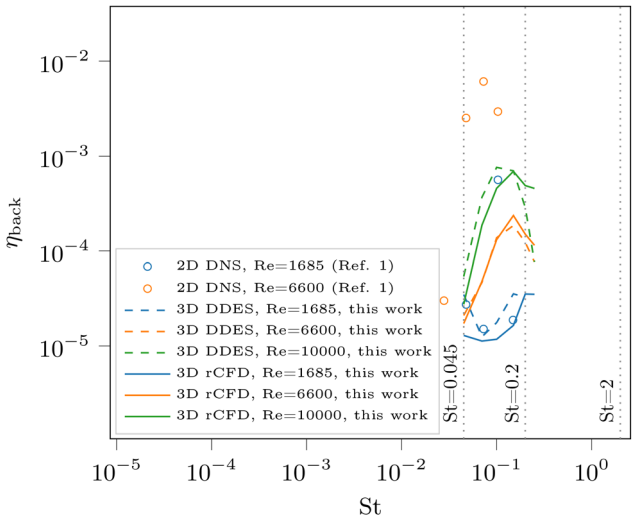

FIG. 10. Back-side deposition results using DDES (dashed lines) and rCFD (solid lines), together with selected DNS results from literature.[1]

number cases Re $\geq$ 1685. Figure 9 presents literature values as markers, whereas Fig. 10 presents results produced in this work, with dashed lines representing DDES results and solid lines representing rCFD results. In Fig. 9, as not much data are available for separate front and backsides, the literature values available for cylinders[1,2,52] are complemented with data for other, similar, geometries such as a square block[50] and a circular disk.[53] For comparison, one set of DNS results from the literature is present in both figures. No backside deposition was observed for the low-Reynolds number cases Re = 20 and Re = 100, in agreement with previous observations.[1]

As is apparent from Fig. 9, the spread in the literature values is very significant, in fact several orders of magnitude. This large spread is due to the difficulty of accurately capturing the strong wake effects present behind the cylinder. It can, however, be noted that the present DDES results in Fig. 10 are within the range of previously predicted backside particle deposition efficiencies for cylinders. The sensitivity of the backside deposition process to accurate descriptions of the fluid turbulence implies that it is vital that the turbulence is either properly resolved or accurately modeled. However, currently available turbulence modeling techniques do not offer sufficient accuracy for particle backside deposition studies. The challenges of modeling can be observed in the presented impact efficiencies from literature, where high-accuracy methods, such as DNS, generally predict lower impact efficiencies than methods with significant degrees of turbulence modeling, such as RANS.[9] Compare, for instance, the impact efficiency results using DNS for flow at Re = 1685[1] and the corresponding RANS results for Re = 1065.[2] In these two cases, we get a difference of about three orders of magnitude for particles with St $\approx$ 0.1. Consequently, properly resolving, not modeling, the turbulence is a well-established way to get reliable particle deposition results on the rear side of a submerged body.

However, one crucial difference between the best available previous high-accuracy DNS study[1] and this work is that the former limits the domain to two dimensions, whereas this work uses three dimensions for cases with Re $\geq$ 1685. In this range of Reynolds numbers, the wake will exhibit increasingly powerful three-dimensional turbulent structures that cannot be captured using a two-dimensional representation. This limitation in the flow description will then lead to different particle deposition results as the three-dimensional turbulent structures become more significant, a fact that is also discussed by the authors of the DNS study.[1] A comparison between the impact efficiencies computed using two-dimensional DNS and the corresponding three-dimensional DDES indicates that there is a clear difference in the calculated deposition efficiencies at Re = 6600. However, for completeness, we also simulated this case with a two-dimensional domain, which yielded results very similar to the 2D-DNS, thus supporting the conclusion that three-dimensional simulations are required for Re $\geq$ 6600.

In Fig. 10, it is immediately clear that the rCFD approximations are very close to the reference CFD simulations, indicating that the approximated solutions still reproduce the main flow structures and fluid–particle interactions. Comparing the CFD and rCFD results from this study to the DNS data from literature,[1] good agreement is found for the case Re = 1685. The sharp peak found in the DNS simulations at St = 0.1 is not observed in our results, but the baseline level is captured accurately. The results for the case Re = 6600 differ between this work and the 2D-DNS,[1] but this difference is attributed to the number of physical dimensions used in representing the system. As previously noted, employing our computational setup, but restricting the domain to a two-dimensional slice, yielded results that are very similar to the 2D-DNS (not shown). The reason behind the significant difference in deposition efficiency is the development of turbulent three-dimensional structures that are not adequately captured by two-dimensional simulations.

The deposition rate on the back of the cylinder has a clear Reynolds number dependency, where more turbulent cases at Re = 6600 and Re = 10 000 correspond to higher particle deposition

efficiencies compared to those for the case Re = 1685. This trend appears as the local boundary layer becomes thinner with increasing Re (so that particles can more easily penetrate it from the wake) in combination with an increase in small-scale vorticity (so that there are eddies that may effectively throw out small particles in directions opposite the main flow).[1] The peak at St ≈ 0.1 is due to the fact that the turbulent eddies characterized by a turnover time equivalent to the particle response time are most likely to effectively throw particles back toward the cylinder.[1] That both these trends are adequately reflected not only by the present DDES but also very well reproduced by the rCFD approximation, is a surprisingly positive observation. Despite the jumps along the rCFD recurrence path, which introduce small artificial glitches in the representation of the turbulent flow field at random intervals, the sensitive particle deposition process on the back side can still be well described.

It is also clear that it is only the low-inertia particles (St≤0.2) that impact the back of the cylinder. Higher-inertia particles pass through without being caught in the wake, preventing them from migrating to the cylinder surface. This critical Stokes number has been found to be 0.125 for potential flow[54] and approximately 0.13 for 2D-DNS.[1]

Note that caution must be exercised when constructing the database. It is vital that the database contains a representative set of flow field snapshots. This is especially important when the length of the database is comparatively short. In this work, we use a fixed database size of three vortex-shedding oscillations, which works well for aerodynamic properties, but may be too short to capture rare events important for particle deposition. An example of this effect can be observed between the two databases in Fig. 11, where the reference DDES simulation is plotted together with two versions of the rCFD approximation. The only difference between the two rCFD approximations is that their corresponding databases were constructed using different segments of the DDES data. Both of the initial flow fields came from the DDES simulation, but at different points in time. Taking different recurrence paths through a single database produces only minor variations to the predicted deposition efficiencies. However, using a different database clearly affects the computed efficiency. This observation indicates that the contents of the database are of major importance when producing accurate rCFD approximations. In the current work, our primary focus is on evaluating the rCFD approach for database sizes sufficient for accurate characterization of the aerodynamics.[51] The database length needed for convergence is therefore not further studied in this work but left for future studies.

Even though rCFD deposition efficiencies are well predicted with the database employed in this work, the deposition locations computed with the rCFD method exhibit some clustering that is not present in the reference DDES results. Figure 12 presents impact locations on the back of the cylinder for the three cases in Fig. 11. The reference DDES results in Fig. 12(a) exhibit an even deposition in a band. The overall trend is similar for the results of the two databases in Figs. 12(b) and 12(c), but with clearer clustering. The clustering patterns are also different between the two databases. These results indicate that the predicted deposition depends on the exact contents of the database, at least for the small database used in this study of only three vortex shedding periods. It is expected that using a larger database would eliminate the effects of database-specific clustering. A small database contains, by definition, only a small number of events that can bring particles to the cylinder surface. A time-extrapolated flow field evolution based

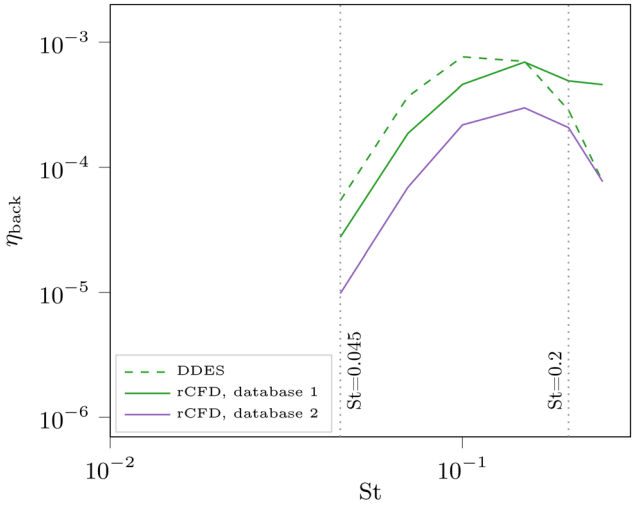

FIG. 11. Comparison of backside deposition results at Re = 10 000 from DDES and two sets of rCFD approximations, using different flow field databases. When using databases of the small size used in this work, it is crucial that the databases contain a representative set of original flow fields, or else the computed deposition results will likely be biased.

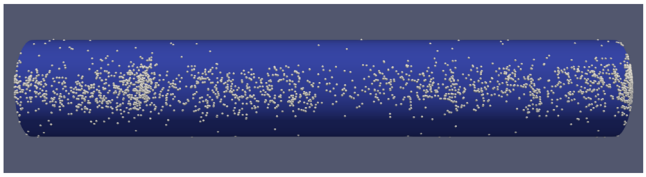

(a) DDES

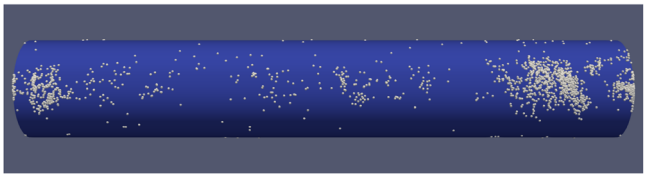

(b) rCFD database 1

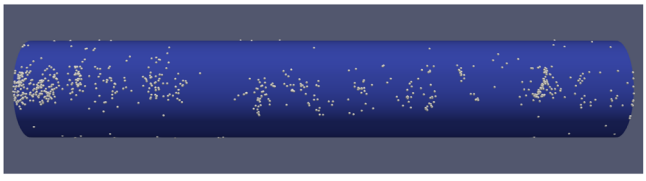

(c) rCFD database 2

FIG. 12. Snapshots of particle deposition locations for the DDES reference (a) and two different rCFD approximations, (b) and (c), using different flow field databases. All cases are for Re = 10 000. The images, here, were captured from a position directly downstream from the cylinder. Particle sizes not to scale. When particle deposition is computed using DDES, the particle deposition locations are evenly spread out across the entire span of the cylinder, with higher particle concentrations at angular locations in the middle, corresponding to directly downstream. Particle deposition calculated using rCFD exhibits more clustering, which is due to the limited size of the flow field database.

only on this small set of events will then naturally lead to clustering. A larger database will contain more fluid flow events that bring particles to the surface, and will therefore reduce the effects of single-event sampling.

Artificial clustering may be a challenge for industrial long-term deposition predictions, as it is less straightforward to differentiate physical from unphysical regions of increased deposition in more complex geometries. At the same time, it is not only the distribution of deposited particles over the surface that is important in assessments of the consequences of deposition but also the temporal deposition characteristics play an important role. For example, if the particles are droplets that enter a liquid film on the surface of the body, not only the location of a deposition event but also its frequency is of utmost importance in accurately determining what will happen next.[55]

To gauge the effects of using the rCFD method on the temporal deposition pattern of particles, a pairwise delay histogram can be computed, such as the one in Fig. 13. In this graph, we plot the delay between each deposition event and all subsequent particle deposition events. The x axis is normalized in terms of the cylinder vortex shedding period and the y axis is normalized so that the integrated quantity sums to unity. It can be seen that the DDES and rCFD temporal deposition patterns show different behavior, with the DDES results exhibiting clear peaks at the first few vortex shedding periods, but then evening out to a steady decay. The rCFD results, on the other hand, show clear periodic correlations at integer multiples of the vortex shedding period, even after long periods of time.

The DDES results show correlation only for the first few vortex shedding periods can be attributed to turbulence-induced variations in the deposition timing. Turbulence causes the vortex shedding process to deviate from the perfectly periodic pattern seen in laminar cases at low Reynolds numbers. The in-sweeps of particles toward the cylinder surface then do not perfectly align with the average vortex shedding period, extinguishing the temporal correlation found for short delays. The clear, long-term correlation found for the rCFD case is instead attributed to the limited number of vortex shedding periods stored in the rCFD database. The method introduces some temporal variation when the recurrence path jumps between snapshots, but these jumps are not very frequent and introduce only small variations in the phase of the oscillation. The individual vortex shedding periods are fixed to the ones recorded in the database. Therefore, the temporal correlation is much stronger and clearly visible, even after several vortex shedding periods.

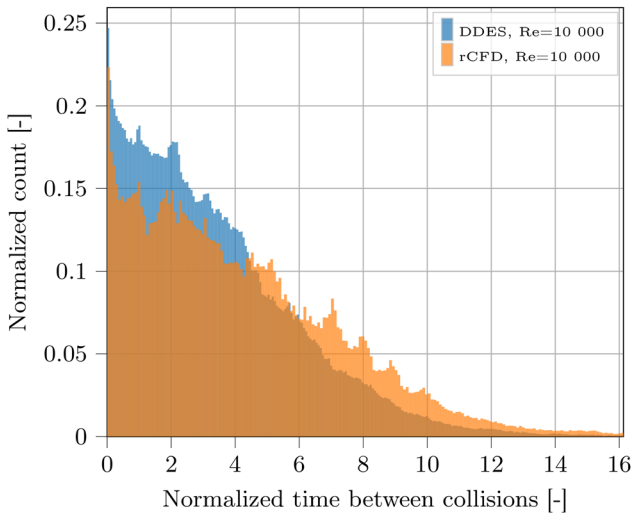

FIG. 13. Normalized delays between each collision and all subsequent collisions.

In conclusion, the rCFD approximation of the deposition process shares the most pronounced characteristics in terms of temporal correlation of deposition events with the full CFD simulation. However, just as a clustering tendency could be discerned in the spatial deposition patterns when the database contains a limited number of vortex shedding periods, an artificially enhanced clustering in time results as the deposition events are necessarily always stitched together using snapshots from the same database. A larger database is expected to further diminish the differences in deposition frequency between CFD and rCFD.

D. Computational performance

A bar chart of the measured wall-clock run time of the different parts of the CFD and rCFD simulations of the Re = 10 000 case is presented in Fig. 14. For the CFD simulations, the measured run time is reported for the fluid and particle solvers, respectively. For the rCFD simulations, the "fluid" time category is the time needed to read the flow fields from database storage. The particle solver times are practically identical due to the two cases using the same particle solver settings. For the rCFD preparation step, the "create database" category represents the time needed for running the fluid solver to generate the CFD database. The second category, "compute recurrence matrix," represents the time needed for calculating the $\mathcal{D}$ matrix.

Looking at the time needed for evolving the fluid, the rCFD approximation is massively faster than the conventional DDES analysis, achieving a speedup of a factor of 942. This is the difference in time needed between calculating the fields and just reading them from

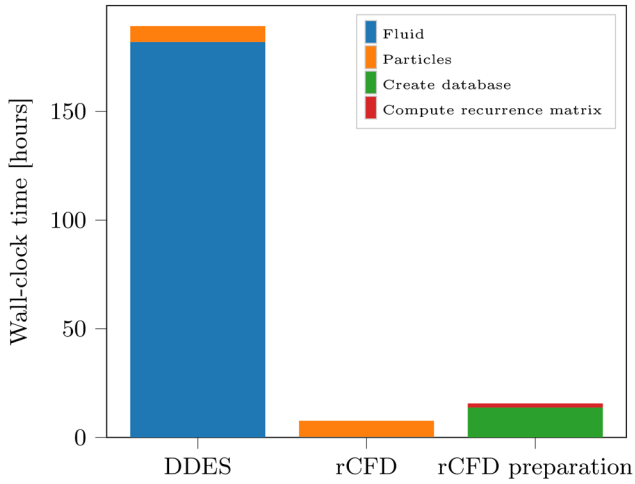

FIG. 14. Computational performance of the DDES and rCFD methods, categorized by type of calculation. The latter method is divided into two phases, "rCFD preparation" and "rCFD." The rCFD preparation phase is only run once for each flow system; the rCFD phase can be run once or several times for parametric studies. Time is measured in wall-clock time.

storage. When also considering the particles, which take the same amount of time to solve in both methods, the rCFD to DDES speedup factor is 25. Finally, comparing the entire rCFD process, including the preparation steps, to the DDES solution time, we get a factor of 8 speedup.

As evidenced by these results, the rCFD process is significantly faster than the conventional DDES solution. The speedups are especially significant if the same fluid flow case can be used, for example, when performing parameter studies on the particles. In these cases, the rCFD preparation step is only run once. The computational cost is reduced to essentially only solving for the particles.

In industrial settings where a fluid flow analysis has already been performed as a previous step of a product development cycle, the rCFD preparation step is reduced to only computing the recurrence matrix. An example would be in predictions of vehicle soiling in the automotive industry, where an aerodynamics characterization step would already have been performed and the particle deposition study can be carried out using an rCFD database of already existing contents.

Furthermore, the computational cost for the particle tracking as performed here depends on the deposition efficiency. As described in Sec. III B, the number of particles tracked for a given Stokes number is inversely proportional to the deposition efficiency, implying that the number of particles tracked to obtain good statistics increases significantly as the deposition efficiency goes to zero. In an industrial setting, it is more likely that the incoming particle size distribution would be part of the problem specification, and the number of injected particles would therefore essentially be independent of deposition efficiency. At the same time, the particles with the lowest deposition efficiency contribute the least to the overall deposition, and they are therefore relatively unimportant in a complete deposition assessment. This effect is even further reinforced by the fact that the lowest deposition efficiencies are seen for the smallest particles, which typically represent an insignificant part of the total mass or volume (as this contribution scales with d_p^3) of a given particle size distribution. In conclusion, the typical particle tracking cost of an industrial deposition study is expected to be much smaller than the cost measured in the present work. The speed-up potential of rCFD as compared to conventional CFD can thus eventually become up to almost three orders of magnitude, depending on the case specifics in question. In this assessment, further speed-up from coarsening the rCFD database in space and time has not been taken into account.

For the simulations presented in this work, fluid and particle solvers ran on eight processor cores. The simulations were performed on a desktop computer with an Intel i7-12700 processor and 64 GB of memory. For the rCFD database, a solid-state drive was used to maximize data throughput rate.

V. CONCLUSIONS AND OUTLOOK

Particle deposition onto objects in turbulent flows is of considerable interest in many applications, such as sensor soiling in the automotive industry, icing on aircraft, and ash build-up in boilers. In these applications, the description of an individual particle deposition event requires access to highly resolved spatiotemporal information of the underlying flow field, which is computationally expensive using conventional CFD. This situation becomes even more challenging as long simulation times are needed to acquire good statistics on the particle deposition process.

In this work, we have shown the potential for using rCFD approximations of the carrier phase flow fields when performing particle tracking and deposition simulations. We have demonstrated that we get very similar deposition efficiencies using both the full CFD method and the new rCFD approximation on the example case of particle–laden flow around a cylinder. This level of accuracy can be achieved with rCFD at only a fraction of the computational cost of a corresponding CFD simulation. Most importantly for predictions of long-term particle deposition, the rCFD technique enables infinite time extrapolation without any deterioration in quality of the flow field description. The discrete temporal jumps introduced by the method have been shown to have a very limited influence on the obtained results, for the complete range of Stokes numbers used, and they therefore do not jeopardize the long-term accuracy of the approach.

For all the Reynolds numbers investigated, the front-side deposition – which dominates the overall deposition onto the cylinder – was very well predicted using rCFD. This result implies that the front-side boundary-layer fluctuations induced by the main vortex shedding are accurately represented in the rCFD approximation. The much less effective back-side deposition could also be well predicted, although some dependence on the database length was discerned for the highest Reynolds number (Re = 10 000). This observation highlights the difficulties in obtaining robust long-term predictions for quantities that are controlled by rare events when using a flow-field database of limited size. Any reduced-order model or machine-learning approach trained on the same CFD dataset should be expected to perform similarly in this regard.

While the rCFD method only provides time extrapolations for the flow at a single Reynolds number, it would be possible to use several rCFD databases at several Reynolds numbers and then interpolate results between them. Previous work in the literature has demonstrated the viability of an interpolation approach between databases with different inlet velocities for the case of gas–solid fluidized beds.[56] For the case of particle deposition on a cylinder, it would be possible to use rCFD to generate particle deposition patterns, and then train a machine learning model to interpolate deposition behavior at intermediate Reynolds numbers. This would extend the applicability of the time-extrapolations provided by the rCFD method.

The results presented in this study show that the rCFD method is capable and efficient in approximating the full CFD solution for Lagrangian particle deposition studies on a cylinder geometry. Future studies should focus on evaluating the rCFD approach for more complex geometries, and on investigating the possibilities to obtain even better computational performance from optimizing the spatiotemporal resolution of the rCFD database by coarse-graining.

ACKNOWLEDGMENTS

This work has been financed by the Swedish Agency for Innovation Systems (Vinnova) under the Strategic Vehicle Research and Innovation collaborative program (FFI) via Grant Agreement No. Dnr 2021-05061. The computations were enabled by resources provided by the National Academic Infrastructure for Supercomputing in Sweden (NAISS), partially funded by the Swedish Research Council through Grant Agreement No. 2022-06725.

AUTHOR DECLARATIONS

Conflict of Interest

The authors have no conflicts to disclose.

Author Contributions

Johannes Hansson: Conceptualization (equal); Formal analysis (equal); Methodology (equal); Resources (equal); Software (equal); Validation (equal); Visualization (equal); Writing – original draft (equal). **Thomas Lichtenegger:** Conceptualization (equal); Methodology (equal); Writing – review & editing (equal). **Stefan Pirker:** Conceptualization (equal); Methodology (equal); Writing – review & editing (equal). **Srdjan Sasic:** Conceptualization (equal); Methodology (equal); Supervision (equal); Writing – review & editing (equal). **Henrik Ström:** Conceptualization (equal); Funding acquisition (equal); Methodology (equal); Project administration (equal); Resources (equal); Software (equal); Supervision (equal); Writing – review & editing (equal).

DATA AVAILABILITY

The data that support the findings of this study are available from the corresponding author upon reasonable request.

REFERENCES

[1]N. E. L. Haugen and S. Kragset, "Particle impaction on a cylinder in a crossflow as function of Stokes and Reynolds numbers," J. Fluid Mech. **661**, 239–261 (2010).
[2]R. Weber, N. Schaffel-Mancini, M. Mancini, and T. Kupka, "Fly ash deposition modelling: Requirements for accurate predictions of particle impaction on tubes using RANS-based computational fluid dynamics," Fuel **108**, 586–596 (2013).
[3]G. Lecrivain, R. Rayan, A. Hurtado, and U. Hampel, "Using quasi-DNS to investigate the deposition of elongated aerosol particles in a wavy channel flow," Comput. Fluids **124**, 78–85 (2016).
[4]E. M. A. Frederix, A. K. Kuczaj, M. Nordlund, A. E. P. Veldman, and B. J. Geurts, "Eulerian modeling of inertial and diffusional aerosol deposition in bent pipes," Comput. Fluids **159**, 217–231 (2017).
[5]U. Kleinhans, C. Wieland, F. J. Frandsen, and H. Spliethoff, "Ash formation and deposition in coal and biomass fired combustion systems: Progress and challenges in the field of ash particle sticking and rebound behavior," Prog. Energy Combust. Sci. **68**, 65–168 (2018).
[6]A. P. Gaylard, K. Kirwan, and D. A. Lockerby, "Surface contamination of cars: A review," Proc. Inst. Mech. Eng., Part D: J. Automob. Eng. **231**(9), 1160–1176 (2017).
[7]A. Soldati and C. Marchioli, "Physics and modelling of turbulent particle deposition and entrainment: Review of a systematic study," Int. J. Multiphase Flow **35**(9), 827–839 (2009).
[8]P. Moin and J. Kim, "Numerical investigation of turbulent channel flow," J. Fluid Mech. **118**, 341–377 (1982).
[9]C. Talebmoustaph, P. Fede, O. Simonin, M. Pallud, and P. Maheshwari, "Stochastic Lagrangian wall deposition model for RANS prediction of deposition in turbulent gas–solid flows," Int. J. Multiphase Flow **178**, 104900 (2024).
[10]P. A. Durbin, "Some recent developments in turbulence closure modeling," Annu. Rev. Fluid Mech. **50**, 77–103 (2018).
[11]K. Duraisamy, G. Iaccarino, and H. Xiao, "Turbulence modeling in the age of data," Annu. Rev. Fluid Mech. **51**, 357–377 (2019).
[12]K. Choudhary, K. A. Krishnaprasad, S. Pandey, N. Zgheib, J. S. Salinas, M. Y. Ha, and S. Balachandar, "Effectiveness of RANS in predicting indoor airborne viral transmission: A critical evaluation against LES," Comput. Fluids **256**, 105845 (2023).
[13]L. Davidson, see https://www.tfd.chalmers.se/~lada/postscript_files/solids-and-fluids_turbulent-flow_turbulence-modelling.pdf for "Fluid mechanics, turbulent flow and turbulence modeling," Chalmers University of Technology, 2024.
[14]S. Dong and G. E. Karniadakis, "DNS of flow past a stationary and oscillating cylinder at Re=10000," J. Fluids Struct. **20**(4), 519–531 (2005).
[15]P. R. Spalart, "Strategies for turbulence modelling and simulations," Int. J. Heat Fluid Flow **21**(3), 252–263 (2000).
[16]A. Ceci, S. Pirozzoli, J. Romero, M. Fatica, R. Verzicco, and P. Orlandi, "Direct numerical simulations of turbulent pipe flow at high Reynolds number," Phys. Rev. Fluids **7**, 110510 (2022).
[17]J. Smagorinsky, "General circulation experiments with the primitive equations," Mon. Weather Rev. **91**, 99 (1963).
[18]Y. Zhiyin, "Large-eddy simulation: Past, present and the future," Chin. J. Aeronaut. **28**(1), 11–24 (2015).
[19]O. Reynolds, "On the dynamical theory of incompressible viscous fluids and the determination of the criterion," Philos. Trans. R. Soc. London, Ser. A **186**, 123–164 (1895).
[20]G. Alfonsi, "Reynolds-averaged Navier–Stokes equations for turbulence modeling," Appl. Mech. Rev. **62**, 040802 (2009).
[21]P. Spalart, W.-H. Jou, M. Strelets, and S. Allmaras, "Comments on the feasibility of LES for wings, and on a hybrid RANS/LES approach," in *Advances in DNS/LES: Direct Numerical Simulation and Large Eddy Simulation* (Greyden Press, Columbus, 1997), pp. 137–148.
[22]T. Eidevåg, M. Eng, D. Kallin, J. Casselgren, Y. Bharadhwaj, T. S. Bangalore Narahari, and A. Rasmuson, "Snow contamination of simplified automotive bluff bodies: A comparison between wind tunnel experiments and numerical modeling," SAE Int. J. Adv. Curr. Pract. Mobility **4**, 2120–2134 (2022).
[23]A. Kabanovs, A. Garmory, M. Passmore, and A. Gaylard, "Computational simulations of unsteady flow field and spray impingement on a simplified automotive geometry," J. Wind Eng. Ind. Aerodyn. **171**, 178–195 (2017).
[24]P. R. Spalart, "Young-person's guide to detached-eddy simulation grids," Technical Report No. NASA/CR-2001-211032, 2001.
[25]P. R. Spalart, "Detached-eddy simulation," Annu. Rev. Fluid Mech. **41**, 181–202 (2009).
[26]P. Justesen, "A numerical study of oscillating flow around a circular cylinder," J. Fluid Mech. **222**, 157–196 (1991).
[27]C. H. K. Williamson, "Vortex dynamics in the cylinder wake," Annu. Rev. Fluid Mech. **28**, 477–539 (1996).
[28]G. Berkooz, P. Holmes, and J. L. Lumley, "The proper orthogonal decomposition in the analysis of turbulent flows," Annu. Rev. Fluid Mech. **25**, 539–575 (1993).
[29]P. J. Schmid, "Dynamic mode decomposition and its variants," Annu. Rev. Fluid Mech. **54**, 225–254 (2022).
[30]A. Solera-Rico, C. S. Vila, M. Gómez-López, Y. Wang, A. Almashjary, S. T. M. Dawson, and R. Vinuesa, "β-variational autoencoders and transformers for reduced-order modelling of fluid flows," Nat. Commun. **15**, 1361 (2024).
[31]T. Lichtenegger and S. Pirker, "Recurrence CFD – A novel approach to simulate multiphase flows with strongly separated time scales," Chem. Eng. Sci. **153**, 394–410 (2016).
[32]S. Abbasi, S. Pirker, and T. Lichtenegger, "Application of recurrence CFD (rCFD) to species transport in turbulent vortex shedding," Comput. Fluids **196**, 104348 (2020).
[33]T. Lichtenegger, S. Abbasi, and S. Pirker, "Transport in turbulent, recurrent flows: Time-extrapolation and statistical symmetrization," Chem. Eng. Sci. **259**, 117795 (2022).
[34]P. R. Spalart, S. Deck, M. L. Shur, K. D. Squires, M. K. Strelets, and A. Travin, "A new version of detached-eddy simulation, resistant to ambiguous grid densities," Theor. Comput. Fluid Dyn. **20**(3), 181–195 (2006).
[35]See https://www.openfoam.com/news/main-news/openfoam-v2312 for "OpenFOAM v2312" (2023).
[36]M. S. Gritskevich, A. V. Garbaruk, J. Schütze, and F. R. Menter, "Development of DDES and IDDES formulations for the k-ω shear stress transport model," Flow Turbul. Combust. **88**(3), 431–449 (2012).
[37]C. Norberg, "Fluctuating lift on a circular cylinder: Review and new measurements," J. Fluids Struct. **17**(1), 57–96 (2003).
[38]B. M. Sumer and J. Fredsøe, *Hydrodynamics around Cylindrical Structures*, Revis Edition (World Scientific, 2006).

[39]F. P. Bretherton, "The motion of rigid particles in a shear flow at low Reynolds number," J. Fluid Mech. **14**, 284–304 (1962).
[40]P. Shi and R. Rzehak, "Lift forces on solid spherical particles in wall-bounded flows," Chem. Eng. Sci. **211**, 115264 (2020).
[41]A. Putnam, "Integratable form of droplet drag coefficient," ARS J. **31**(10), 1467–1468 (1961).
[42]A. A. Amsden, P. J. O'Rourke, and T. D. Butler, "KIVA-II: A computer program for chemically reactive flows with sprays," Technical Report No. LA-11560-MS (Los Alamos National Lab., 1989).
[43]L. Tian and G. Ahmadi, "Particle deposition in turbulent duct flows—Comparisons of different model predictions," J. Aerosol Sci. **38**, 377–397 (2007).
[44]S. Pirker and T. Lichtenegger, "Efficient time-extrapolation of single- and multiphase simulations by transport based recurrence CFD (rCFD)," Chem. Eng. Sci. **188**, 65–83 (2018).
[45]S. Bouhairie and V. H. Chu, "Two-dimensional simulation of unsteady heat transfer from a circular cylinder in crossflow," J. Fluid Mech. **570**, 177–215 (2007).
[46]H. Schlichting and K. Gersten, *Boundary-Layer Theory* (Springer, 2000).
[47]U. Kleinhans, C. Wieland, S. Babat, and H. Spliethoff, "Large eddy simulation of a particle-laden flow around a cylinder: Importance of thermal boundary layer effects for slagging and fouling," Fuel **241**, 585–606 (2019).
[48]L. Davidson, "Large eddy simulations: How to evaluate resolution," Int. J. Heat Fluid Flow **30**(5), 1016–1025 (2009).
[49]V.-T. Nguyen and H. H. Nguyen, "Detached eddy simulations of flow induced vibrations of circular cylinders at high Reynolds numbers," J. Fluids Struct. **63**, 103–119 (2016).
[50]A. Li, G. Ahmadi, R. G. Bayer, and M. A. Gaynes, "Aerosol particle deposition in an obstructed turbulent duct flow," J. Aerosol Sci. **25**(1), 91–112 (1994).
[51]S. Jakirlic, L. Kutej, P. Unterlechner, and C. Tropea, "Critical assessment of some popular scale-resolving turbulence models for vehicle aerodynamics," SAE Int. J. Passenger Cars Mech. Syst. **10**, 235–250 (2017).
[52]X. Li, H. Zhou, and K. Cen, "Influences of various vortex structures on the dispersion and deposition of small ash particles," Fuel **87**(7), 1379–1382 (2008).
[53]J. H. Vincent and W. Humphries, "The collection of airborne dusts by bluff bodies," Chem. Eng. Sci. **33**(8), 1147–1155 (1978).
[54]S. K. Friedlander, *Smoke, Dust and Haze: Fundamentals of Aerosol Behavior* (Wiley, 1977).
[55]Y. Zhang, G. Vinay, A. Poubeau, and Q. V. Hoang, "Development of a hybrid Lagrange–Euler transition model for the film formation and dynamics of an impinging liquid spray," Comput. Fluids **251**, 105756 (2023).
[56]T. Lichtenegger, P. Kieckhefen, S. Heinrich, and S. Pirker, "Dynamics and long-time behavior of gas–solid flows on recurrent-transient backgrounds," Chem. Eng. J. **364**, 562–577 (2019).