When Particle Shifting is Expected to Improve Convergence of Lagrangian Smoothed Particle Hydrodynamics

Lennart Beck
dive solutions GmbH,
Berlin, Germany
beck@dive-solutions.de

Ivo F. Sbalzarini
Technische Universität Dresden, Faculty of Computer Science,
Max Planck Institute of Molecular Cell Biology and Genetics,
Center for Systems Biology Dresden,
Dresden, Germany,
ivo.sbalzarini@tu-dresden.de

***Abstract*—We introduce a metric to determine when particle shifting is expected to improve the convergence of weakly compressible Lagrangian SPH. Shifting regularizes particle positions by biasing them toward a more homogeneous distribution. However, shifting also introduces an additional discretization error. To quantify these two opposing effects, we analyze the two-dimensional Taylor-Green case and show that particle shifting is advantageous when there is strong-enough distortion in the particle distribution. We propose a dimensionless metric of particle distortion to predict this break-even point, which we show to behave qualitatively similar across simulation cases. Since the proposed metric can be computed in a running simulation, it can be used to control the onset of shifting.**

I. Introduction

In Lagrangian particle methods that use a single global interpolation kernel across the entire computation domain, like Smoothed Particle Hydrodynamics (SPH), convergence depends on the regularity of the particle distribution [6]. In typical SPH applications to fluid flow, the main causes of particle distortion are shear velocities and locally compressible expansion. Since particles follow Lagrangian material trajectories, an initially uniform particle distribution distorts over the course of a simulation, eventually limiting the convergence of the method. Therefore, maintaining sufficient regularity of the particle distribution throughout an SPH simulation is necessary to ensure convergence.

One approach to regularizing the particle distribution is the popular particle shifting technique (PST) [11]. Its straightforward implementation makes it applicable to many use cases of SPH, including free-surface flows. Particle shifting regularizes the particle distribution by "slightly" shifting the particles in the opposite direction of the local particle density gradient. When applied with sufficient frequency, shifting maintains a more regular particle distribution. However, the shifted particles no longer follow Lagrangian trajectories, introducing an additional discretization error. It is therefore unclear how shifting affects convergence and if or when the shifting error is amortized by the gained regularity in the particle distribution.

The error introduced by shifting has been analyzed by second-order Taylor expansion of the physical particle quantities around their shifted positions [11, 2]. In weakly compressible SPH, the shifting magnitude is limited by the small time steps required to resolve the artificial speed of sound. The shifting error is therefore usually assumed to be negligible [8]. The conditions under which this assumption holds are not well studied, though.

Here, we systematically study the validity of the shifting assumption and the effect of shifting on the numerical convergence of the SPH method. We show that while particle shifting increases the simulation error for short times, it becomes advantageous beyond a break-even time, which we characterize. Since particle distortion is caused by advection, we re-normalize time in terms of the Reynolds number, and we characterize the regularity of the particle distribution by an order parameter. Using this dimensionless quantity, we provide a guideline for when shifting is expected to improve convergence of SPH.

II. The discretization model

We consider the continuity equation in its compressible form to permit small variations in density. The pressure can directly be derived from an equation of state. To prevent larger density variations, an artificial speed of sound c_0 is introduced. This artificial speed of sound is chosen to depend on the expected flow field as

$$c_0 = 10 \max\left(v_{\max}, \sqrt{\frac{p_{\max}}{\rho_0}}\right), \tag{1}$$

with the expected maximum velocity $v_{\max}$ and the maximum pressure $p_{\max}$ in the simulation, and the initial reference density of the fluid ρ_0. Setting c_0 according to Eq. (1) limits density fluctuations to about 1% of ρ_0. In addition, the continuity equation is complemented by an artificial diffusive term, which dampens the spurious pressure oscillations occurring when evaluating pressure and velocity at the same locations in space. We use the formulation proposed by Molteni and Colagrossi [3] with artificial diffusivity $\delta = 0.1$. Momentum diffusion is approximated using the Morris viscosity operator

[4]. Altogether, this results in the following system of ODEs governing the positions and properties of each particle i:

$$\begin{cases} \dfrac{\mathrm{d}\rho_i}{\mathrm{d}t} & = \rho_i \displaystyle\sum_j \frac{m_j}{\rho_j} (\mathbf{v}_i - \mathbf{v}_j) \cdot \nabla W_{h,ij} \\ & \quad + 2c_0 \delta h \displaystyle\sum_j \frac{m_j}{\rho_j} \frac{(\rho_i - \rho_j)}{\|\mathbf{x}_i - \mathbf{x}_j\|} W'_{h,ij} \\ \dfrac{\mathrm{d}\mathbf{v}_i}{\mathrm{d}t} & = -\dfrac{1}{\rho_i} \displaystyle\sum_j m_j \left(\frac{p_i}{\rho_i^2} + \frac{p_j}{\rho_j^2} \right) \nabla W_{h,ij} \\ & \quad + 2\nu_i \displaystyle\sum_j \frac{m_j}{\rho_j} \frac{(\mathbf{v}_i - \mathbf{v}_j)}{\|\mathbf{x}_i - \mathbf{x}_j\|} W'_{h,ij} \\ \dfrac{\mathrm{d}\mathbf{x}_i}{\mathrm{d}t} & = \mathbf{v}_i \\ p_i & = \dfrac{c_0^2 \rho_0}{\gamma} \left(\left(\dfrac{\rho_i}{\rho_0}^{\gamma} \right) - 1 \right). \end{cases} \tag{2}$$

For the SPH kernel, we use the shorthand notation $W_{h,ij} := W_h(\mathbf{x}_i - \mathbf{x}_j)$, and we use the fifth-order Wendland kernel [10] throughout this study:

$$W_h(\mathbf{x}) = \begin{cases} \alpha_d \left(1 - \dfrac{\|\mathbf{x}\|}{2h} \right)^4 \left(1 + \dfrac{2\|\mathbf{x}\|}{2h} \right) & \text{if } \dfrac{1 + 2\|\mathbf{x}\|}{2h} \leq 2 \\ 0 & \text{else}. \end{cases} \tag{3}$$

The normalization factor α_d is necessary to ensure zeroth-order consistency of the kernel and, consequently, depends on the space dimension d. For two and three dimensions, the values of α_d are:

$$\alpha_2 = \frac{7}{4\pi h^2}, \quad \alpha_3 = \frac{21}{16\pi h^3}. \tag{4}$$

The kernel gradient for the fifth-order Wendland kernel is derived from the first derivative:

$$\nabla W_{h,ij} = \frac{\mathbf{x}_i - \mathbf{x}_j}{\|\mathbf{x}_i - \mathbf{x}_j\|} W'_{h,ij}. \tag{5}$$

The smoothing length h defines the kernel radius and thus how many neighbors any particle interacts with. It is set relative to the particle diameter Δx and the smoothing ratio σ as:

$$h = \sigma \Delta x. \tag{6}$$

The parameter γ in the equation of state is set to $\gamma = 7$ throughout this paper. We solve the set of ODEs in Eq. (2) using the explicit Euler scheme to discretize time. Thus, physical properties are advanced in time as

$$\begin{cases} \mathbf{v}_i^{(t+\Delta t)} & = \mathbf{v}_i^{(t)} + \Delta t \dfrac{d\mathbf{v}_i^{(t)}}{dt} \\ \rho_i^{(t+\Delta t)} & = \rho_i^{(t)} + \Delta t \dfrac{d\rho_i^{(t)}}{dt} \\ \mathbf{x}_i^{(t+\Delta t)} & = \mathbf{x}_i^{(t)} + \Delta t \mathbf{v}_i^{(t)}. \end{cases} \tag{7}$$

To analyze convergence in space, the same time-step size $\Delta t = 10^{-5}$ is used for all spatial resolutions. This is small enough to not dominate the error even for the finest resolution tested. Smaller time steps have been tested and did not show a significant effect on the results.

a) Particle Shifting: We use the particle shifting formulation by Lind et al. [2]. This adds to the particle movements from Eq. (7) a small non-physical displacement, referred to as *shift*. The direction of the shift is given by the negative concentration gradient. The concentration C is defined using the SPH formulation for gradients:

$$\nabla C_i = \sum_j \frac{m_j}{\rho_j} \nabla W_{h,ij}. \tag{8}$$

The negative concentration gradient then points toward under-resolved areas, and particles will be shifted in this direction. The magnitude of the shift is controlled by the CFL number of the time stepper and the smoothing length h as:

$$\tilde{\mathbf{x}}_i = -0.4 C_{\text{CFL}} h^2 \nabla C_i. \tag{9}$$

These parameters ensure that the shift magnitude is small compared to the advective displacement. Additionally, they ensure by construction that shifting converges with second order in h, as the magnitude of the shifting displacement scales with $O(h^2)$. In the literature, an upper limit for the shifting magnitude is often additionally imposed [2, 5]. Here, we do not do so, but we verified that none of the limits were ever reached during our benchmarks. The modified position update in the explicit Euler scheme of Eq. (7) with shifting then becomes:

$$\mathbf{x}_i^{(t+\Delta t)} = \mathbf{x}_i^{(t)} + \Delta t \mathbf{v}_i^{(t)} + \tilde{\mathbf{x}}_i^{(t)}. \tag{10}$$

In the original shifting paper [2], the particle properties were corrected after shifting by a second-order Taylor expansion to approximate the velocity and pressure at the shifted positions. In the weakly compressible SPH literature, however, such correction is usually omitted because time steps are much smaller than in incompressible Smoothed Particle Hydrodynamics (ISPH). In either case, particle shifting introduces an additional error. In the following, the error induced by shifting is compared with the expected improvement in convergence.

III. Convergence analysis

We quantify the convergence of weakly compressible SPH with and without shifting and determine the degree of particle distortion beyond which shifting is beneficial in the sense that the shifting errors are amortized by the gain in convergence.

The numerical experiment: As a benchmark case, we consider the two-dimensional incompressible Taylor-Green vortex in the domain $\Omega = [-0.5, 0.5]^2$ with periodic boundary conditions in both directions. For this case, the analytical solution is:

$$\begin{aligned} \mathbf{v}_0(\mathbf{x}, t) &= U \begin{bmatrix} \sin(2\pi \mathbf{x}_1) \cos(2\pi \mathbf{x}_2) \\ \cos(2\pi \mathbf{x}_1) \sin(2\pi \mathbf{x}_2) \end{bmatrix} F(t) \\ p_0(\mathbf{x}, t) &= \frac{\rho_0}{4} \left(\cos(4\pi \mathbf{x}_1) + \cos(4\pi \mathbf{x}_2) \right) F^2(t), \end{aligned} \tag{11}$$

with $F(t) := \mathrm{e}^{-8\pi^2 \nu t}$ [9]. This allows us to accurately quantify the convergence of the numerical error. The initial condition

is taken to be the analytical solution at time $t = 0$. The case parameters are set to $U = 1$ and $\rho_0 = 1$. The characteristic length of the Taylor-Green vortex is set to $L = 1$, as is the width of the domain Ω. The kinematic viscosity ν is used to control the Reynolds number

$$Re = \frac{LU}{\nu} = \frac{1}{\nu}\,. \tag{12}$$

Particle methods often use a Cartesian initialization of particle positions. The particle distribution that minimizes the shifting gradient flow, however, is not Cartesian. This typically leads to large initial shifts. We therefore equilibrate the initial particle distribution by performing several rounds of shifting until a threshold on the maximum concentration gradient is reached. The initial condition is then discretized on the resulting set of particles. In this way, shifting during the simulation is determined by the Lagrangian particle displacements from the equations of motion, and not by the initial particle placement. Pre-equilibration of particle positions also improves comparability across different resolutions, as otherwise initial shifting would vary between resolutions.

Convergence Analysis: We quantify convergence in terms of the following norms of the errors e_i on the particles i:

$$\begin{aligned} \|e\|_{L_1} &= \Delta x^d \sum_i \|e_i\|\,, \\ \|e\|_{L_2} &= \left(\Delta x^d \sum_i \|e_i\|^2\right)^{\frac{1}{2}}\,, \\ \|e\|_{L_\infty} &= \max_i \|e_i\|\,. \end{aligned} \tag{13}$$

The error is computed in the L_2 and L_∞ norms of the velocity vectors

$$\|e_i(t)\|_{2,\infty} = \|\mathbf{v}_i(t) - \mathbf{v}_0(\mathbf{x}_i, t)\|_{L_2, L_\infty} \tag{14}$$

against the analytical solution $\mathbf{v}_0(t)$ from Eq. (11). In order to evaluate the impact of particle shifting on the accuracy of the simulation results, we analyze the convergence behavior for four different Reynolds numbers. For the two lower Reynolds numbers (50 and 100), the viscosity rapidly dissipates the initial momentum of the vortex, and the distortion of the particle distribution remains small. For the higher Reynolds numbers (1000 and 2000), momentum dissipation has a lesser influence, and particles move further causing higher distortion. In addition, an effect specific to the SPH model used here appears for Reynolds numbers above 200: the particle distribution starts to form voids at the centers of the vortices. This is caused by the low (sometimes negative) pressures that can occur in these areas and has been extensively studied [7]. This effect disappears when shifting is enabled. Distortion of vortex cores is not driven by shear forces, as particle speeds in those regions are relatively small. Instead, the observed distortion can be interpreted as the result of an expansion process. Even though the Taylor-Green vortex, from a physical perspective, only models effects of shear forces, the numerical errors in the SPH model cause the discretized Taylor-Green vortex to also contain expansion effects. This renders the present test case well-suited for quantifying the effect of particle shifting over a range of lower and higher Reynolds numbers.

Theoretical error analysis of SPH reveals two components of the numerical error: First, the approximation error due to the moments of the differential operators only being reproduced up to a certain order. Second, the discretization error stemming from the fact that the moment conditions are enforced in the continuous but the operator is applied to a discrete domain. The discretization error is independent of particle density and therefore becomes limiting eventually at high-enough resolutions [6].

These two error components are visible in Fig. 1, showing convergence plots in the L_2 norm for the Taylor-Green vortex with and without shifting at different Reynolds numbers and time $\tilde{t} = t\frac{U}{L} = 0.1$, i.e., after the first 10^4 time steps. The error without shifting is smaller in almost all cases. For the lower two Reynolds numbers, the errors meet at the same the characteristic error plateau [6] for higher resolutions. This indicates that at high resolutions the discretization error dominates the shifting error. For the two higher Reynolds numbers, however, the error with shifting drops below the error without shifting for the highest resolution. For $Re = 2000$, the error without shifting starts to increase again for the highest resolution, as particle separation occurs at the vortex cores. The plateau with shifting appears later and at a lower error level, showing a first benefit from the regularized particle distribution.

At $\tilde{t} = 1$ (10^5 time steps) the situation is different, as shown in Fig. 2. The error with shifting is now lower in most cases. Without shifting, also, the convergence order for the higher Reynolds numbers is reduced, whereas it remains quadratic with shifting. The error plateau is not observed when using shifting for the higher Reynolds numbers. This has two possible explanations: Either the improved particle distribution created by shifting reduces the discretization error and thus pushes the plateau beyond the resolutions tested, or the error of the SPH operators is not the dominant error any more, but rather the quadratic shifting error. Since the early-time convergence analysis on pre-shifted particle distributions showed a plateau, the latter explanation is more likely. This means that the second-order error introduced by shifting starts to dominate over the SPH errors at later times and higher Reynolds numbers. This is consistent with the observation in Fig. 2 that the error grows for larger Reynolds numbers.

We verify the results also for a smaller kernel with $\frac{h}{\Delta x} = 2$. The discrepancy between the two kernel sizes is largest for the earliest time point. With the smaller kernel, the results with and without shifting already differ at the earliest time. Still, as with the larger kernel, shifting is not beneficial to the error for low levels of particle distortion. At time $\tilde{t} = 1.0$, shifting again reduces the error for all resolutions and Reynolds numbers, like it did for the larger kernel. As expected, though, the break-even point for shifting occurs earlier for smaller kernels that introduce less inter-particle smoothing and are thus more

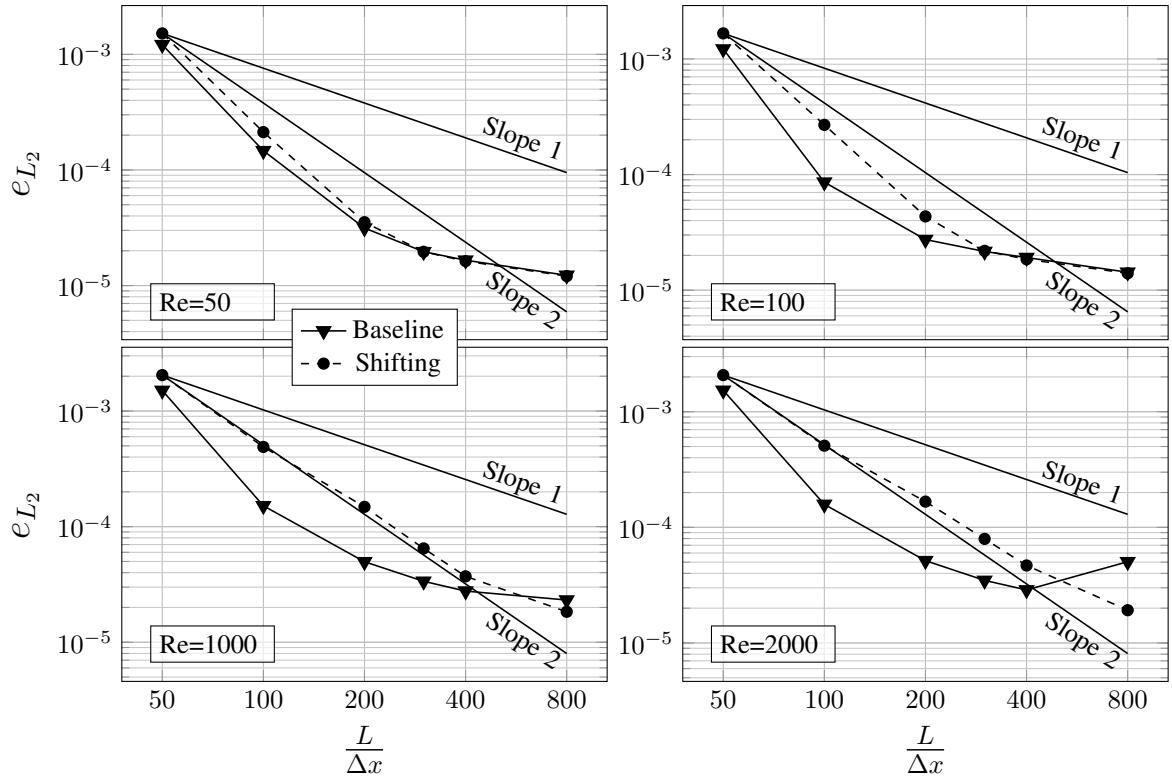

Figure 1: Convergence with and without shifting in the L_2 norm for different Reynolds numbers at time $\tilde{t} = 0.1$ for $\frac{h}{\Delta x} = 6$.

vulnerable to distortions in the particle distribution.

Taken together, these results confirm that the SPH model in Eq. (2) shows the expected convergence behavior. Additionally, they show that shifting initially increases the velocity error when particle distortion is still low, but eventually becomes beneficial for higher distortion.

IV. Analysis of the shifting error

Beyond the individual time points considered so far, we analyze how the two competing error terms—the shifting error and the error from particle distortion—evolve over time. This will allow us to quantify the break-even point for shifting by considering the error difference

$$\Delta E(t) = e^{-}_{L_\infty}(t) - e^{+}_{L_\infty}(t)\,, \tag{15}$$

where $+$, $-$ represent the errors with and without shifting, respectively. Negative values of ΔE indicate that shifting is not beneficial to the overall error, whereas positive ΔE indicate that shifting reduces the overall error. The time evolution of ΔE across Reynolds numbers and resolutions is plotted in Fig. 3. It confirms the conclusions drawn from the individual time points in the previous section. At early times, shifting always increases the overall error, as shown by the negative ΔE values. After a certain break-even time, corresponding to a certain accumulated particle distortion, the sign of ΔE flips in all cases, and shifting starts to reduce the overall error. The break-even point is always followed by a sharp increase in ΔE and a peak. But even after the peak, ΔE remains positive in all cases. The amplitude of the ΔE peak, as well as the break-even times, however, differ between resolutions and Reynolds numbers, reflecting differences in the accumulated particle distortion.

Specifically, for $Re = 50$, the simulation without shifting has a lower error for approximately the first 0.3 seconds. Beyond that time, shifting improves the results for the remaining simulation time. The greatest improvement from shifting can be seen at around 0.55 s for all resolutions. The behavior for $Re = 100$ follows a similar pattern, but shifting improves the error already after about 0.2 s. The maximum in ΔE also occurs earlier, at around 0.45 s. Additionally, a higher maximum is reached for all resolutions. Both observations can be explained by stronger advection leading to more particle distortion. Consequently, the shifting error amortizes earlier and quicker.

This trend remains consistent also for the two higher Reynolds numbers, with break-even times reached earlier and the peaks becoming more pronounced. However, the peaks no longer occur at the same time for all resolutions. This is due to the phenomenon of particle separation around vortex cores. At the higher two Reynolds numbers, particle separation becomes the dominant source of particle distortion, and indeed the peaks coincide with the times when particle separation first occurs. For coarser resolutions, separation occurs later. This is in contrast to the lower Reynolds numbers, where particle distortion is only caused by viscous forces and hence the peak

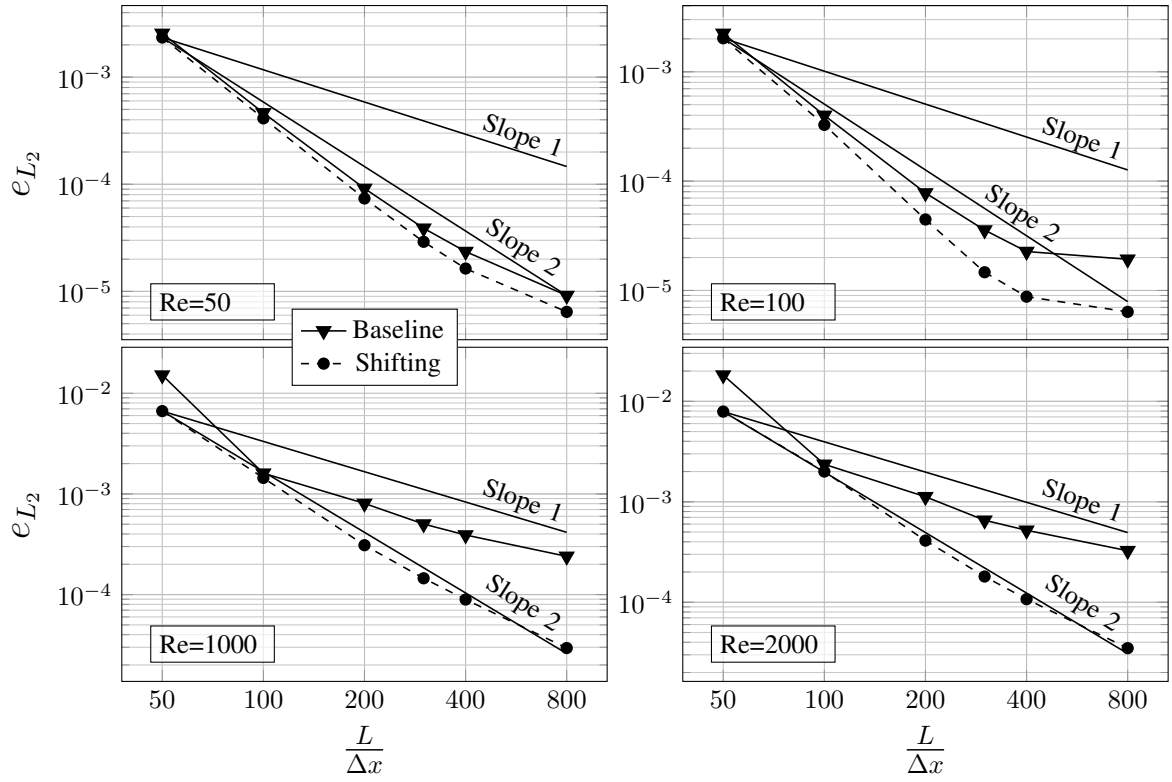

Figure 2: Convergence with and without shifting in the L_2 norm for different Reynolds numbers at time $\tilde{t} = 1.0$ for $\frac{h}{\Delta x} = 6$.

always occurs at the same time.

In order to understand the origin and scaling of the shifting error, we consider the shifting displacements as defined in Eq. (9), expressed in terms of the particle diameter Δx instead of h:

$$\tilde{\mathbf{x}}_i = K\Delta x^2 \nabla C_i \,. \tag{16}$$

The constant $K = -0.4 C_{\text{CFL}} \sigma^2$ controls the distance a particle is shifted. The accumulation of all shifting displacements up until a time t is called *global shifting distance*, and it is defined as:

$$R(t) := \sum_i \|\tilde{\mathbf{x}}_i(t)\|_2 = |K|\Delta x^2 \sum_i \|\nabla C_i(t)\|_2 \,. \tag{17}$$

Therefore, $R(t) = |K| \|\nabla C(t)\|_{L_1}$ for the two-dimensional L_1-norm defined in Eq. (13). Normalizing this by the particle number N and diameter, we obtain the quantity

$$\bar{R}_d(t) := \frac{R(t)}{N\Delta x} \,, \tag{18}$$

which describes the average global shifting distance per particle in units of particle diameters. The number of particles in d dimensions is $N = \frac{L}{\Delta x^d}$. For $L = 1$, as used in the above benchmarks, the average global shifting distance per particle can thus be expressed as:

$$\bar{R}_d(t) = \frac{\Delta x^d R(t)}{\Delta x} = \Delta x^{d-1} R(t) \,, \tag{19}$$

which for the two-dimensional case becomes $\bar{R}_2(t) = \Delta x R(t) = \Delta x |K| \|\nabla C(t)\|_{L_1}$.

This quantity can be used to compare different cases across varying particle diameters. It describes how many particle diameters a particle has been shifted on average by time t. It behaves qualitatively similar across resolutions and Reynolds numbers of the Taylor-Green test case, as shown in Fig. 4. For all Reynolds numbers and resolutions, a similar magnitude of $\bar{R}$ is observed. Thus, the magnitude of the particle displacement from shifting is comparable across cases. Another similarity among all cases is that $\bar{R}$ increases significantly in the beginning of the simulation, reaching a first local maximum around 0.1 seconds. From there, the value decreases monotonously over time for the smaller Reynolds numbers. This is in agreement with the viscous forces decaying over time and causing less particle distortion. For the larger Reynolds numbers, $\bar{R}$ remains more or less constant beyond 0.1 seconds, which is consistent with particle separation from vortex core expansion dominating the shifting.

Another difference can be observed between the lower and higher Reynolds numbers: For the former, the average shifting distance is practically independent of the resolution, whereas for the latter higher resolutions lead to less shifting. This is consistent with the results in Fig. 3, where finer resolutions benefited earlier from shifting.

Taken together, across all Reynolds numbers and resolutions, shifting adds additional error initially, but becomes amortized beyond a certain break-even point and remains ad-

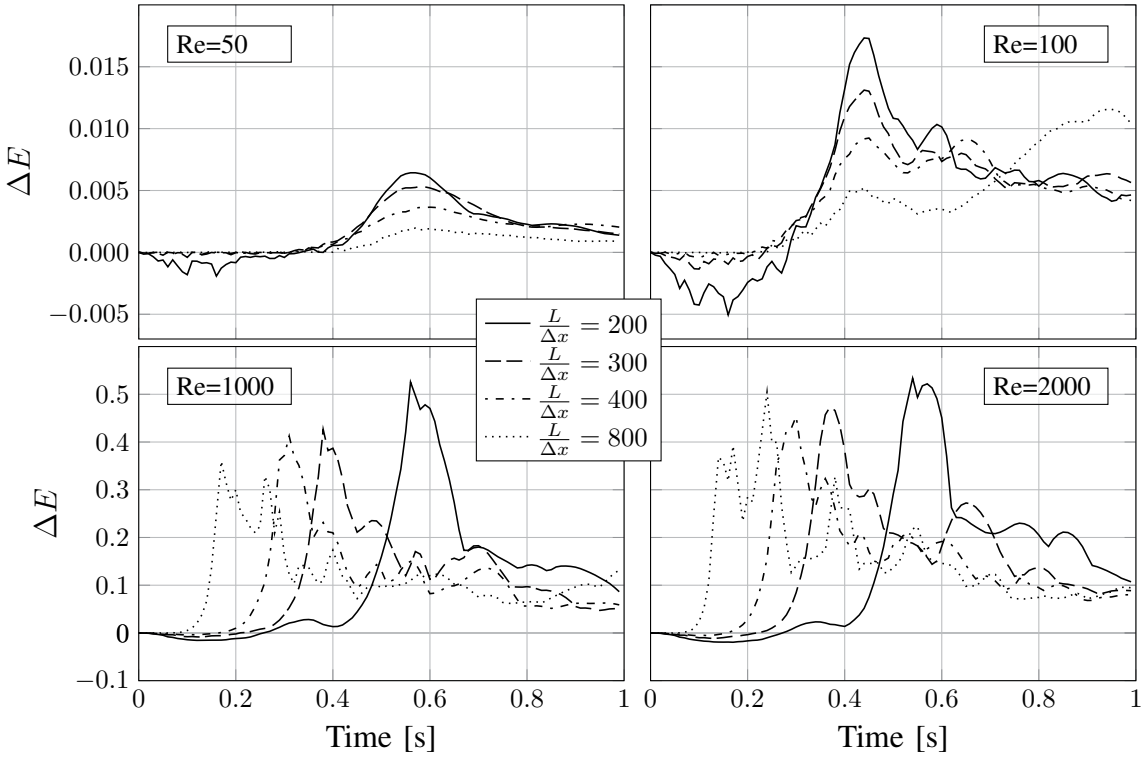

Figure 3: Time evolution of the error difference ΔE between the L_∞ errors without shifting and with shifting for different resolutions (line styles, inset legend) and Reynolds numbers (panels). Positive values indicate that shifting reduces the overall error.

vantageous thereafter. In all cases, the normalized global shifting distance per particle behaves similarly across Reynolds numbers and resolutions. The break-even times, however, depend on the resolution and the Reynolds number in a case-specific manner.

V. Break-even point for shifting

From the observations in the previous section, we design a case-independent metric for when shifting is beneficial for the accuracy of SPH. This metric should be independent of Reynolds number and resolution, hence predicting shifting break-even times across cases.

As we know from the previous section, such a metric should be proportional to time at lower Reynolds numbers, where break-even times do not depend on resolution, but capture the resolution-dependent particle expansion behavior at higher Reynolds numbers. All of these effects are captured in a problem-independent manner by the accumulated Lagrangian particle distortion. Since shifting acts on the particle distortion, a distortion metric should provide a universal and problem-independent way of quantifying the effect of shifting on convergence.

We quantify the accumulated Lagrangian particle distortion using an order parameter for the particle distribution, which is independent of the cause of distortion. In two dimensions, we start by applying Delauney triangulation to the positions of the particles. The resulting triangulation, denoted by $\mathcal{D}(t)$, contains k triangles with areas $A_k(t)$ at time t. In a perfectly uniform particle distribution, the reference area A_0 would be:

$$A_0 = \frac{\Delta x^2}{2}\,. \tag{20}$$

We hence use the maximum relative deviation of A_k from A_0 to measure Lagrangian particle distortion in the L_∞ sense:

$$\max_{k\in\mathcal{D}(t)} \frac{|A_k(t) - A_0|}{A_0}\,. \tag{21}$$

The normalization with A_0 renders the order parameter dimensionless. This provides a way of expressing the shifting error in a problem-independent way. However, simply expressing ΔE in terms of these quantities is not particularly insightful, since both the errors with and without shifting accumulate over time steps. Therefore, the trajectory of the particle distortion until the current time has to be taken into account, and not just the magnitude of the current distortion. We thus define the distortion metric

$$A_D(t) = \sum_{t'\in[0,t]} \max_{k\in\mathcal{D}(t')} \frac{|A_k(t') - A_0|}{A_0} \tag{22}$$

over simulation time steps.

Figure 5 shows the shifting error difference ΔE as a function of A_D. Consistently, simulations without shifting

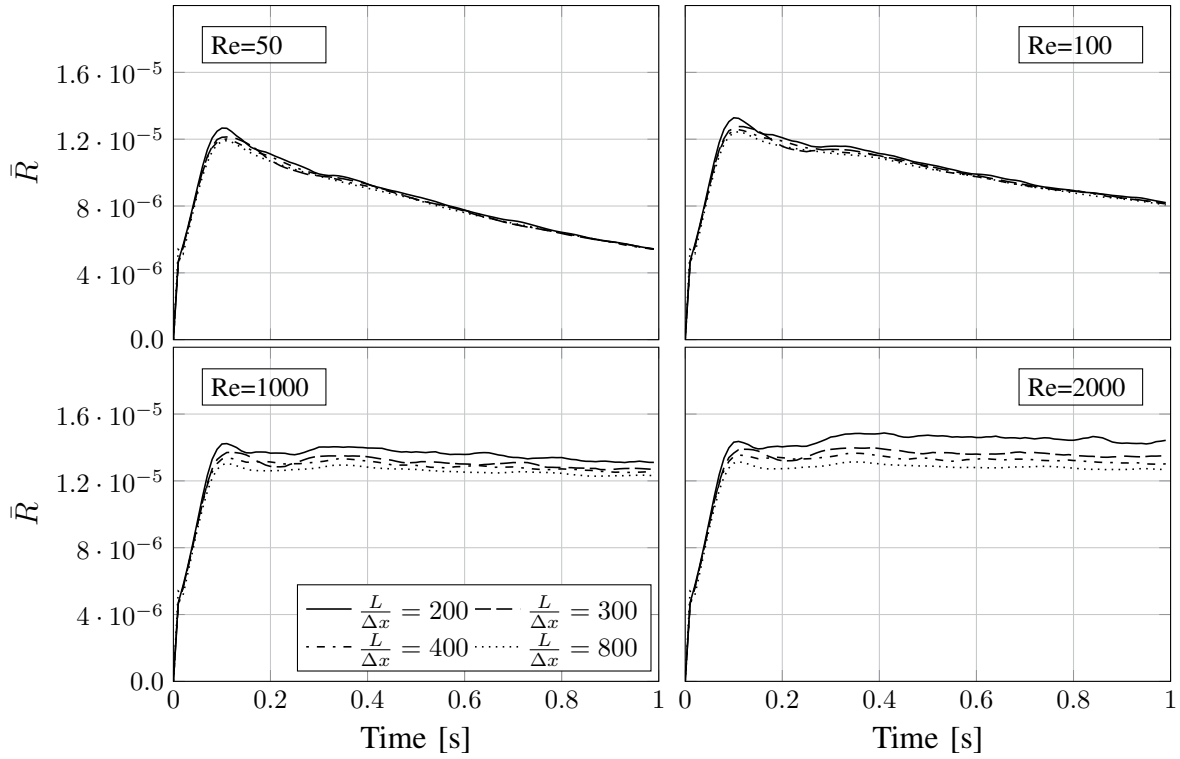

Figure 4: The average global shifting distance per particle as a function of time for the two-dimensional Taylor-Green vortex at different resolutions (line styles, inset legend) and Reynolds numbers (panels).

have lower errors for small A_D. At a certain distortion, there is a break-even point followed by a pronounced peak. Unsurprisingly, A_D reaches higher values for larger Reynolds numbers, rendering the results consistent with the previous observations. Importantly, however, both the break-even points and the locations of the peaks are similar across problem cases, and the behavior is qualitatively similar for both the lower and higher Reynolds numbers. This indicates that indeed the dimensionless accumulated particle distortion A_D can be used to determine when shifting becomes advantageous, independent of the cause of the distortion.

Table I: Break-even points in A_D for the Reynolds numbers and resolutions tested.

	$\frac{L}{\Delta x} = 200$	$\frac{L}{\Delta x} = 300$	$\frac{L}{\Delta x} = 400$	avg.
$Re = 50$	131.71	116.80	129.0	125.84
$Re = 100$	91.63	76.85	62.59	77.02
$Re = 1000$	93.03	69.42	49.71	70.72
$Re = 2000$	102.58	72.93	46.09	73.87

In order to check to what extent the universality in the behavior of the shifting error difference with A_D is also quantitative, we report the A_D values of the break-even points in Table I for the different Reynolds numbers and resolutions tested. While there is still a trend for higher resolutions to have smaller break-even A_D, especially at high Re, the overall recommendation is to use shifting when A_D exceeds a value of about 70 to 80.

VI. Conclusion

We have studied how particle shifting affects the convergence behavior of SPH. We found that shifting introduces an additional second-order error to the simulation, which can restore convergence when it becomes dominant. We examined the properties of the shifting error across different Reynolds numbers and resolutions in the two-dimensional Taylor-Green vortex case, for which an analytical solutions is available. We observed distinct behavior between smaller and larger Reynolds numbers. The disparities were attributed to two different dominant drivers of Lagrangian particle distribution distortion: viscous shear forces at lower Reynolds numbers and particle separation at higher Reynolds numbers. Regardless of the cause of particle distortion, however, the shifting error was always amortized for long-enough simulated times.

Based on our observations, we proposed a case-independent recommendation for when shifting is expected to improve the accuracy of SPH. This was based on a dimensionless metric of the accumulated Lagrangian particle distortion. When plotted in this metric, the difference in error with and without shifting was qualitatively similar across resolutions and Reynolds numbers. Finally, we found that the proposed distortion metric is approximately quantitative and can thus be used to predict when shifting should be beneficial.

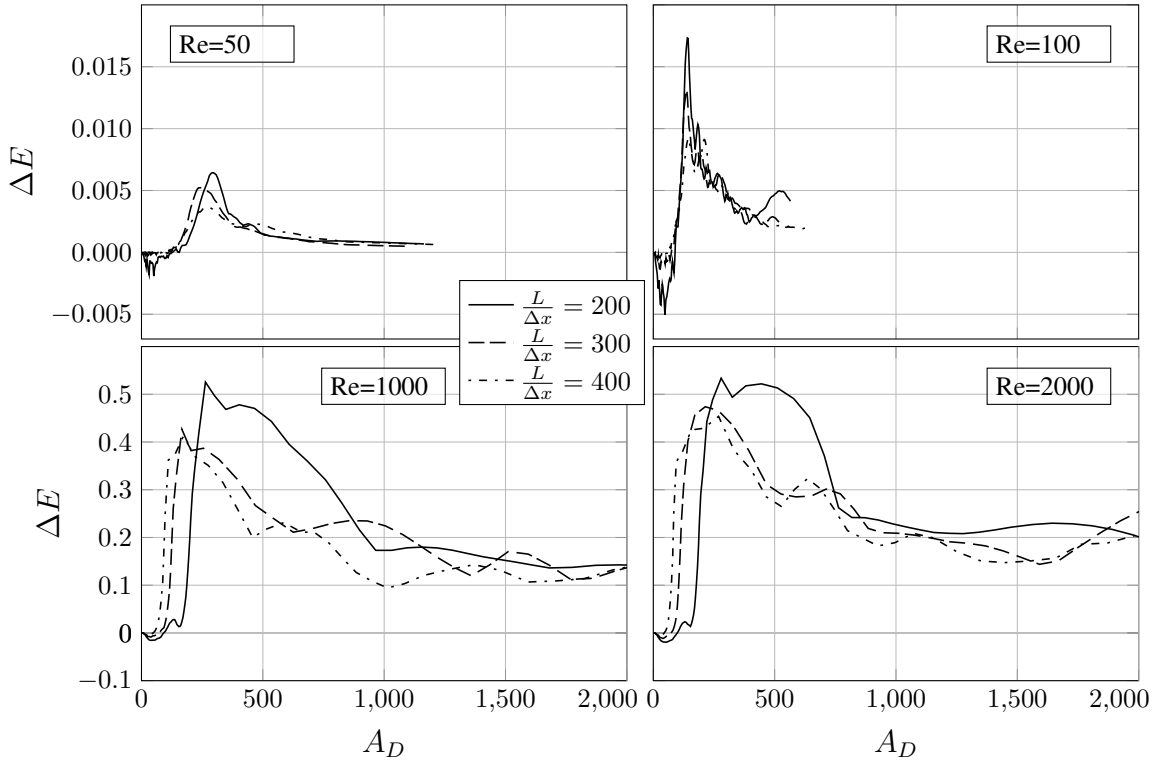

Figure 5: L_∞ shifting error difference ΔE as a function of the dimensionless Lagrangian particle distortion metric A_D for different resolutions (line styles, inset legend) and Reynolds numbers (panels).

In the future, the present results could be extended to three-dimensional problems and to free-surface flows. They could also be compared with other measures of particle distortion [1].

References

[1] M. Antuono, B. Bouscasse, A. Colagrossi, and S. Marrone. A measure of spatial disorder in particle methods. *Computer Physics Communications*, 185(10):2609–2621, 2014.

[2] S. Lind, R. Xu, P. Stansby, and B. Rogers. Incompressible smoothed particle hydrodynamics for free-surface flows: A generalised diffusion-based algorithm for stability and validations for impulsive flows and propagating waves. *Journal of Computational Physics*, 231:1499–1523, 2012.

[3] D. Molteni and A. Colagrossi. A simple procedure to improve the pressure evaluation in hydrodynmaic context using the SPH. *Computer Physics Communications*, 180:861–872, 2009.

[4] J. Morris, P. Fox, and Y. Zhu. Modeling low reynolds number incompressible flows using SPH. *Journal of Computational Physics*, 136:214–226, 1997.

[5] G. Oger, S. Marrone, D. Le Touzé, and M. De Leffe. SPH accuracy improvement through the combination of a quasi-lagrangian shifting transport velocity and consistent ALE formalisms. *Journal of Computational Physics*, 313:76–98, 2016.

[6] N. Quinlan, M. Basa, and M. Lastiwka. Truncation error in mesh-free particle methods. *International Journal for Numerical Methods in Engineering*, 66(13):2064–2085, 2006.

[7] P. Sun, A. Colagrossi, S. Marrone, M. Antuono, and A. Zhang. Multi-resolution delta-plus-SPH with tensile instability control: Towards high reynolds number flows. *Computer Physics Communications*, 224:63–80, 2018.

[8] P. Sun, A. Colagrossi, S. Marrone, and A. Zhang. The δ-plus-SPH model: Simple procedures for a further improvement of the SPH scheme. *Computer Methods in Applied Mechanics and Engineering*, 315:25–49, 2017.

[9] G. Taylor and A. Green. Mechanism of the production of small eddies from large ones. *Proceedings of the Royal Society of London. Series A, Mathematical and Physical Sciences*, 158:499–521, 1937.

[10] H. Wendland. Piecewise polynomial, positive definite and compactly supported radial functions of minimal degree. *Advances in Computational Mathematics*, 4:389–396, 1995.

[11] R. Xu, P. Stansby, and D. Laurence. Accuracy and stability in incompressible SPH (ISPH) based on the projection method and a new approach. *Journal of Computational Physics*, 228(18):6703 – 6725, 2009.